ASIA SUMMER INSTITUTE IN BEHAVIORAL ECONOMICS

National University of Singapore, July 26 to August 6, 2010

Lectures on Strategic Thinking

Vincent P. Crawford, University of Oxford

(with special thanks to Miguel Costa-Gomes and Nagore Iriberri)

This version 22 July 2010

Outline

1. Introduction: Why Study Strategic Thinking?

2. Alternative Models of Strategic Thinking

3. Keynes' Beauty Contest: Experimental Evidence from Guessing and Other Normal-Form Games

4. M. M. Kaye's Far Pavilions: Responding to Payoff Asymmetries in Outguessing Games

5. Groucho's Curse: Zero-Sum Betting and Auctions with Incomplete Information

6. Kahneman's Entry Magic: Coordination via Symmetry-Breaking

7. Bank Runs: Coordination via Assurance

8. Non-Equilibrium Econometrics: Structural Alternatives to Incomplete Models

9. Yushchenko and Lake Wobegon: Non-neutral Framing in Outguessing Games

10. Mr. Schelling Goes to Chicago: Coordination via Payoff Asymmetries and Non-neutral Framing

11. Huarangdao and D-Day: Communication of Intentions in Outguessing Games

12. Alphonse and Gaston: Communication of Intentions in Coordination Games

13. October Surprise: Communication of Private Information in Outguessing Games

1. Introduction: Why study strategic thinking?

Strategic thinking pervades human interaction.

As soon as children develop enough "theory of mind" to model other people as independent decision makers, they must be *taught* to look both ways before crossing one-way streets—suggesting that they instinctively rely on rationality assumptions to predict others' decisions.

Adult attempts to predict other people's responses to incentives are shaped by similar—though usually more subtle—rationality-based inferences.

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

The canonical model of strategic thinking is the game-theoretic notion of Nash equilibrium, defined as a combination of strategies, one for each player, such that each player's strategy maximizes his expected payoff, given the others' strategies.

(Equilibrium can be defined and applied without reference to its interpretation, but it is best thought of as an "equilibrium in beliefs," in which players who are rational in the decision-theoretic sense have beliefs about each other's strategies that are correct, given the rational strategy choices they imply.)

In games, rationality alone seldom restricts behavior enough to be useful.

Even common knowledge of rationality implies only that players' strategies are rationalizable (Bernheim 1984 and Pearce 1984), which leaves behavior completely unrestricted in many games.

Equilibrium makes more definite predictions by augmenting rationality with the "rational-expectations" assumption that players' beliefs are correct.

Because many interesting games have multiple equilibria, equilibrium is often further augmented by refinements, with the goal of deriving unique predictions.

The structure equilibrium and refinements add to beliefs often gives a precise and plausible account of strategic behavior; and the generality, simplicity, and tractability of equilibrium analysis have made it the method of choice in strategic applications.

Although equilibrium is almost always assumed in applications, it is better justified in some than others.

When players have abundant prior experience with perfectly analogous games, both theory and experiments suggest that under mild assumptions about cognition and behavior, learning has a strong tendency to converge to equilibrium.

(In reinforcement learning, for example, players need not even know that they are playing a game. My statement omits qualifications that are unimportant for normal-form games.)

If only long-run outcomes matter, and if equilibrium is unique or if there are multiple equilibria but selection does not depend on the details of learning, then such applications can safely rely entirely on equilibrium, and there is no need to study strategic thinking.

However, in many settings players' current interaction has only imperfect precedents, or none at all, so that the learning justification for equilibrium is unavailable.

Then, if assuming equilibrium is justified, it must be via strategic thinking rather than learning.

The literature on epistemic game theory provides conditions under which reasoning based on iterated knowledge of rationality and beliefs focuses rational players' beliefs on a particular equilibrium, even in their initial responses to a game.

But in many games such reasoning is complex enough to make the thinking justification for equilibrium behaviorally implausible.

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.

Assuming equilibrium is then less justified for initial responses than when learning is possible.

(Note that there is no plausible "as if" thinking justification for equilibrium: If players' models of other players do not accurately reflect their thinking, equilibrium will predict their decisions in general games accurately only by chance.)

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 long-run outcome is selected from multiple equilibria via history-dependent learning dynamics. (My lecture slides on learning give detailed examples of this.)

(Analyses of "long-run equilibria" study ergodic dynamics that rule out the history-dependence that is evident in the data by assumption, and so do not solve this problem.)

All such applications depend on reliably predicting initial responses to games.

And in games where the thinking justification for equilibrium is behaviorally implausible, this may require a non-equilibrium model of strategic thinking.

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.

(The robustness then resembles that established by a rationalizability-based analysis, but better models can add useful structure, and the results may deviate from equilibrium-based conclusions.)

It can challenge conclusions based on equilibrium or refinements in 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 from imperfect analogies, where assumptions about cognition determine which analogies between current and previous games players recognize and elucidate the structure of learning, where assumptions about cognition distinguish reinforcement from beliefs-based and more sophisticated rules.

But even those who grant the desirability of improving upon equilibrium models of initial responses may doubt its feasibility.

How can any model systematically out-predict a rational-expectations notion such as equilibrium?

And how can one identify simple models that allow such improvements among the huge number and variety of logically possible non-equilibrium models?

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

When equilibrium is assumed in such cases, or when the scope of its learning justification is overestimated, it may be because analysts hope that equilibrium will still be correct on average, or fear that without equilibrium there can be no basis for analysis.

Experiments

Yet there is now a large body of experimental research that studies strategic thinking by eliciting initial responses to games with a variety of structures, which suggests that in many applications neither the hope nor the fear just mentioned is justified.

(Although most empirical work in economics has used econometrics to analyze field data, experiments have played a leading role in empirical work on behavior in games.

Such behavior is notoriously sensitive to the details of the environment, so that strategic models carry a heavy informational burden, which is often compounded in the field by an inability to observe all relevant variables.

Experimental methods allow a control that often gives experiments a decisive advantage in identifying the relationship between behavior and the environment.

I discuss clear evidence from field data below whenever possible, but the bulk of my discussion focuses on experimental data.)

The experimental research shows, with progressively increasing clarity, that people's initial responses to novel games often deviate systematically from equilibrium.

Importantly, the results also show that the deviations have a large structural component that can be modeled in a simple way.

This is possible because subjects' thinking systematically avoids the fixed-point or indefinitely iterated dominance reasoning that equilibrium sometimes requires.

In Selten's words:

> "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 stationary setting, learning can't make real people converge to steady states that *we* would need fixed-point reasoning to characterize, just that such reasoning doesn't directly describe people's thinking.)

If the structural component of subjects' thinking avoids fixed-point or indefinitely iterated dominance reasoning, then what does it consist of?

Much of the experimental evidence suggests that subjects tend instead to follow rules of thumb that anchor beliefs in an instinctive reaction to the game and then adjust their beliefs via a small number of iterated best responses.

These rules of thumb—called "types" in this context by historical accident; no relation to private-information variables!—are cognitively simple, and have strong intuitive appeal.

Although subjects' thinking is typically heterogeneous, their types are drawn from a stable population distribution concentrated on one to three best-response iterations.

The results identify a class of "level-*k*" or "cognitive hierarchy" ("CH") models that share the generality and much of the simplicity and tractability of equilibrium analysis, but which in many settings can systematically out-predict equilibrium.

Although level-*k*/CH models are alternatives to equilibrium analysis, they generalize equilibrium rather than replacing it.

Level-*k* types are rational in the sense of best-responding to some beliefs; they depart from equilibrium only in that their beliefs are derived from simplified, non-equilibrium models of other players.

Level-*k* type *k* (though not its CH counterpart beyond $k = 1$) respects *k*-rationalizability, the condition that corresponds in two-person games to the result of *k* rounds or iterated deletion of dominated strategies.

In sufficiently simple games, the low-level types that describe most subjects' behavior mimic equilibrium strategy choices, even though they deviate from equilibrium thinking.

But in more complex games, some or all such types may deviate systematically from equilibrium choices.

Importantly, a level-*k*/CH model not only predicts that such deviations will sometimes occur.

Given estimates of the population type frequencies, it also identifies which settings are likely to evoke deviations; what forms they will take; and with what frequencies.

Although level-*k*/CH models predict a sizeable fraction of deviations from equilibrium in many settings, by no means do they predict all deviations in all interesting settings.

They seem to predict half or more of the deviations in a majority of normal-form settings.

This should not be disappointing: It is encouraging that such simple and tractable models can predict half or more of anything as elusive as deviations from equilibrium.

Moreover, the experimental results also suggest that the strategic thinking-related deviations that level-*k*/CH models do *not* predict have little discernable structure.

Thus, level-*k*/CH models generalize equilibrium analysis in a way that is likely to be useful in settings where deviations from equilibrium are important, while ignoring little that cannot reasonably be modeled as errors.

Viewed this way, there is a strong case for adding level-*k*/CH models to our toolkit.

Folk GameTheory

Throughout the lectures I will link the experimental evidence to clear informal descriptions of strategic thinking that I will call folk game theory, which vividly illustrate the need for non-equilibrium models of strategic thinking, the issues that a successful model must address, and the range of potential applications.

I use the term folk game theory to suggest an analogy with folk physics, untrained people's intuitive beliefs about the laws of physics.

Why study folk instead of "real" game theory?

Folk physics yields insight into human cognition, but only limited insight into real physics.

Just as folk physics is an imperfect reflection of real physics, folk game theory is an imperfect reflection of traditional game theory.

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

The lessons about strategic thinking from folk game theory reinforce the lessons from experiments that study strategic thinking in more conventional ways.

This close correspondence is powerful further evidence for level-*k*/CH models.

The instinctive reactions people use to anchor their beliefs follow different principles in different settings, as I will illustrate:

● uniform randomness (reflecting either the principle of insufficient reason or payoff sampling uninformed by strategic structure) in most normal-form games

● attraction to salient labels or payoffs in context where these are important

● truthfulness in games where players can communicate

The finite iteration of best responses by which people adjust their beliefs is common to all settings, although the number of iterations may vary across people and/or settings.

These lessons are representative of folk game theory.

One can also find clear informal descriptions of strategic thinking that reflect one or two steps of iterated (strict or weak) dominance in the normal form, or of one or two steps of iterated (weak) dominance via forward or backward induction in the extensive form.

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

It is more difficult (impossible? counterexamples?) to find descriptions that illustrate the fixed-point reasoning that underlies equilibrium in games without dominance.

Organization

The rest of these lectures follow the outline given at the start, interweaving experimental evidence with strategic and economic applications, and ordering topics by strategic rather than economic issues.

(Level-*k*/CH aficionados will see that the topics are grouped by the principles used to specify the anchoring *L0* type—uniform randomness for topics 3-8, attraction to salience for topics 9-10, or truthfulness for topics 11-13. The topics are also ordered to facilitate teaching non-aficionados how the models work and their economic implications.)

I assume throughout that players have accurate models of the games and that, except for errors, their strategy choices are rational responses to some beliefs about others' choices.

I also focus on normal-form games, including extensive-form games only to study communication.

Finally, I focus mainly on two-person games.

2. Alternative Models of Strategic Thinking

Until recently, the choices for modeling non-equilibrium initial responses to games were quite limited, but now there are several alternatives.

This section sets the stage by reviewing them, their cognitive requirements, and how they are implemented in applications.

The leading models of strategic thinking all allow players' strategy choices to be in equilibrium, but do not assume equilibrium in all games.

The models to be discussed include:

- equilibrium plus noise
- finitely iterated dominance and *k*-rationalizability (Bernheim 1984 and Pearce 1984)
- quantal response equilibrium ("QRE"; McKelvey and Palfrey 1995)
- level-*k* models (Nagel 1995, Stahl and Wilson 1994, 1995, Costa-Gomes, Crawford, and Broseta 2001, Costa-Gomes and Crawford 2006)
- cognitive hierarchy models ("CH"; Camerer, Ho, and Chong 2004)
- noisy introspection models ("NI"; Goeree and Holt 2004).

A. Equilibrium plus noise

Any notion that is to be taken to data must allow for errors in some way.

Equilibrium plus noise adds to equilibrium predictions mean-zero noise with a payoff-sensitive error distribution (usually logit) and an estimated precision parameter.

Although a player's error distribution is sensitive to the payoff costs of errors, those costs are evaluated assuming (unlike in most other models discussed here) that other players play their equilibrium strategies without errors.

In games with multiple equilibria, equilibrium plus noise is "incomplete" in that it does not specify a unique (even if probabilistic) prediction conditional on the values of its behavioral parameters (in this case, the precision parameter or parameters).

To put equilibrium plus noise on an equal footing with the other models to be considered here, which with one exception are usually complete, one can add a coordination refinement such as Harsanyi and Selten's (1987) risk- or payoff-dominance, or possibly maximin behavior, which in symmetric coordination games functions like a refinement.

In many games equilibrium plus noise describes observed behavior well.

But even in games with unique equilibria, equilibrium plus noise sometimes misses systematic patterns in deviations from equilibrium.

Such deviations are often sensitive to out-of-equilibrium payoffs, in ways that equilibrium plus noise can account for only via a payoff-sensitive error distribution.

But this channel is limited by its assumption that a player's deviation costs are evaluated assuming that other players play their equilibrium strategies without errors.

QRE and level-*k*/CH and NI models each seek to track the sensitivity of deviations to a richer set of a player's payoffs, in different ways.

B. Finitely iterated (strict) dominance and *k*-rationalizability

Finitely iterated strict dominance and *k*-rationalizability are incomplete models, yielding set-valued restrictions on individual players' strategies that capture the implications of common knowledge (sometimes now called common belief) or iterated knowledge of players' rationality without further restrictions on their beliefs.

k-rationalizability reflects the implications of *k* levels of mutual knowledge of rationality.

The more familiar notion of rationalizability is equivalent to *k*-rationalizability for all *k*.

(Equilibrium and QRE, by contrast, restrict the *relationship* among players' strategies.)

(Finitely iterated strict dominance and *k*-rationalizability are equivalent in the two-person games I mostly focus on, and their differences in *n*–person games are unimportant here.

A 1-rationalizable strategy (the sets R1 on the next slide) is one for which there is a profile of others' strategies that makes it a best response; a 2-rationalizable strategy (the sets R2 on the next slide) is one for which there exists a profile of others' 1-rationalizable strategies that make it a best response; and so on.

		R1,R2 L	R1,R2,R3,R4 C	 R
R1,R2,R3	T	0 / 7	5 / 0	3 / 0
R1,R2,R3,R4	M	0 / 5	2 / 2	0 / 5
R1	B	7 / 0	5 / 0	3 / 7

Dominance-solvable game

		Rk for all k L	Rk for all k C	Rk for all k R
Rk for all k	T	0 / 7	5 / 0	7 / 0
Rk for all k	M	0 / 5	2 / 2	0 / 5
Rk for all k	B	7 / 0	5 / 0	0 / 7

Unique equilibrium without dominance

Each game has a unique equilibrium (M, C). In the first game M and C are the only rationalizable strategies; in the second game all strategies are rationalizable.

Equilibrium reflects the implications of common knowledge of rationality *plus* common beliefs: Any equilibrium strategy is *k*-rationalizable for all *k*, but not all combinations of rationalizable strategies are in equilibrium.

In games that are dominance-solvable in *k* rounds, *k*-rationalizability implies that players have the same beliefs—with a qualification for mixed-strategy equilibrium (that a player may know more than others about his own pure strategy) that is unimportant here—so any combination of *k*-rationalizable strategies is in equilibrium as in the first game above.

In other games, *k*-rationalizability and rationalizability allow deviations from equilibrium, as in the second game above, where there is a "tower" of beliefs, consistent with common knowledge of rationality, to support any combination of strategies.

(But except for the equilibrium beliefs (M, C), the tower beliefs differ across players.)

As we will see, finitely iterated dominance and *k*-rationalizability are often consistent with systematic patterns in subjects' deviations from equilibrium.

Importantly, the beliefs that support many rationalizable outcomes are behaviorally implausible in that (as in the tower above) they rest on rationality-based inferences at unrealistically high levels and/or they vary from level to level in an implausible way.

A possible remedy is to combine rationality with empirically based restrictions on beliefs, as in the level-*k*/CH models discussed below.

C. Quantal response equilibrium ("QRE") and logit QRE ("LQRE")

QRE seeks to capture the payoff-sensitivity of deviations from equilibrium that equilibrium plus noise sometimes misses by allowing a player's deviations to respond to a richer set of out-of-equilibrium payoffs with a more nuanced model of their likelihoods.

In a QRE, as in equilibrium plus noise, players' strategy choices are noisy, with the probability density of each choice increasing in its expected payoff.

A QRE is a fixed point in the space of choice probability distributions, with each player's choice distribution a noisy best response to others' distributions.

Thus, unlike in equilibrium plus noise models—or in level-*k* or, usually, CH models—the payoff costs of deviations are evaluated taking the noisiness of others' choices into account.

A QRE model is closed by specifying a response distribution, which is logit in almost all applications. E.g. for a matrix game, "logit QRE" or "LQRE" implies:

$$P_{ij} = \frac{exp(\lambda EU_{ij}(P_{-i}))}{\sum_k exp(\lambda EU_{ik}(P_{-i}))}$$

where P_{ij} is the probability of player i choosing strategy j, $EU_{ij}(P_{-i})$) is the expected utility to player i of choosing strategy j given that other players are playing according to the probability distribution P_{-i}, and λ is the logit precision parameter.

As $\lambda \rightarrow 0$, players play each strategy with equal probability, and as $\lambda \rightarrow \infty$, players' strategies approach a particular equilibrium (often with clear implications for equilibrium selection).

The resulting logit QRE or LQRE for short implies error distributions that respond to out-of-equilibrium payoffs, often in plausible ways.

In applications LQRE’s precision is estimated econometrically or calibrated from previous analyses.

Like equilibrium plus noise, QRE is a general model of strategic behavior, applicable to any game, with a small number of behavioral parameters.

With estimated precision, LQRE’s sensitivity to a richer set of out-of-equilibrium payoffs and its more nuanced model of their likelihoods often allows it to fit subjects’ initial responses better than an equilibrium plus noise model.

But in some settings LQRE fits worse than equilibrium, even making systematic qualitative errors.

A LQRE player faces the greatest cognitive difficulty among players in the models considered here.

He must both (noisily) best respond to a probability distribution of other players' responses, and find a generalized equilibrium that is a fixed point in a very large space of possible response distributions.

If equilibrium reasoning is cognitively taxing enough to make equilibrium behaviorally implausible in a game, then LQRE reasoning is doubly taxing.

From the analyst's point of view, the mathematical complexity of LQRE means that it must usually be solved for computationally and is not usually easily adapted to theoretical analysis.

By contrast with level-*k* and CH models, QRE's structure is not directly grounded in experimental evidence.

Moreover, Haile, Hortaçsu, and Kosenok (2008) have shown that QRE's distributional assumptions are crucial in that without such assumptions a distribution can be constructed to make QRE exactly fit any given dataset with one observation per game-player pair.

The associated sensitivity of QRE's mean predictions to distributional assumptions stems from the fact that—unlike in quantal response models of individual decisions or (aside from NI) other models of strategic thinking—QRE's choice probabilities respond to players' expected payoffs evaluated taking the noisiness of others' decisions into account.

As Goeree, Holt, and Palfrey (2008) note, QRE does have some distribution-free testable implications within certain games and across game.

Yet in the vast majority of applications, which involve games and datasets with one observation per game-player pair, QRE's power to describe deviations from equilibrium rests on distributional assumptions which there is little theory to guide and which Haile, Hortaçsu, and Kosenok's critique shows are untestable.

The almost universal use of the logit in QRE analyses has been guided more by fit and custom than by evidence.

D. Level-*k* models

A different class of models treat deviations from equilibrium as an integral part of the structure rather than as errors or responses to errors.

Although the number of possible non-equilibrium structural models seems daunting, both experimental evidence and folk game theory are broadly consistent with, and sometimes directly suggest, a class of models called level-*k* models.

In a level-*k* model players follow rules of thumb (called "types") that anchor beliefs in an instinctive reaction to the game and then adjust them via iterated best responses.

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

(The fact that *Lk* is assumed to best respond to *Lk*-*1* rather than a distribution of lower level types is the main difference between level-*k* and CH models, and I will use the term "level-*k*" as shorthand for this distinction.)

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

Lk's only departure from equilibrium is in replacing its perfect model of others' decisions with a simplified non-equilibrium model that avoids the complexity of equilibrium.

L1 and higher types make undominated decisions, and a level-*k* *Lk* (though not a CH *Lk* beyond $k = 1$) respects *k* rounds of iterated dominance and *k*–rationalizability.

On the other hand, in games that are not dominance-solvable, *Lk*'s choices can oscillate forever as *k* increases, with no tendency toward increasing rationality (unlike a CH *Lk*).

(Such comparisons may have limited relevance, however, because empirically *k* seldom exceeds 3.)

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

Unlike in an LQRE, however, an *Lk* type does not respond to the noisiness of others' choices.

Instead the probability density of a type's choice is increasing in its expected payoff, evaluated using the type's model of others' decisions: *Lk* makes errors whose distribution is sensitive to the payoff costs of deviations, evaluated assuming other players are *Lk-1*.

Despite types' failure to respond to the noisiness of others' choices, the deterministic structure of a level-*k* model captures the sensitivity of players' deviations from equilibrium to out-of-equilibrium payoffs as illustrated below.

Level-*k* models allow behavior to be heterogeneous, but assume that each player follows a rule drawn from a common distribution over *Lk* types.

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

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

Given these empirical regularities, a level-*k* model limits the radical variation of beliefs from level to level that rationalizability allows, without relying on the behaviorally implausible cross-player interactions that drive epistemic justifications of equilibrium.

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 the instinctive reactions that define it may follow one of several principles, such as uniform randomness, attraction to salience, or truthfulness.

There is an emerging consensus about how to specify *L0* 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.

Because *Lk* complies with *k* rounds of iterated dominance and *k*–rationalizability, 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 may deviate systematically from equilibrium in 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, without assuming that players respond to the noisiness of others' choices (and so avoiding Haile, Hortaçsu, and Kosenok's critique).

Like LQRE and CH, level-*k* models often fit initial responses better than equilibrium plus noise.

Level-*k* (and CH) models are cognitively much simpler than QRE models, both for the players and for the analyst.

Instead they have a simple recursive structure, which avoids the common criticism of LQRE that finding a fixed point in the space of distributions is too taxing for a realistic model of strategic thinking.

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

Thus, unlike the other models discussed here except for equilibrium plus noise, level-*k* models' predictions do not depend on distributional assumptions or estimated precisions.

The fact that level-*k* models can provide structural, distribution-free explanations of deviations from equilibrium is an important advantage over the alternatives.

E. Cognitive hierarchy (“CH”) models

CH models are close relatives of level-*k* models in which *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 recalling that people often use rules of thumb that are simpler than econometric models, which specification better describes people’s strategic thinking remains an empirical question (on which the jury is still out).

(The fact that *Lk* is assumed to best respond to a distribution of lower level types rather than *Lk-1* alone is the main difference between CH and level-*k* models, and I will use the term “CH” as shorthand for this distinction.)

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, a CH *Lk* might not comply with *k* rounds of iterated dominance and *k*–rationalizability.

A CH *Lk* type (unlike a level-k *Lk* type) becomes progressively more rational as *k* increases, in that as the population frequency of types with the same or higher *k* declines its beliefs converge to correct beliefs.

The increasing rationality of CH *Lk* types is important in some settings, as illustrated below; but because empirically *k* seldom goes above 3, I view the *k*-rationalizability of level-k types as the more important notion of strategic rationality in practice.

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 necessarily has positive frequency in nontrivial estimated Poisson distributions, 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.

Like a level-*k* model, given the Poisson distribution a cognitive hierarchy model makes point or mean predictions that do not depend on its estimated precision.

But unlike a level-*k* model, and to some extent like QRE, the form of the distribution influences the model's point predictions.

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 (typically estimated with unconstrained distribution); but in others the Poisson constraint seems overly restrictive.

F. Noisy Introspection "NI" models

Although LQRE has so far been the most popular model of initial responses, not all researchers consider it suitable for that purpose.

McKelvey and Palfrey (1995) suggest using LQRE for both initial responses and limiting outcomes, in the latter case with precision increasing over time as a reduced-form model of learning.

But Goeree and Holt (2004) *GEB* suggest using LQRE for limiting outcomes, instead proposing an NI model for initial responses.

NI relaxes LQRE's equilibrium assumption while maintaining its assumption that players best respond to a probability distribution of others' responses:

Players form beliefs by iterating best responses roughly as in a level-*k* model, but with higher-order beliefs reflecting increasing amounts of noise.

For a given noise distribution, NI makes probabilistic predictions that depend on how fast the noise grows:

In the extreme case in which the noise does not grow with the number of iterations, NI mimics LQRE.

Other extreme cases of NI mimic level-*k* types:

If the noise jumps immediately to ∞, NI beliefs are *L0*.

If the noise is zero for one iteration and then jumps immediately to ∞, NI beliefs are *L1*; and so on.

Thus, in a sense, NI nests level-*k* and LQRE (see also Rogers, Camerer, and Palfrey 2009).

(Although Goeree and Holt motivate NI as a kind of noisy rationalizability, because it builds on iterated best responses it is more akin to level-*k* and CH models.)

In applications GH assume that the noisiness of higher-order beliefs grows geometrically with the number of iterations, which yields beliefs similar but by no means identical to *Lk*'s; slower noise growth is like a higher *k*.

The resulting NI model is more flexible than LQRE, and cognitively less taxing because it requires no fixed-point reasoning.

But it is more taxing than a level-*k* or CH model because players' choices are indefinitely iterated best responses to noisy higher-order beliefs (although for computational purposes Goeree and Holt truncate the iteration to ten rounds).

By contrast with level-*k* and CH models, NI's structure, like LQRE's, is not directly grounded in experimental evidence.

3. Keynes' Beauty Contest: Experimental Evidence from Guessing and Other Normal-Form Games

"...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*

"…imagine you are partners in a private business with a man named Mr. Market. Each day, he comes to your office or home and offers to buy your interest in the company or sell you his [the choice is yours]. The catch is, Mr. Market is an emotional wreck. At times, he suffers from excessive highs and at others, suicidal lows. When he is on one of his manic highs, his offering price for the business is high as well…. His outlook for the company is wonderful, so he is only willing to sell you his stake in the company at a premium. At other times, his mood goes south and all he sees is a dismal future for the company. In fact… he is willing to sell you his part of the company for far less than it is worth. All the while, the underlying value of the company may not have changed - just Mr. Market's mood."

—Warren Buffett's intellectual hero Benjamin Graham (of Graham and Dodd's *Security Analysis*), in Graham's *The Intelligent Investor*

The Keynes and Graham quotations evoke simultaneous-move *n*-person guessing or perhaps "outguessing" games, possibly with multiple equilibria.

Like the folk game theory quotations below, they concern games played without clear precedents.

The key issue is anticipating others' strategic responses, for Keynes to a "landscape" of personal judgments about prettiness, which is otherwise payoff-irrelevant; and for Graham to the psychology of a representative uninformed investor's reaction to news.

Equilibrium is not very helpful in anticipating others' responses in such settings.

Instead the quotations explicitly suggest thought processes in which players anchor beliefs in a model of others' instinctive reactions and then iterate best responses a finite number of times, processes whose heterogeneity and finiteness closely resemble a level-*k* or cognitive hierarchy model.

Keynes' "fourth, fifth and higher degrees" is somewhat more than the evidence suggests is realistic, but it may be only a coy reference to himself.

There is now a large body of experimental research that studies strategic thinking by eliciting initial responses to games with a variety of structures.

The most important studies of normal-form complete-information games with neutral framing include Stahl and Wilson (1994, 1995); Nagel (1995); Ho, Camerer, and Weigelt (1998); Costa-Gomes, Crawford, and Broseta (2001); Bosch-Domènech et al. (2002); Camerer, Ho, and Chong (2004); Chong, Camerer, and Ho (2005); Costa-Gomes and Crawford (2006); and Costa-Gomes and Weizsäcker (2008).

I first discuss Nagel's (1995); Ho, Camerer, and Weigelt's (1998); and Bosch-Domènech et al.'s (2002) analyses of *n*-person guessing games directly inspired by Keynes' (1936, Chapter 12) beauty contest analogy, which give a simple introduction to the evidence.

I then discuss Costa-Gomes and Crawford's (2006) analysis of two-person guessing games, whose design is more powerful and comes closest to letting the data reveal subjects' thinking directly, without an econometric "middleman".

Costa-Gomes and Crawford's (2006) conclusions are consistent with and representative of the conclusions of other carefully done studies of initial responses to normal-form games with neutral framing, just more precise. With adjustments described below, their conclusions are also consistent with those of studies of the other kinds of games.

A. Nagel's (1995); Ho, Camerer, and Weigelt's (1998); and Bosch-Domènech et al.'s (2002) experiments

In Nagel's and Ho, Camerer, and Weigelt's *n*-person guessing games, *n* subjects (*n* = 15-18 in Nagel, *n* = 3 or 7 in Ho, Camerer, and Weigelt) made simultaneous guesses between lower and upper limits (0 and 100 in Nagel, 0 and 100 or 100 and 200 in HCW).

In Bosch-Domènech et al. (2002) essentially the same games were played in the field, by more than 7500 volunteers recruited from subscribers of the newspapers *Financial Times*, *Spektrum der Wissenchaft*, or *Expansión*.

In each case the subject who guessed closest to a target (*p* = 1/2, 2/3, or 4/3 in Nagel; *p* = 0.7, 0.9, 1.1, or 1.3 in Ho, Camerer, and Weigelt; and *p* = 2/3 in Bosch-Domènech et al.) times the group average guess won a prize.

There were several treatments, each with identical targets and limits for all players and games. The structures were publicly announced, to justify comparing the results with predictions based on complete information.

For definiteness, consider Nagel's leading treatment:

- 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 are one of the few people in the world who 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?

Although Nagel's and Ho, Camerer, and Weigelt's subjects each played a game repeatedly, their first-round guesses can be viewed as initial responses if they treated their own influences on future guesses as negligible, which is plausible for all but Ho, Camerer, and Weigelt's three-subject groups.

Bosch-Domènech et al.'s subjects played only once.

The results vividly illustrate the failure of equilibrium as a descriptive model of initial responses, and the heterogeneity and discreteness of strategic thinking.

Nagel's subjects never made equilibrium guesses initially; Ho, Camerer, and Weigelt's rarely did so, and Bosch-Domènech et al.'s (who had much more time to reflect, and who could consult with others) fairly rarely did so.

In each case most subjects' initial guesses respected from 0 to 3 rounds of iterated dominance, in games where 3 to an infinite number are needed to reach equilibrium.

Here we reproduce part of Nagel's Figure 1 and Bosch-Domènech et al.'s Figure 1, which illustrate these points most clearly.

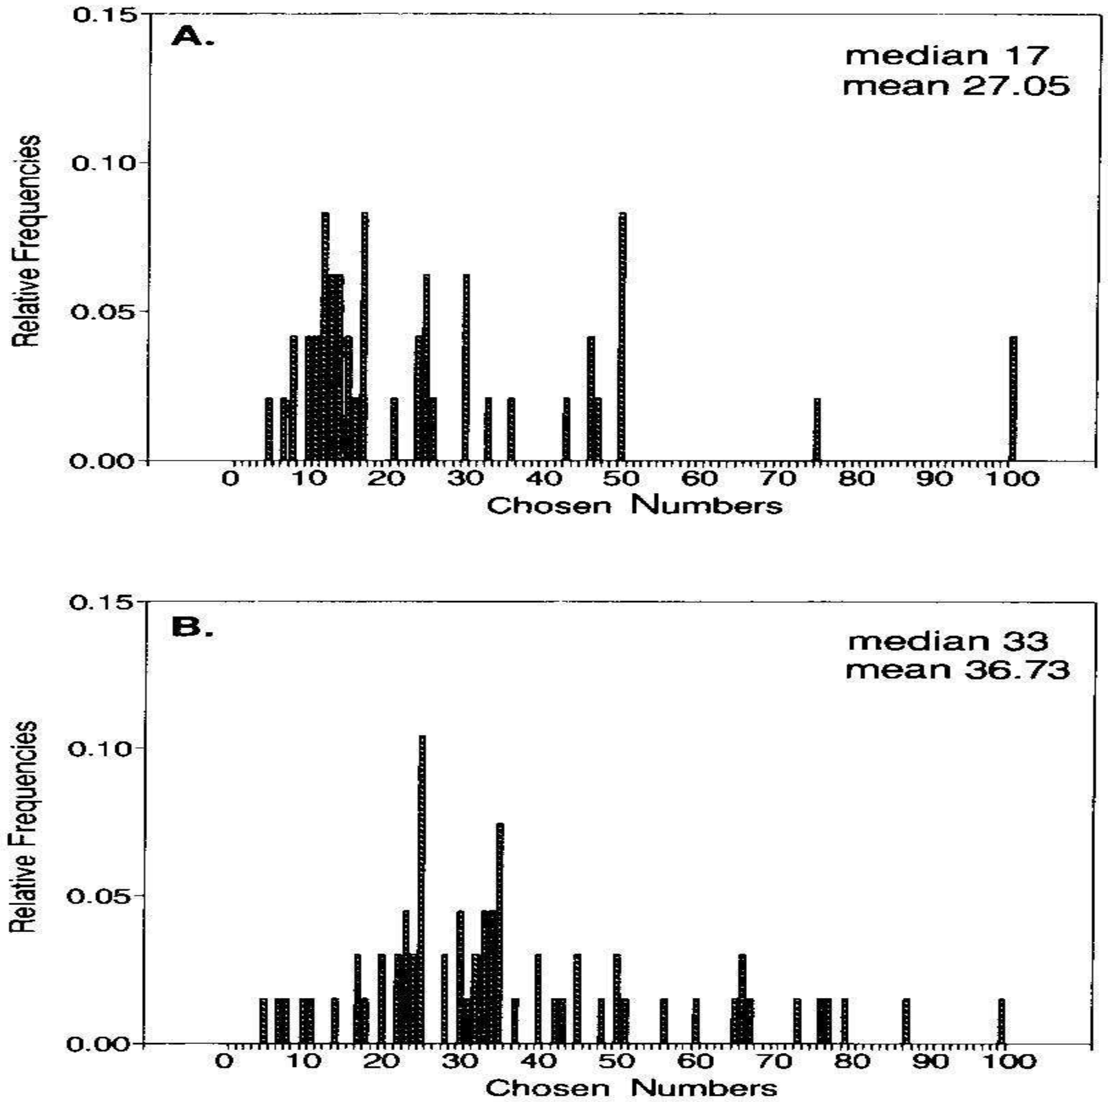

Part of Nagel's Figure 1: top of figure *p* = 1/2, bottom of figure *p* = 2/3.

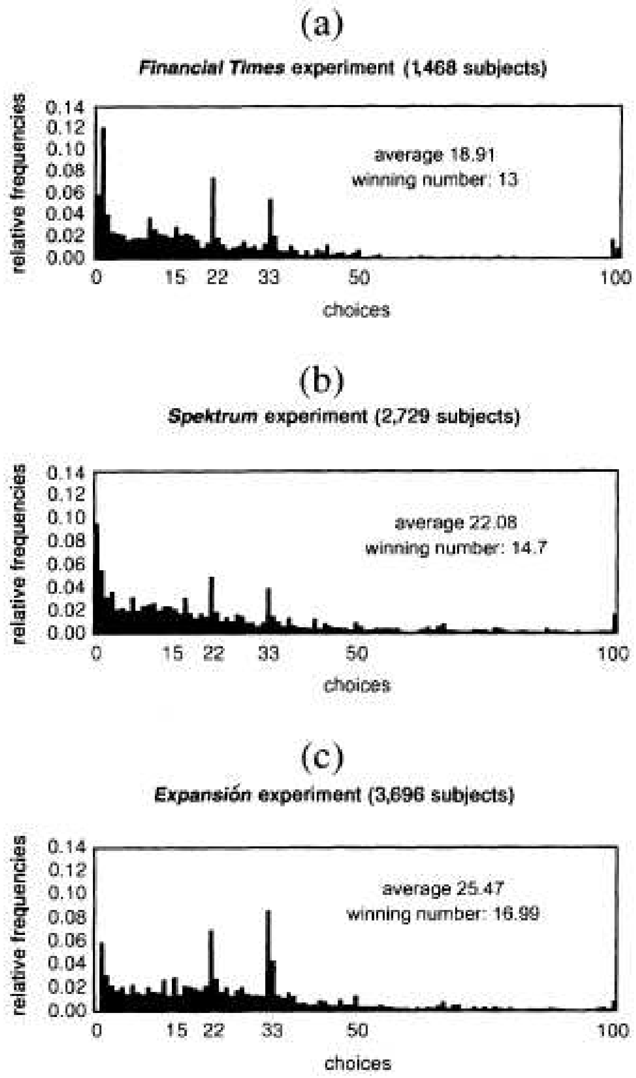

FIGURE 1. RELATIVE FREQUENCIES OF CHOICES IN THREE NEWSPAPER EXPERIMENTS

Figure 1. Bosch-Domènech et al.'s (2002) Figure 1

These data resemble neither equilibrium plus noise nor "equilibrium taking noise into account" as in QRE (for any reasonable error distribution, though by Haile, Hortaçsu, and Kosenok's (2008) result we could make the data an exact QRE for an unreasonable one).

They do suggest that subjects' deviations from equilibrium have a coherent structure.

In each case the distributions of guesses have spikes that track $50p^k$ for $k = 1, 2, 3$ across the targets p in various treatments, thus respecting 0 to 3 rounds of iterated dominance.

Like the spectrograph peaks that foreshadow the existence of chemical elements, these spikes are evidence of a partly deterministic structure, one that is discrete and individually heterogeneous.

It is already clear that no model that imposes homogeneity of strategic thinking (as all on our list but level-*k* and CH do) will do justice to subjects' behavior.

(Allowing heterogeneity of thinking is essential for the explanations proposed below of Kahneman's Entry Magic, Yushchenko, Lake Wobegon, and Huarongdao.)

Also, subjects do not respect indefinitely iterated dominance or indefinitely iterated best responses; instead their decisions respect *k*-rationalizability for at most small values of *k*.

But what about the spikes, whose consistency is the most remarkable part of the results?

Many theorists have interpreted Nagel's results as evidence that subjects explicitly performed finitely iterated dominance, the way we teach students to solve such games.

In this interpretation, which I will call *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 (completing *k*-rationalizability by adding a specific selection): thus in Nagel's games *Dk* guesses $([0+100p^k]/2)p \equiv 50p^k$.

But there is another interpretation of the spikes, which has the same implications for choice behavior in Nagel's game but which can differ in important ways in other settings, and which I shall argue is very likely the correct interpretation.

In this interpretation, called "level-*k*" or "*Lk*," a player starts with a uniform prior *L0* over the strategy space and then iterates best responses *k* times, with $k = 1, 2,$ or 3: thus in Nagel's games *Lk+1* guesses $[(0+100)/2]p^{k+1}$, which equals $([0+100p^k]/2)p \equiv 50p^k$.

Note that it is *Lk*+1 that is *Dk*'s cousin, not *Lk*. The difference in indices is only a quirk of notation, without further significance.

Both *Lk*+1 and *Dk* yield *k*-rationalizable strategies, though not always the same ones in other games. In games without dominance *Dk*, $k = 1,2,\ldots$ coincides with *L1*.

Aside on specifying *L0*

A uniform prior *L0* follows an emerging consensus in the level-*k*/CH literature for normal-form games with neutral framing. Why is this reasonable?

First, it can be viewed literally, as *L1* having a uniform prior about his partners' choices.

But it may be more plausible to take it as an analysis ignoring strategic considerations, and therefore invoking the principle of insufficient reason regarding partners' choices.

It can also be viewed as a plausible approximation of the outcome of a player's randomly sampling his payoffs for various strategies, unstratified by partners' choices.

(Later I discuss experiments on normal-form games with non-neutral framing or incomplete information, and on extensive-form games with preplay communication, in which the level-*k*/CH literature has used different, also natural specifications of *L0*.)

Another issue arises in specifying a level-*k* or CH model, to which I shall return later.

In these lectures I focus mainly on two-person games, but in *n*-person games like those discussed here it matters whether *L0* is independent across players or correlated.

The limited evidence that is available (Ho, Camerer, and Weigelt 1998 and Costa-Gomes, Crawford, and Iriberri 2009) suggests that most people have highly correlated models of others.

In analyzing Nagel's data, in keeping with the available evidence and the literature, I simply take *L0* to directly model the distribution of all others' average guess.

End of aside

But the complete lack of separation of *Dk*'s and *Lk*+1's guesses in Nagle's design shows that the inference that subjects performed finitely iterated dominance is premature.

In other experiments, including some of Ho, Camerer, and Weigelt's and Costa-Gomes, Crawford, and Broseta's, *Dk*'s and *Lk*+1's guesses are weakly separated, and the results are inconclusive on this point.

But in Costa-Gomes and Crawford's experiment discussed next, *Dk*'s and *Lk*+1's guesses are strongly separated, and the results very strongly favor *Lk+1* over *Dk* rules.

Thus, subjects' guesses respect *k*-rationalizability for small *k* not because they explicitly perform iterated dominance, but because they follow rules that implicitly respect it.

B. Costa-Gomes and Crawford's (2006) ("CGC") experiments

Nagel's and related designs are distinguished by very large strategy spaces, which greatly increase the informativeness of their results.

But from the point of view of studying strategic thinking it is a weakness that a subject played only one game (although there was between-subjects variation across treatments).

Even though most subjects played their game repeatedly, their later choices confound strategic thinking with learning, so there was in effect only one observation per subject.

(Recall that first-round choices can still be viewed as initial responses to a game played as if in isolation if subjects treat their own influences on future choices as negligible.)

Even with very large strategy spaces, one observation yields limited information, and the results leave considerable ambiguity of interpretation regarding subjects' types.

By contrast, Stahl and Wilson's (1994, 1995) designs have the great advantage of series of different but related games, run to suppress learning and repeated-game effects.

But Stahl and Wilson's games have small strategy spaces, only three choices per player.

CGC’s design combines the variation through a series of different but related games of Stahl and Wilson’s designs with the large strategy spaces of Nagel’s design.

Subjects were randomly and anonymously paired to play a series of 16 different two-person guessing games, with no feedback.

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.

The design suppresses learning and repeated-game effects 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.)

In CGC’s guessing games, each player has his own lower and upper limit, both strictly positive to make the games finitely dominance-solvable.

(Players are not required to guess between their limits. Guesses outside the limits are automatically adjusted up to the lower 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 varied at most between subjects across treatments.)

CGC’s guessing games have essentially unique equilibria (“essentially” due to the automatic adjustment), determined (but 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 enhances the separation of equilibrium from other types and stress-tests equilibrium, which responds much more strongly to the product of targets than alternative rules do.

(It also reveals other interesting patterns; see Crawford 2008.)

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 the equilibrium is therefore determined by players' upper limits. (When the product is < 1, the equilibrium is determined by the lower limits.)

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.

CGC'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 type, in all 16 games.

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

Most of the analysis restricted attention to a list of behaviorally plausible types:

- *L0*, *L1*, *L2*, and *L3* as defined above, with *L0* uniform random between the limits
- *D1* and *D2* as defined above
- *Equilibrium*, which makes its equilibrium decisions (because CGC's games are all finitely dominance-solvable, traditional equilibrium refinements are not relevant)
- *Sophisticated*, which best responds to the probability distributions of others' decisions, estimated from the observed frequencies (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 as explained below, 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.

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

()

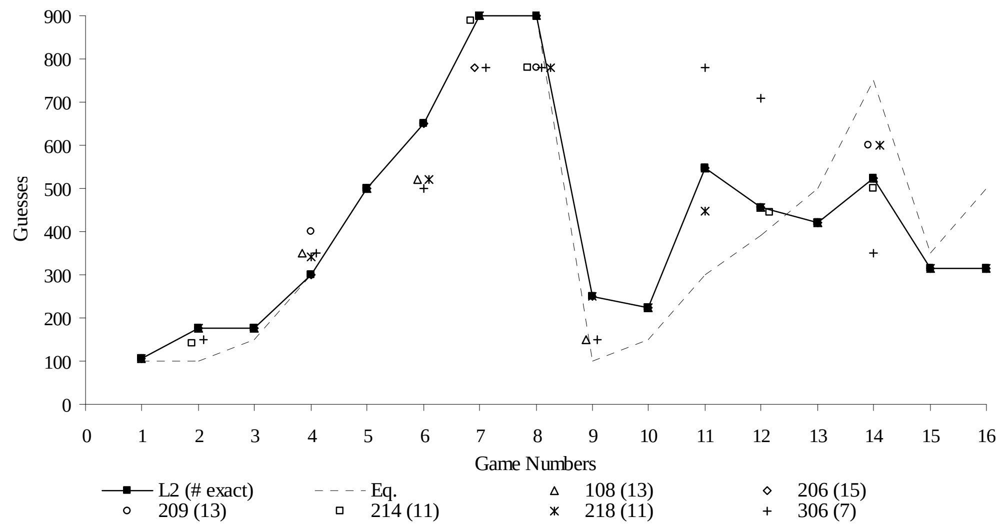

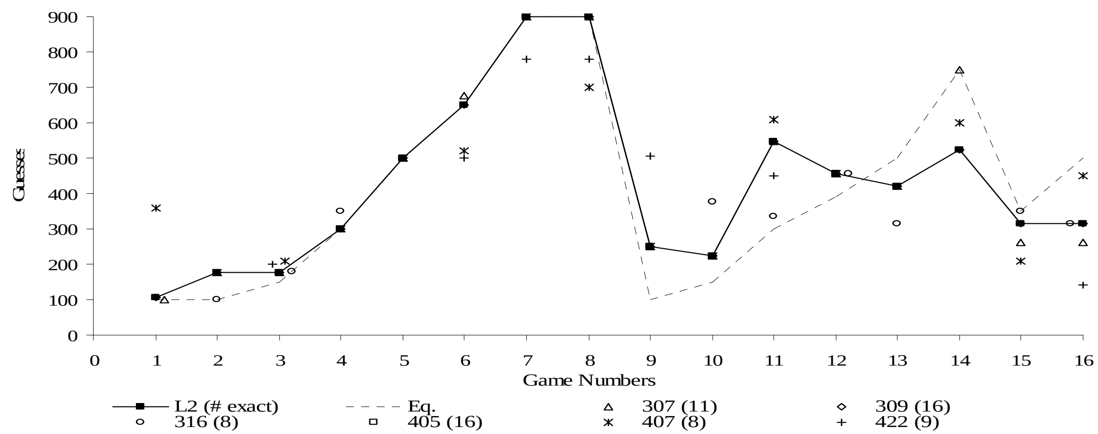

CGC's Figure 2. "Fingerprints" of 12 Apparent *L2* Subjects (Only deviations from *L2*'s guesses are shown.)

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 a type:

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.)

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

CGC's other 45 subjects made guesses that conformed less closely to a type.

But for all but 14 of them, violations of simple dominance were comparatively rare (less than 20%, versus 38% for random guesses), suggesting that their behavior was coherent, even though less well described by the types.

And 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.

CGC's Figure 1.

Aside on CGC's specification test

For the 45 subjects whose types must be estimated econometrically, there is some room for doubt about whether CGC's specification omits relevant types and/or overfits by including irrelevant types.

To test for overfitting and omission of relevant types, CGC conducted a specification test, which compares the likelihood of each subject's econometric type estimate with the likelihoods of estimates based on 88 *pseudotypes*, each constructed from one subject's guesses in the 16 games.

For a subject's type estimate to be credible, it should have higher likelihood than at least as many pseudotypes as it would at random: with 8 types, assuming approximately i.i.d. likelihoods, this makes $87/8 \approx 11$.

Some subjects' type estimates do not pass this test, and so they are left unclassified in columns 5 and 6 of CGC's Table 1.

With regard to omitted types, imagine that CGC had omitted a relevant type, say *L2*.

The pseudotypes of CGC's estimated *L2* subjects would then outperform the non-*L2* types estimated for them with *L2* omitted, and make approximately the same guesses.

Finding such a *cluster*, CGC diagnosed an omitted type, and studied what its subjects' guesses had in common to reveal its decision rule.

CGC found five clusters involving 11 subjects, who were also left unclassified in Table 1.

The paper's online appendix discuss what these 11 subjects seemed to be doing; most of it appears quite idiosyncratic.

Because a cluster must contain at least two subjects, in a larger sample there might be more clusters; but because such clusters did not reach the two-subject threshold in a sample of 88, they unlikely to be more than 2% of any larger sample.

CGC concluded from these tests that the types estimated to be in the population are truly relevant and that omitted types are at most 2% of the population, so not worth modeling.

End of aside

Lessons for modeling strategic behavior

CGC's 2006 *AER* analysis reinforces the conclusions already reached from Nagel's and related analyses:

No model that imposes homogeneity of thinking will do justice to subjects' behavior.

Subjects do not respect indefinitely iterated dominance or indefinitely iterated best responses; instead their decisions respect *k*-rationalizability for at most small values of *k*.

To these we can now add:

There are few if any no *Dk* subjects. People respect iterated dominance to the extent that their *Lk* types do, not because they explicitly perform it.

(This is reinforced by CGC's data on subjects' searches for hidden payoff information (Crawford 2008) and their data on "robot/trained subjects," which show that most subjects are capable of learning to follow *Dk*, but suggest that they find *Dk* rules far less natural than the analogous *Lk* rules.)

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.)

A hybrid level-*k* or CH 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:

Although about half of CGC's subjects' deviations from equilibrium remain unexplained by such a model, the specification test suggests that the deviations have little discernable structure.

Thus it may still be optimal to treat the remaining deviations as errors, and the part of the structure that can be identified can provide a secure basis for unbiased modeling of initial responses to games.

I stress that although CGC's evidence and analysis are more precise than previous studies, their conclusions are consistent with the results of earlier studies.

Model comparison: level-*k* versus equilibrium plus noise or LQRE

CGC's econometric analysis allows enough heterogeneity to nest equilibrium plus noise, represented by the *Equilibrium* type, with logit errors.

Only 11 of the 88 subjects in CGC's main treatments are estimated to be *Equilibrium*, and there is clear evidence that even they are following rules that only mimic *Equilibrium*, and that only in some of the games (CGC, pp. 1753-1754).

Aside on CGC's "near-*Equilibrium*" subjects

Focus on the eight subjects whose fingerprints are closest to equilibrium, and order the games by strategic structure, with the eight games with mixed targets on the right.

CGC's Figure 4 (next slide) then shows that these subjects' deviations from equilibrium almost always occur with mixed targets.

Thus it is nonparametrically clear that these subjects, whose equilibrium compliance is off the scale by normal standards, are actually following a rule that only mimics *Equilibrium*, and that only in games without mixed targets.

Yet all the ways we teach people to identify equilibria (equilibrium checking, best-response dynamics, iterated dominance) work as well with or without mixed targets.

Thus, whatever these subjects are doing, it's something we haven't thought of yet.

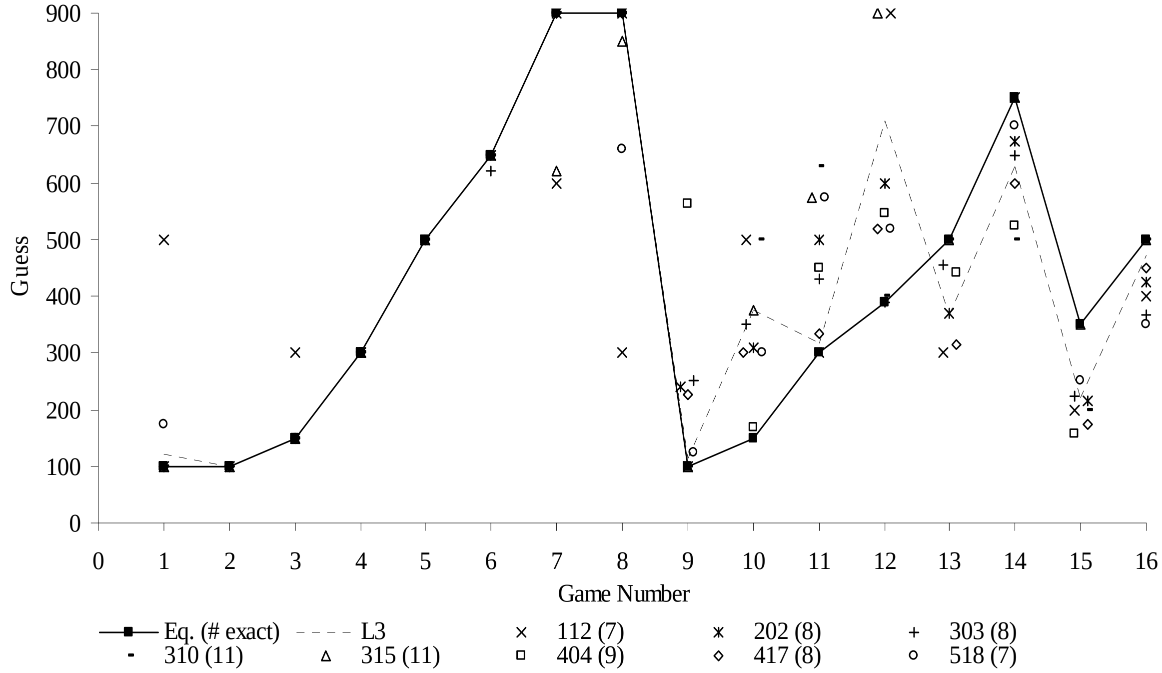

CGC's Figure 4. "Fingerprints" of 8 Apparent Baseline *Equilibrium* Subjects

(Only deviations from *Equilibrium*'s guesses are shown.
69 (54%) of these subjects' 128 guesses were exact *Equilibrium* guesses.)

End of aside

Model comparison: level-*k* versus equilibrium plus noise or LQRE continued

In CGC's design, LQRE choices are weakly separated from equilibrium plus logit noise choices. (CGC's (footnote 34, p. 1763) "median voter" result shows that they coincide except for small asymmetries in the payoff function due to automatic adjustment).

For CGC's 77 non-*Equilibrium* subjects, LQRE misses clear patterns in the data, as equilibrium plus logit noise did.

But these subjects' "errors" neither center on 0 nor exhibit the sensitivity to deviation costs assumed in a logit specification, perhaps because the errors are cognitive or structural, reflecting misspecification rather than an effort-accuracy trade-off.

Instead the errors have a clear deterministic structure, which is well described by the level-*k* model that emerges from CGC's estimates.

The evidence also suggests the absence of exotic types like Stahl and Wilson's (1995) *Worldly* that best respond to estimated mixtures of other types, some of them noisy with estimated precisions (CGC 2006, Section II.D).

This is indirect evidence that the best-response to others' noise aspect of QRE is unlikely to be empirically reliable.

Model comparison: level-*k* versus CH

Partly by a quirk of CGC's design, level-*k* and CH types' choices are not separated:

- Level-*k* and CH *L1*'s are identical by definition

- By CGC's (footnote 34 and 36, p. 1763) "median voter" result (which stems from the piecewise linearity and symmetry of the payoff function), for empirically plausible type distributions, a CH *L2* and *L3* are both identical to CGC's *L2*.

However, CGC's subjects' searches for hidden payoff information (Crawford 2008) are much more consistent with the implications of a level-*k* model than a CH model.

Further, to fit the data a CH Poisson parameter τ (roughly, the average *k*) must be approximately 1.5, which constrains the frequency of *L0* to 0.22.

By contrast, CGC's and other unconstrained estimates almost always assign *L0* a far lower frequency, usually 0.

This is a good sign, in that high frequencies 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.

The Poisson constraint is strongly binding in CGC's dataset, and with comparable error structures (though possibly not with the popular CH structure in which a random *L0* doubles as the errors for higher types), level-*k* will have an advantage in fit over CH.

Estimating an unconstrained type distribution as in a level-*k* model also provides a useful diagnostic: If the data can only be fitted by a weird type distribution—non-hump-shaped (in a homogeneous population) or with implausibly high frequencies of higher types—then the explanation is not credible.

Model comparison: Level-*k* versus NI

Given that NI is a flexible parameterization that includes LQRE as a special case, Haile, Hortaçsu, and Kosenok's (2008) critique applies to it a fortiori.

Preliminary analyses suggest that Goeree and Holt's (2004) favored NI specification, in which the noisiness of higher-order beliefs grows geometrically, the model is heavily overparameterized, yielding a range of possible decisions spanning level-*k* types as well as *Equilibrium*.

NI may therefore risk overfitting even in datasets that span multiple games.

4. M. M. Kaye's Far Pavilions: Responding to Payoff Asymmetries in Outguessing Games

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

—Koda Dad* to Ash/Ashok in M.M. Kaye's *The Far Pavilions* (1978, p. 97)

*played by Omar Sharif in the HBO miniseries
(his greatest contribution to game theory)

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

Early in *The Far Pavilions*, the main male character, Ash/Ashok, is trying to escape from his pursuers along a north-south road.

Both Ash and his pursuers must choose between north and south. Although Ash moves first, the pursuers must make their choice irrevocably before they learn Ash's choice, so their choices are strategically simultaneous.

South is warm, but north lie the Himalayas, with winter coming.

Imagine that if the pursuers catch Ash, they gain two units of payoff and Ash loses two, and that they both gain one extra unit for choosing South whether or not Ash is caught.

This yields the payoff matrix:

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

***Far Pavilions* Escape**

Examples like this are as common in experimental game theory as they are in fiction, but fiction sometimes more clearly reveals the thinking behind a decision.

Ash's mentor Koda Dad advises Ash to ride north, for the reason given in the quotation.

Ash overcomes his fear of freezing and follows this advice, the pursuers unimaginatively go south, and Ash escapes.

Koda Dad is advising Ash to choose the *L3* response to a uniform random *L0*.

(*L3* ties my personal best *k* for a clearly explained level-*k* type in fiction. I suspect even postmodern fiction may have nothing higher than *L3*: it wouldn't be credible.)

If the pursuers expect Ash to go south because it's "kinder", they must be modeling Ash as an *L1* response to a uniform random *L0*.

For, the payoff asymmetry on which this inference rests is decisive only if north and south do not differ in the probability of being caught, which is more important.

Thus, Koda Dad must be modeling the pursuers as *L2* and advising Ash to choose the *L3* response to a uniform random *L0*.

We could take the inference that Ash will go south because it's "kinder" literally as a response to a uniform random *L0*.

But there is a behaviorally more plausible interpretation in which the inference is a model of the pursuers' model of Ash's instinctive reaction ignoring strategic considerations, and given this, plausibly based on the principle of insufficient reason.

In a more complex game a uniform random *L0* could plausibly approximate random sampling of payoffs unstratified by the other player's strategy choices.

How does the level-k model compare in predictive success with an equilibrium model?

Escape has a unique equilibrium in mixed strategies, in which

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

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

Thus Ash's Pr{South}, $p^* = 1/4$, and the Pursuers' Pr{South}, $q^* = 3/4$.

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 $p = 1/4 < 1/2$).

In equilibrium the novel's observed outcome {Ash North, Pursuers South} has probability $(1 - p^*)q^* = 9/16$: much better than a random 25%.

By contrast, the level-*k* model implies decisions as follows:

Type	**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**

Thus the level-*k* model exactly predicts the outcome provided that Ash is either *L2* or (as we know from the quotation) *L3*, and the Pursuers are either *L1* or (as Koda Dad believes) *L2*.

Applications seldom come with an omniscient narrator telling us how players think.

Even so, if the game is clearly defined and we have data, we can specify a level-*k* model, derive its implications, and use them to estimate the type frequency distribution.

Alternatively, we can calibrate the model using previous estimates from similar settings.

Suppose, for example, that we calibrate a level-*k* model by assuming that each player role in Escape is filled from a 50-30-20 mixture of *L1*s, *L2*s, and *L3*s and there are no errors.

Then the predicted frequency with which Ash goes North is 1/2 and the frequency with which the Pursuers go South is 4/5.

Assuming independence, this implies that the observed outcome {Ash North, Pursuers South} has probability 2/5: less than the equilibrium predicted frequency of 9/16, but still noticeably better than a random 25%.

More importantly, the level-*k* model gracefully explains a puzzling divergence between observed aggregate behavior patterns and equilibrium predictions.

In outguessing games like Escape, and in closely related perturbed Matching Pennies games, the unique mixed-strategy equilibrium responds to the payoff asymmetry between South and North in a decision-theoretically intuitive way for the Pursuers' role ($q^* = 3/4 > 1/2$, the probability with which Pursuers go south in the analogous game with no north-south payoff asymmetry); but in a counterintuitive way for Ash's role ($p^* = 1/4 < 1/2$).

Yet experimental subjects' aggregate choices in initial responses to games like this reflect decision-theoretic intuition in both player roles. (Ash's counterintuitive choice would not contradict this pattern if he were a subject, because his revealed type is in the minority.)

In such games the level-*k* and CH models' predictions "quasi-purify" something roughly like a mixed-strategy equilibrium via the predictable heterogeneity of players' strategic thinking, while avoiding some implausible implications of equilibrium.

Consider the perturbed Matching Pennies games with countervailing payoff perturbations in both roles studied in Rosenthal, Shachat, and Walker's (2003) experiments.

They considered two versions of such games, one with deterministic and the other with stochastic money payoffs, but constructed to have the same normal form, with both player roles playing Left in equilibrium with the same probability, 2/3.

In each case the observed strategy frequencies deviated from the equilibrium frequencies in the decision-theoretically intuitive direction in both player roles (Figure 2, p. 281).

Rosenthal, Shachat, and Walker show (2003, Figure 3, p. 285) that LQRE with precision estimated for the same dataset can explain this intuitive pattern qualitatively, but that it underpredicts the frequency of Left play in both player roles.

A calculation like that done above for Escape shows that a calibrated level-k model with a 50-30-20 mixture of *L1*s, *L2*s, and *L3*s in each role and no errors can also explain the qualitative pattern, and even without reestimating the type frequencies comes approximately as close to the observed frequencies, overpredicting the frequency of Left play in one role and underpredicting it in the other.

5. Groucho's Curse:
Zero-Sum Betting and Auctions with Incomplete Information

"I sent the club a wire stating, 'Please accept my resignation. I don't want to belong to any club that will accept people like me as a member'."

—Groucho Marx (1959, p. 321), Telegram to the Beverly Hills Friar's Club

"Son…One of these days in your travels, a guy is going to show you a brand-new deck of cards on which the seal is not yet broken. Then this guy is going to offer to bet you that he can make the jack of spades jump out of this brand-new deck of cards and squirt cider in your ear. But, son, do not accept this bet, because as sure as you stand there, you're going to wind up with an ear full of cider."

—Sky Masterson, quoting his father in Damon Runyon (1934)

Although most laboratory evidence on strategic thinking comes from symmetric-information designs, most field evidence and applications involve settings with informational asymmetries.

It is therefore of great importance to extend whatever can be learned about strategic thinking in complete-information games to incomplete-information games.

I now discuss laboratory and field evidence on games with informational asymmetries, focusing on cases where the games allow clear inferences about strategic thinking.

A. Experiments on zero-sum betting

Experiments on zero-sum betting build on Milgrom and Stokey's (1982) no-trade theorem, which shows that if traders are weakly risk-averse and have concordant beliefs, and the initial allocation is Pareto-efficient relative to the information available at the time, then even if traders receive new private information, no weakly mutually beneficial trade is possible; and if traders are strictly risk-averse, no trade at all is possible.

For, any such trade would make it common knowledge that both traders had benefited, contradicting the hypothesis that the original allocation was Pareto-efficient.

This result has been called the Groucho Marx theorem because its logic resembles that of our Marx quotation.

By contrast with the conclusions of the theorem, speculative zero-sum trades are common in real markets. This fact has a number of possible explanations, of which one is nonequilibrium strategic thinking.

The experiments by Brocas, Carrillo, Camerer, and Wang (2009) (see also Camerer, Ho, and Chong 2004, Section VI) and Rogers, Camerer, and Palfrey 2009) we now discuss have the control required to distinguish between such explanations and those based on other factors such as hedging or the joy of gambling.

Brocas et al.'s design used simple three-state betting games, including this one:

player/state	**A**	**B**	**C**
1	25	5	20
2	0	30	5

Zero-Sum Betting Game

The rules of the game and the information structure were publicly announced, with the goal of inducing common knowledge.

Each of two players, 1 and 2, is given information about which of three ex ante equally likely states has occurred, A, B, or C. As indicated in Figure 1, player 1 learns either that the state is {A or B} or that it is C; player 2 learns that the state is A or that it is {B or C}.

Once informed, the players choose simultaneously between two decisions: Bet or Pass.

A player who chooses Pass earns 10 no matter what the state. If one chooses Bet while the other chooses Pass, they both earn 10.

If both players choose Bet, they get the payoffs in the table, depending on which state has occurred.

This game has a unique trembling-hand perfect Bayesian equilibrium.

In this equilibrium, player 1 told C will Bet because 20 > 10, and player 2 told A will Pass because 0 < 10.

Given this, player 1 told {A or B} will Pass, because player 2 will Pass if told A, so betting given {A or B} yields player 1 at most 5 < 10.

Given this, player 2 will Pass if told {B or C}, because player 1 will Pass if told {A or B}, so betting given {B or C} yields player 2 at most 5 < 10.

This covers all contingencies and completes the characterization of equilibrium, which shows that the game is weakly dominance-solvable in three rounds.

No betting takes place in equilibrium in any state.

Despite this clear conclusion, in Brocas et al.'s and similar experiments approximately half of the subjects bet.

To explain this Brocas et al. proposed a level-*k* model with a specification like those discussed above, but with *L0* adapted to allow for incomplete information.

Following Camerer, Ho, and Chong (2004, Section VI), Brocas et al. assumed that *L0* bids uniformly randomly, independent of its private information.

In judging this specification, bear in mind that *L0* is meant to describe a player's model of the instinctive starting point of others' strategic thinking, from which point of view it is behaviorally plausible that *L0* ignores others' private information.

As in previous level-*k* analyses, Brocas et al. took their *L1* to best respond to their *L0*, and their *L2* to best respond to their *L1*.

Following Crawford and Iriberri (2007a), we call an *L1* that best responds to a random *L0* a "random *L1*" even though it is not itself random; and we call an *L2* that best responds to a random *L1* a "random *L2*".

Given this, random *L1* player 1s will Bet if told {C} because it yields 20 > 10 if player 2 Bets, a random *L0* player 2 will bet with probability one-half in either contingency, and Betting is otherwise costless.

Unlike in equilibrium, Random *L1* player 1s will Bet if told {A or B} because it yields 25 in state {A} and a random *L0* player 2 will bet with probability one-half in {A}, it yields 5 in state {B} and a random *L0* player 2 will Bet with probability one-half in {B}, the two states are equally likely ex ante, so Betting if told {A or B} yields expected payoff (25 + 5)/2 = 15 > 10.

Random *L1* player 2s will Pass if told {A}, because it yields 0 < 10. Unlike in equilibrium, random *L1* player 2s will Bet if told {B or C}, because it yields 30 in state {B} and a random *L0* player 1 will bet with probability one-half in {B}, it yields 5 in state {C} and a random *L0* player 1 will Bet with probability one-half in {C}, the two states are equally likely ex ante, so Betting if told {B or C} yields expected payoff (30 + 5)/2 = 17.5 > 10.

Thus, if all subjects were random *L1s*, 100% of player 1s and 67% of player 2s would Bet, too much in each role; and betting would occur only in states B and C, not true in the data.

Brocas et al.'s data analysis finds clusters corresponding to random *L1*s, *L2*s, and *L3*s (who in this setting correspond to *Equilibrium* players), and an additional cluster of apparently irrational players, which mixture of types fits more exactly.

B. Auction experiments

There is a rich literature on sealed-bid incomplete information auction experiments, which despite similar goals and methods has developed largely independently of the literature on game experiments.

In sealed-bid auction experiments, subjects' initial responses tend to exhibit overbidding, relative to the risk-neutral Bayesian equilibrium, in first- or second-price, independent-private-value or common-value auctions (Goeree, Holt, and Palfrey 2002, Kagel and Levin 1986).

The literature on auction experiments has proposed different explanations of overbidding in private- and common-value auctions: "joy of winning" and/or risk-aversion for private-value auctions, and the winner's curse for common-value auctions.

Moreover, these explanations are only loosely related to explanations that have been proposed for deviations from equilibrium in other games.

Kagel and Levin (1986) and Eyster and Rabin (2005) sought to unify the explanations of nonequilibrium behavior in auctions and other games.

Also noteworthy, but not discussed here, are Charness and Levin's (2009) experiments on Samuelson and Bazerman's (1985) *Acquiring a Company* game.

Kagel and Levin formalize the intuition behind the curse in models in which "naïve" bidders do not adjust their value estimates for the information revealed by winning (essentially, random *L1* bidding), but otherwise follow equilibrium logic.

Eyster and Rabin's notion of "cursed equilibrium", in which people underestimate the correlation between others' decisions and private information but otherwise behave as in a Bayesian equilibrium, generalizes Kagel and Levin's model to allow intermediate levels of value adjustment, ranging from standard equilibrium with full adjustment to "fully-cursed" equilibrium with no adjustment; and from auctions and bilateral exchange to other kinds of incomplete-information games.

Both models allow players to deviate from equilibrium only to the extent that they do not draw correct inferences from the outcome. Thus their predictions coincide with equilibrium in games in which such inferences are not relevant, and they do not help to explain non-equilibrium behavior in independent-private-value auctions.

Crawford and Iriberri (2007a) propose a level-*k* analysis that provides an alternative way to unify the explanation of results for initial responses in auction experiments, without invoking joy of winning or risk-aversion.

Their analysis makes it possible to explore the robustness of equilibrium auction theory to failures of the equilibrium assumption, and establishes a connection between large bodies of experiments on auctions and on strategic thinking.

The key issue is how to specify *L0*; in auctions there are two natural possibilities:

Random L0, analogous to the *L0* in level-*k* analyses of zero-sum betting, bids uniformly on the interval between the lowest and highest possible values (even if above own realized value, in keeping with the view that *L0* represents a player's model of others' instinctive, nonstrategic responses to the game).

Truthful L0, which is meaningful in auctions though not in all incomplete-information games, bids its expected value conditional on its own signal.

Crawford and Iriberri build separate type hierarchies on these *L0*s: *Random* (*Truthful*) *Lk* is defined by iterating best responses from *Random* (*Truthful*) *L0*; and allow each subject to be one of the types, from either hierarchy.

For a given *Lk* type, and just as in an equilibrium analysis, the optimal bid must take into account value adjustment for the information revealed by winning (only in common-value auctions), and the bidding trade-off between the higher price paid if the bidder wins and the probability of winning (only in first-price auctions).

Crawford and Iriberri show that their level-*k* model allows a tractable characterization of these aspects of the bidder's problem, which closely parallels the equilibrium characterization.

With regard to value adjustment, Random *L1* does not condition on winning because Random *L0* bidders bid randomly, hence independently of their values; thus Random *L1* is "fully cursed" in Eyster and Rabin's sense.

The other types do condition on winning in various ways, but this conditioning tends to make bidders' bids strategic substitutes, in that the higher others' bids are, the greater the (negative) adjustment.

Thus, to the extent that Random *L1* overbids, Random *L2* tends to underbid (relative to equilibrium): If it's bad news that you beat equilibrium bidders, it's even worse news that you beat overbidders.

The bidding tradeoff, by contrast, can go either way, as in an equilibrium analysis.

The question, empirically, is whether an estimated mixture of Random *L1* overbidding and Random *L2* underbidding fits the data better than alternative models.

In three of the four leading cases Crawford and Iriberri study, a level-*k* model does better than equilibrium plus noise, cursed equilibrium, and/or LQRE.

For the remaining case (Kagel and Levin's first-price auction), the most flexible cursed equilibrium specification has a small advantage.

Except in Kagel and Levin's second-price auctions, the estimated type frequencies are similar to those found in other experiments:

Random and Truthful *L0* have low or zero estimated frequencies, and the most common types are (in order of importance) Random *L1*, Truthful *L1*, Random *L2*, and sometimes *Equilibrium* or Truthful *L2*.

Crawford and Iriberri estimated large frequencies (59-65%) of random *L1* bidders and much smaller but significant frequencies of random *L2* (4-9%), truthful *L1* (9-18%), and truthful *L2* (1-16%).

C. Field studies: Movie Opening and Lowest Unique Positive Integer Games

Similar methods have now been used in innovative field studies.

Brown, Camerer, and Lovallo (2010) use field data to study an incomplete-information signaling game with verifiable signals.

Film distributors face a choice between “cold opening” a movie and pre-releasing them to critics in the hope that favorable reviews will increase profits.

In perfect Bayesian equilibrium, cold-opening should not be profitable, because moviegoers will infer low quality for cold-opened movies and the process will unravel.

Yet distributors sometimes cold-open movies, and in a set of 856 widely released movies, cold opening produced a 15% increase in domestic box office revenue, consistent with the hypothesis that some moviegoers did not infer low quality from cold opening.

However, movie distributors appear not to be rational, since only 7% of movies are opened cold despite the profit advantage.

After preliminary tests that rule out more conventional explanations, Brown, Camerer, and Lovallo try variants of cursed equilibrium and level-*k*/CH models.

In the latter variants, *Lk* best responds to *Lk-1* rather than an estimated mixture of all lower-level types as it would in a CH model; but *Lk* responds to *Lk-1*'s decision noise as in an LQRE model, a choice that is not standard in level-*k* or CH models.

Further, L0 for moviegoers assumes a uniform distribution over the whole range of possible qualities; although sample-mean quality might seem more natural here, the authors say that that does not work well either.

The main findings are that the best fitting cursed-equilibrium model has moviegoers almost fully cursed (drawing no inferences regarding cold-opened movies) but studios not cursed at all; given the specification employed, the resulting model is like a partially "cursed" (moviegoers but not distributors) version of LQRE.

The best fitting level-*k*/CH model again has moviegoers almost fully cursed (τ, the average k, is 1.12 where 1 is fully cursed, which given the assumed Poisson distribution for k implies that the population frequency of L0s is 33%) but studios very sophisticated ($\tau = 8.5$). The best fitting level-*k*/CH model has a significant likelihood advantage over the best fitting cursed-equilibrium or LQRE model.

Overall, neither model really explains the behavior of studios. This may be unsurprising, because studio overoptimism is plausible and there has been a huge recent trend in the percentage of cold-opened movies, complicated by changes in technology.

Östling, Wang, Eileen Chou, and Camerer (2009) study a novel set of field data from a Swedish gambling company, which ran a competition for a short period of time involving a “lowest unique positive integer” or LUPI game.

In the LUPI game, players pick positive integers and the player who chose the lowest unique (not chosen by anyone else) number wins a prize.

Except for the uniqueness requirement, the game closely resembles a first-price auction.

The game would have complete information except that participants had no way to know how many others would enter in a given week.

The authors deal with this by adapting Myerson’s (2000) “Poisson games” model, in which fully rational players face Poisson-distributed uncertainty about the number of players.

The paper characterizes the LUPI game’s unique symmetric Poisson-Nash equilibrium, and compares it to the predictions of versions of QRE and CH models, using both the field data and data from experiments using a scaled-down version of the LUPI game.

Both the field and laboratory data show participants choosing very low and very high numbers too often, relative to the Poisson-Nash equilibrium, and avoiding round numbers.

However, participants' initial responses are surprisingly close to the equilibrium, given that the setting makes it almost inconceivable that they could be computing it, and learning enhances this correspondence over time.

The authors propose a CH model in which the types best respond to the noise in others' decisions and there is an unusual power error distribution.

This model explains why participants' initial responses are so close to the equilibrium, and tracks some of the other main patterns in the data.

An LQRE model, by contrast, gets the pattern of deviations from equilibrium qualitatively wrong, but this could be an artifact of different distributional assumptions.

D. Level-k auction design

A number of recent papers reconsider mechanism design taking a "behavioral" view of individual decisions or probabilistic judgment, but to date there are very few analyses of design outside the equilibrium paradigm.

Yet design inherently involves the creation of new games, and it may be important for an application to work the first time.

Further, assuming equilibrium may yield theoretically optimal designs that are too complex for confidence in equilibrium behavior.

Replacing equilibrium with a model that better describes people's responses to new and/or complex games should allow us to design more effective mechanisms.

It also suggests an evidence-based way to assess the robustness of mechanisms, something previously left to intuition.

In a level-*k* analysis, a "robust" mechanism that implements desired outcomes in dominant strategies or is dominance-solvable in one or two rounds may have an actual advantage over a more complex mechanism that theoretically implements better outcomes, but only in equilibrium.

Crawford, Kugler, Neeman, and Pauzner (2009) explored relaxing equilibrium in mechanism design by conducting a level-*k* analysis of optimal auction design.

They considered the leading case of an optimal (expected-revenue maximizing) single-object sealed-bid auction with two symmetric bidders who have independent private values, for which there is a complete equilibrium-based analysis (Myerson 1981).

To focus sharply on strategic behavior, they maintained the standard rationality assumptions regarding decisions and judgment.

They modeled strategic behavior via a level-*k* model that follows Crawford and Iriberri's (2007a) analysis of data from leading auction experiments, with either a random *L0* that bids uniformly over the natural range of bids or a truthful *L0* that bids its private value.

They assumed that bidders are drawn from a given population of level-*k* types, known to the designer.

In representative examples, they considered what reserve prices are optimal and how much revenue they yield in first-price auctions.

They also considered the optimality of auction forms and the use of exotic auctions that exploit bidders' non-equilibrium beliefs to exceed Myerson's revenue bound.

Crawford et al. show, trivially and unsurprisingly, that with independent private values, revenue-equivalence breaks down.

Because a second-price auction makes the equilibrium bid a dominant strategy, level-*k* bids coincide with equilibrium bids, hence a second-price auction yields only the equilibrium expected revenue.

By contrast, level-*k* bidders in a first-price auction can deviate from equilibrium, and they give an example to show that such an auction with a suitable reserve price can yield higher expected revenue than the best second-price auction.

Crawford et al. also give examples in which the optimal reserve price is large with equilibrium bidders but small with level-*k* bidders, and vice versa.

Interesting open questions are when a reserve induces more aggressive bidding for equilibrium than level-*k* bidders, and the extent to which this makes optimal level-*k* reserves higher than optimal equilibrium reserves.

Finally, Crawford et al. give an example to show that in theory, a designer can use exotic auction forms to exploit level-*k* bidders' non-equilibrium beliefs to obtain very large expected revenues.

They note, however, that their formulation of the design problem takes the level-*k* model's specification as given, independent of the auction design, just as the standard formulation assumes that bidders will play an equilibrium for any design.

Although the specification is based on substantial experimental evidence, there is reason to doubt the exogeneity assumption, particularly for exotic auctions that which go beyond the evidence on which our specification is based.

A general formulation of the design problem must take a position on how the design influences the rules that describe bidders' behavior and develop new methods to deal mathematically with that influence.

Even without such influences, the heterogeneity of level-*k* beliefs and behavior greatly complicates the characterization of optimal auctions.

In the standard analysis there is no loss of generality in using the revelation principle to restrict attention to direct mechanisms because, if equilibrium is assumed (with a selection rule in case of multiple equilibria), a bidder's private value is all that is needed to predict his behavior.

Given the restriction to direct mechanisms, the design problem is well-behaved enough that it is guaranteed to have a solution.

The example given in the paper shows that this is no longer the case with level-*k* bidders, even if their level-*k* types are all the same, and even if this is known to the designer.

With a heterogeneous population of types, the problem becomes more complex.

Bidders with the same private values but different level-*k* types have different beliefs and will generally behave differently.

It appears that Myerson's (1981) methods can be used to characterize an optimal auction if the designer knows that the population is homogeneous, and knows its type; and if the class of possible designs is restricted to rule out those that are too exotic for an optimal auction to exist.

But if the population is heterogeneous the problem becomes multidimensional and much more difficult; and the high-dimensional reporting mechanisms one would consider for this case complicate the specification of *L0* and the influence of design on behavior.

6. Kahneman’s Entry Magic: Coordination via Symmetry-Breaking

“…to a psychologist, it looks like magic.”

—Kahneman 1988, quoted in Camerer, Ho, and Chong (2004)

Kahneman’s “magic” refers to the fact that subjects in his own and others’ market-entry experiments (see also Rapoport et al. 1998 and Rapoport and Seale 2002) achieve systematically better coordination ex post than in the natural equilibrium benchmark.

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

In these experiments, n subjects choose simultaneously between entering (“In”) and staying out (“Out”) of a market with given capacity.

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.

For simplicity, Out yields 0 profit, no matter how many subjects enter.

In these games, efficient coordination requires breaking the symmetry of players' roles.

But because players cannot distinguish their roles, it is not sensible to predict systematic differences across roles in behavior.

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

This mixed-strategy equilibrium yields an expected number of entrants approximately equal to market capacity, but there is a positive probability that either too many or too few will enter.

Even so, subjects in market-entry experiments have systematically better coordination ex post (number of entrants stochastically closer to market capacity) than in the symmetric equilibrium.

A. A level-k Analysis of Two-Person Entry/Battle of the Sexes Games

Camerer, Ho, and Chong (2004, Section III.C) explain Kahneman's magic via a cognitive hierarchy model, in which the heterogeneity of strategic thinking allows some players to mentally simulate others' entry decisions and accommodate them.

The more sophisticated players become somewhat like Stackelberg followers, which in entry games yields coordination benefits for all.

Here I illustrate their analysis in a two-person Battle of the Sexes game, which is like a two-person market-entry game with capacity one, and which makes the central points as simply as possible.

For simplicity, I also substitute a level-*k* model for their CH model.

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

Consider a two-person Battle of the Sexes game with $a > 1$:

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

Battle of the Sexes

The unique symmetric equilibrium is in mixed strategies, with $p \equiv \Pr\{\text{In}\} = a/(1+a)$ for both players.

The mixed-strategy equilibrium expected coordination rate is $2p(1 - p) = 2a/(1+a)^2$, and players' equilibrium expected payoffs are $a/(1+a)$.

This expected coordination rate is maximized when $a = 1$, where it takes the value ½.

With $a > 1$ these expected payoffs $a/(1+a) < 1$: worse for each player than his worst pure-strategy equilibrium. As $a \to \infty$, $2a/(1 + a)^2 \to 0$ like $1/a$.

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.

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

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**
Outcomes in Battle of the Sexes				

*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.

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 near ½.

For more extreme values the rate is worse, converging to 0 as $v \to 0$ or 1.

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

The level-*k* analysis suggests a view of tacit coordination profoundly different from the traditional view, and illustrates the importance of the heterogeneity of strategic thinking the model allows.

With level-*k* thinking, equilibrium and selection principles such as risk- or payoff-dominance (Harsanyi and Selten 1987) play no direct role in players' thinking.

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

Even though players' 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.

Camerer, Ho, and Chong (2004, Section III.C) and Chong, Camerer, and Ho (2005, Section 2.1) argue that in this context, CH models fit better than level-*k* models because they yield smooth monotonicity of entry rates as market capacity increases, as seen in the data, while a level-*k* model implies a step function; and because CH models imply increasingly rational expectations as *k* increases, unlike level-*k* models which cycle perpetually in these games.

However in most of the datasets they consider, unlike in their stylized CH model, there are congestion effects that allow payoff-sensitive logit errors like those in a typical level-*k* analysis smooth things as well.

Further, cycling and the limiting argument for rationality increasing with *k* have little or no relevance when k seldom goes above 3.

One question I do not consider here is whether a level-*k* model can explain the pattern in the data that entry rates are too high for low capacities and too low for high capacities, which the CH model explains by estimating a high frequency of a random *L0* type.

6.2. Field Studies: CH Analyses of Entry Games

The same issues arise in field settings such as those studied using complete-information models by Goldfarb and Yang (2009) and Goldfarb and Xiao (2010), which both use CH models.

These studies provide only limited comparison of alternative models of strategic thinking, but they are of particular interest because they are among the first studies of nonequilibrium models of strategic thinking using field data.

Goldfarb and Yang (2009) apply a complete-information CH model to explain choices by managers at 2,233 Internet Service Providers (ISP) in 1997 whether or not to offer their customers access through 56K (versus the standard then, 33K) modems.

There were two possible 56K technologies: Rockwell Semiconductor’s K56Flex and US Robotics’s X2. Thus an ISP manager could make one of four choices: (i) adopt neither technology, (ii) adopt Rockwell’s, (iii) adopt US Robotics’s or (iv) adopt both.

Controlling for market and ISP-specific characteristic, Goldfarb and Yang adapted the CH model to describe the heterogeneity in strategic sophistication or ability among the SPI managers in these decisions.

They assumed (departing from the standard specification) that an *L0* manager maximizing profits on the assumption that he will be a monopolist.

An *L1* manager assumes that his competitors will be *L0*s. An *L2* manager assumes its competitors will be an estimated mixture of *L0*s and *L1*s, and so on.

Goldfarb and Yang found significant evidence of heterogeneity of strategic sophistication among managers, with an estimated τ, the average k in a CH model, of 2.67—a bit higher than most previous estimates, but reasonable.

The CH model fits no better than an equilibrium plus noise model, but CH estimates have interesting and plausible implications.

Managers behaved more strategically, in the sense of having higher estimated *k*s, if they competed in larger cities, with more firms, or in markets with more educated populations.

Those managers estimated to be more strategic in 1997 (always advantageous in their CH model) were more likely to survive through April 2007.

Interestingly, the estimated CH models suggests that heterogeneity of strategic thinking, relative to equilibrium, slowed the diffusion of the new 56K technology, in that more strategic managers were *less* likely to adopt, anticipating more competition.

Goldfarb and Xiao (2010) applied a complete-information CH model to explain managers' choices whether or not to enter local U. S. telecommunications markets after the *Telecommunications Act* of 1996, which allowed free competition in such markets.

They use Goldfarb and Yang's (2009) specification of *L0*.

They found that more experienced and/or better educated managers did better, in the sense of entering markets with fewer competitors, on average; having better survival rates; and having higher revenues, conditional on survival.

Estimated strategic thinking goes up from 1998 to 2002.

The CH model fits much better than an equilibrium plus noise model in 1998, but only slightly better in 2002, in keeping with the view that models like CH are well suited to initial responses to novel situations, but are less relevant once players have had enough experience to converge to equilibrium.

7. Bank Runs: Coordination via Assurance

“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 (2000).

Here Summers has in mind an *n*-person coordination game with Pareto-ranked equilibria.

In such games models that make unique predictions despite multiple equilibria have important advantages, discussed below.

The traditional path to unique predictions is via models of equilibrium selection, in which players first identify the set of equilibria and then—other refinements having no power in coordination games—discriminate among them via coordination refinements such as risk- or payoff-dominance (Harsanyi and Selten 1987).

A popular modern analog of such refinements is Carlsson and Van Damme's (1993) notion of "global games" (see also Morris and Shin 1998), which uses iterated dominance in a stochastically perturbed version of the original game to derive unique predictions, which in simple games coincide with risk-dominance.

This simple global-games view has been a tractable workhorse model of behavior in Stag Hunt and bank runs games.

But viewed as a model of initial responses—as it must be to describe bank runs—its assumptions are subject to doubt because the perturbed game is arbitrarily chosen and because indefinite reliance on iterated dominance is behaviorally questionable

By contrast, in level-*k*/CH models players use the same decision rules to choose their strategies with or without multiple equilibria; and both equilibrium and equilibrium selection are accidental but statistically predictable by-products of how those rules interact with the game.

This makes level-*k*/CH models of coordination consistent with the leading of decisions in other settings and behaviorally more plausible.

In symmetric coordination games the higher payoffs of equilibria attract level-*k*/CH as well as equilibrium players, so the likely level-*k*/CH outcome is some equilibrium, and that equilibrium is likely to be the risk-dominant one in simple examples, just as a global games analysis predicts.

But a level-*k*/CH model predicts the likelihood of coordination failure and the forms it may take, and in more complex games equilibrium selection may not follow predictions based on risk-dominance or global games (Crawford 1995; Costa-Gomes, Crawford, and Iriberri 2009).

A level-*k*/CH analysis also highlights an issue that is not considered in the traditional literature, how players model the correlation of others' strategy choices, on which the experimental evidence gives some guidance that departs from tradition.

The game Summers describes can be represented by a payoff table (not a matrix!):

		Summary statistic In	Summary statistic 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.)

What will happen?

In this example the coordination refinement of payoff-dominance uniquely favors the all-In equilibrium, for any value of the population size *n*.

This again seems behaviorally unlikely even for small *n*.

The basic idea of risk-dominance (the precise formalization is controversial, and is fully agreed on only in two-action games) 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 payoff assumptions risk-dominance favors the all-Out equilibrium for any $n > 2$, even if only two-thirds need to stay In.

A global games analysis yields the same conclusion as risk-dominance here.

Now consider a level-*k* model.

In this context an *L0* in the style of Graham's Mr. Market is behaviorally plausible, but that would require a complex discussion of market psychology.

To illustrate how the model works, I assume instead a uniform random *L0*.

Recall that 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 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 an *L0* in the style of Mr. Market.

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.

By contrast, a global games analysis relies on indefinitely iterated dominance in a game with payoff uncertainty artificially grafted onto its structure in a particular way; and the empirical support even for finitely iterated dominance is weak.

8. Non-Equilibrium Econometrics: Structural Alternatives to Incomplete Models

How might the availability of structural non-equilibrium models that reliably describe initial responses to games change the way we think about data?

Although some empirical applications concern games that are dominance-solvable in small numbers of rounds, many involve games that are not, and many others involve games with multiple equilibria.

In games that are not sufficiently dominance-solvable, finitely iterated dominance and *k*-rationalizability are incomplete in that they do not specify a unique (though possibly probabilistic) prediction conditional on the value of the behavioral parameters.

In games with multiple equilibria, equilibrium plus noise but without refinements is incomplete in the same general sense.

In the empirical literature, such incompleteness has been dealt with in one of two ways:

- By accepting a theory's set-valued restrictions as the model's only implications and testing them (for *k*-rationalizability—they call it "level-*k* rationality"—in Aradillas-Lopez and Tamer 2008; or unrefined equilibria in Echenique and Komunjer 2009).
- By estimating an unrestricted probability distribution over the set of equilibria (Bresnahan and Reiss 1991).

However, just as experimental results suggest that equilibrium is too strong to be descriptive of people's responses to novel or complex games, it also suggests that *k*-rationalizability and even rationalizability are too weak.

Rationalizability sometimes agnostically allows beliefs that are behaviorally outlandish, even though consistent with common knowledge of rationality.

Because CGC's experimental results suggest that to the extent that people respect *k*-rationalizability, they do so not because they perform finitely iterated dominance leading to a set of *k*-rationalizable decisions, but because they follow a level-*k* decision rule that selects a specific such decision, it seems behaviorally natural to replace *k*-rationalizability (and equilibrium) by a structural level-*k* model that has the advantage of being complete.

In settings where this can be done without risking serious misspecification, it seems likely to yield significantly more useful econometric models.

Aradillas-Lopez and Tamer (2008) provide some indirect evidence on the potential benefits of structural non-equilibrium models by comparing the identification powers of equilibrium and *k*-rationalizability in two-person entry games without or with privately observed payoff perturbations; and in first-price auctions with incomplete information and independent private values.

In entry games attention centers on identification and estimation of payoff parameters, which are normally unobservable in the field.

In auctions attention centers on identification and estimation of bidders' value distributions, which are again normally unobservable in the field.

The standard approach assumes equilibrium and shows that the parameters of interest are identified (parametrically or nonparametrically).

Aradillas-Lopez and Tamer show that weakening equilibrium to *k*-rationalizability implies weaker identifying restrictions—sometimes much weaker, for low values of *k*—and that individuals' *k*'s are not fully identified.

In entry games, 1-rationalizability only slightly restricts the payoff parameters:

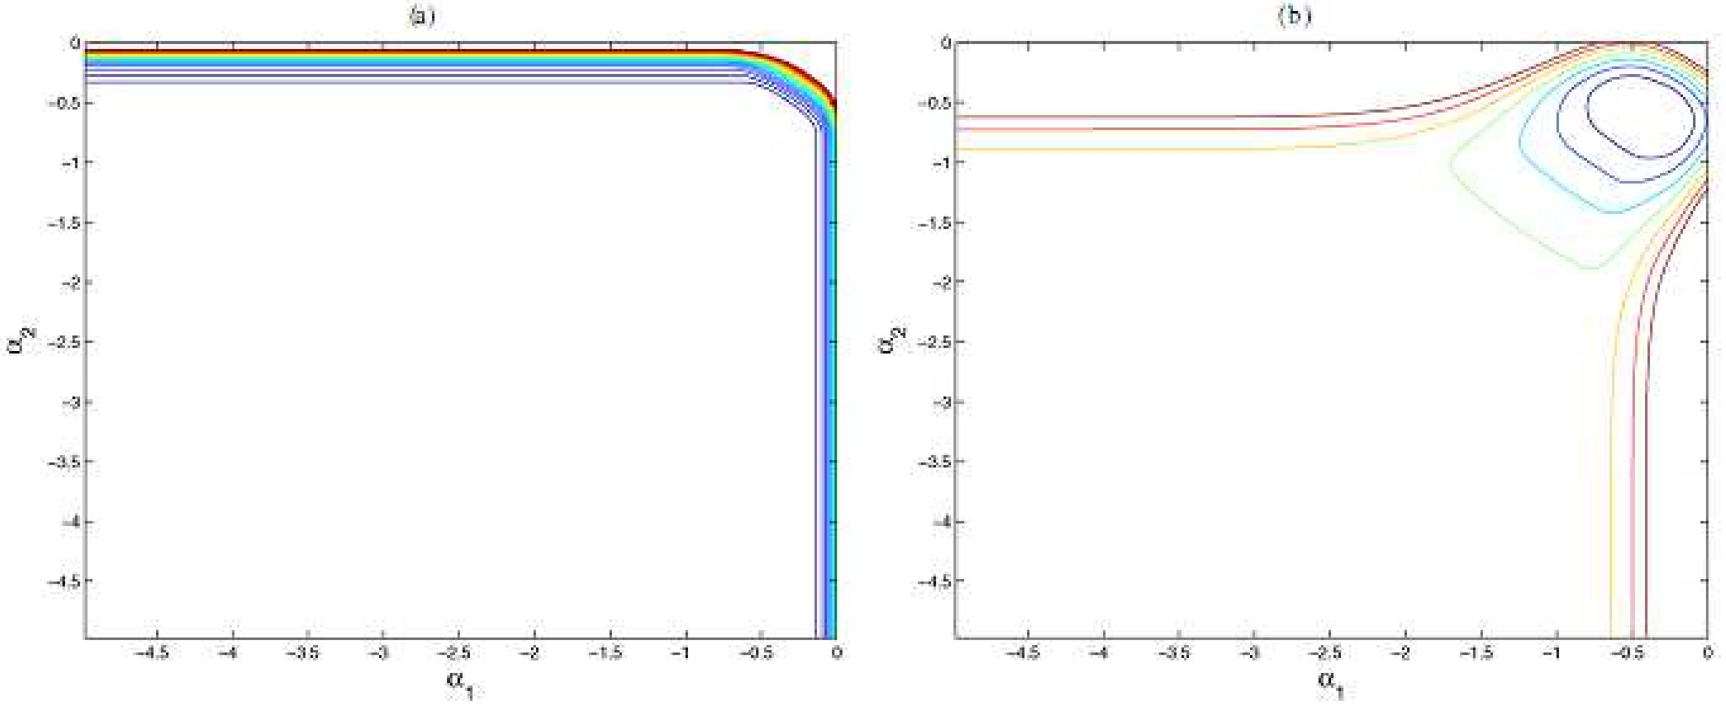

Figure 3. Identification set under Nash and 1-level rationality. Shown the identified regions for (α_1, α_2) under $k = 1$ rationality (a) and Nash (b). We set in the underlying model $(\alpha_1, \alpha_2) = (-.5, -.5)$. The model was simulated assuming Nash with (0, 1) selected with probability one in regions of multiplicity. Note that in (a), the model only places upper bounds on the alphas, whereas in (b) (α_1, α_2) are constrained to lie a much smaller set (the inner "circle").

Aradillas-Lopez and Tamer's Figure 3

In first-price auctions Aradillas-Lopez and Tamer note (following Battigalli and Siniscalchi 2003) that *k*-rationalizability implies only a weak upper bound on bids, which shrinks with *k* but for any *k* allows bids both above and below equilibrium; with correspondingly weak bounds on value distributions.

Benjamin Gillen, "Identification and Estimation of Level-*k* Auctions" provides additional evidence on the benefits of structural non-equilibrium models.

He shows that in a level-*k* model (based on Crawford and Iriberri's 2007 model), under a reasonable (but not completely unrestrictive) assumption on the separation of different types' bidding functions, with variation in the number of bidders the value distributions and individual bidders' *k*s are identified, parametrically or nonparametrically.

The difference arises because Gillen's level-*k* model completes Aradillas-Lopez and Tamer's *k*-rationalizability model, which with enough data makes it possible to estimate the level-*k* model's additional structural parameters with bidders' value distributions.

CGC (footnote 42, p. 1766) makes a similar point in a different way.

CGC note that in their maximum likelihood estimation of a model of subjects' guesses and searches for hidden payoff information, the guess part of the log-likelihood is nearly six times larger than the search part.

This occurs because their theory of subjects' decisions makes very precise predictions of a subject's decisions, conditional on his type.

By contrast, CGC's theory of cognition and search imposes (via filters described in the paper) only weak, set-valued restrictions on a subject's searches, conditional on his type:

Although CGC's theory of decisions is complete, their theory of search is incomplete.

As a result, the search restrictions are much more likely than the decision restrictions to be satisfied by chance, which causes the disparity in likelihood weights.

Turning to games with multiple equilibria, the freedom that assuming rationalizability or estimating an unrestricted probability distribution over the set of equilibria can yield severe overfitting and/or very weak tests.

Costa-Gomes, Crawford, and Iriberri (2009) address this issue for Van Huyck, Battalio, and Beil's (1990, 1991) coordination games, in which any of the seven pure strategies is both rationalizable and consistent with one of the seven pure-strategy equilibria.

Using VHBB's data to estimate an unrestricted probability distribution over equilibria yields good fits, but it also yields estimates that vary incoherently across games and don't inspire confidence for beyond-sample prediction.

CGCI complete equilibrium plus noise by adding coordination refinements risk- or payoff-dominance (in turn), to put it on a more equal footing with LQRE, level-*k*, CH, and NI models, which are already complete.

9. Yushchenko and Lake Wobegon: Non-neutral Framing in Outguessing Games

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

> —comment, quoted in Chivers (2004), on the poisoning of Ukrainian presidential candidate—now ex-president—Viktor Yushchenko.

"…in Lake Wobegon, the correct answer is usually 'c'."

> —Garrison Keillor (1997) on multiple-choice tests (quoted in Attali and Bar-Hillel 2003)

Both quotations refer to simultaneous-move zero-sum two-person games with unique mixed-strategy equilibria.

In the first, the players are a government assassin choosing one of several dinners at which to try to poison Yuschenko, only one of which is with officials of the government suspected of wanting to poison him; and an investigator who has the resources to check only one of the dinners.

In the second, the players are a test designer deciding where to hide the correct answer and a clueless test-taker trying to guess the hiding place.

Although there is nothing as uniquely salient as the dinner with government officials, psychologists 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; see Christenfeld 1995 and Rubinstein, Tversky, and Heller 1996).

In each case 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 quotations is plainly strategic, but non-equilibrium:

To the first, for example, 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.

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 (2007b).

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?

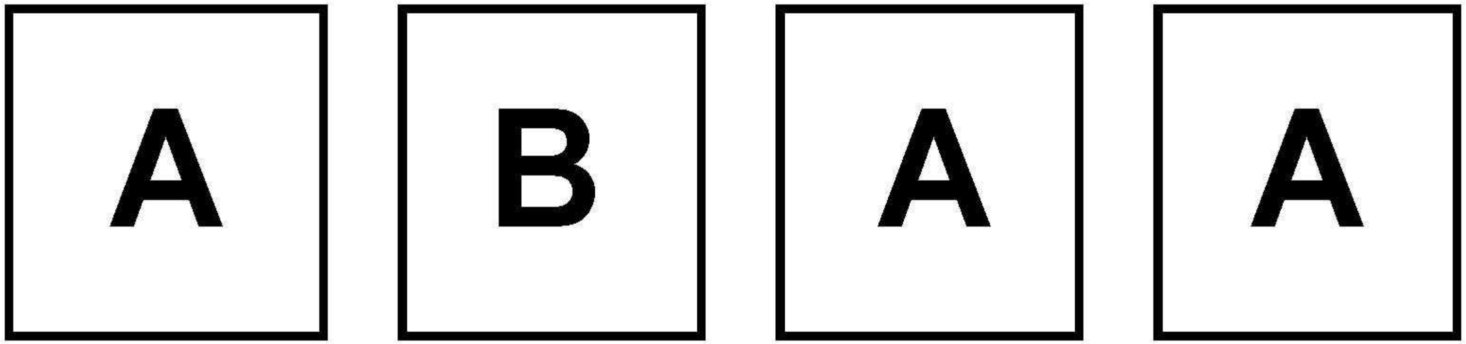

RTH’s framing of the hide-and-seek game is non-neutral in two ways:

- 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 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.

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 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."

But strategic thinking need not be equilibrium thinking.

Crawford and Iriberri's analysis suggests that RTH's subjects were actually 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.

Their 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.

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%)?

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

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

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), 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, discussed below, also suggests that allowing *L0* to vary across roles leads to overfitting.

Aside: Evaluating the model's explanation: Overfitting and portability

Although prior intuitions about the likely hump shape and location of the type distribution impose some discipline in specifying a level-*k* model, the freedom to specify *L0* leaves room for doubts about overfitting and portability, the extent to which a model estimated from responses to one game can be extended to predict or explain responses to different games.

To see if the proposed level-*k* explanation of RTH's results is more than a "just-so" story, Crawford and Iriberri compared it on the overfitting and portability dimensions with the leading alternatives:

- Equilibrium with intuitive payoff perturbations (salience lowers hiders' payoffs, other things equal; while salience raises seekers' payoffs).

- LQRE with similarly intuitive payoff perturbations.

- Alternative level-*k* specifications (for example, with role-asymmetric *L0* or an *L0* that avoids salience, as in Table 3).

Crawford and Iriberri tested for overfitting by re-estimating each model separately for each of RTH's six treatments and using the re-estimated models to "predict" the choice frequencies of the other treatments.

Their favored level-*k* model, with a role-symmetric *L0* that favors salience, has a modest prediction advantage over equilibrium and LQRE with perturbations models, with mean squared prediction error 18% lower and better predictions in 20 of 30 comparisons.

LQRE with payoff perturbations (in different cases) either gets the patterns in the data qualitatively wrong or estimates an infinite precision and thereby turns itself back into an equilibrium model (Crawford and Iriberri's online Appendix).

A more challenging test regards portability.

Crawford and Iriberri tested for portability by using the leading alternative models, estimated from RTH's data, to "predict" subjects' initial responses in the two closest relatives of RTH's games in the literature:

- O'Neill's 1987 *PNAS* famous card-matching game, and
- Rapoport and Boebel's 1992 *GEB* closely related game.

These games both raise the same kinds of strategic issues as RTH's games, but with more complex patterns of wins and losses, different framing, and in the latter case five locations.

I focus here on Crawford and Iriberri's analysis of O'Neill's game.

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.

Discussions of O'Neill's data (which we did not have before we carried out the analysis), 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".

Yet no behaviorally plausible level-*k* model can make a row player play A more than the equilibrium 20%.

Crawford and Iriberri's 2007b, online appendix Tables A3 and A4 show that, excluding *L0*s as normally having 0 estimated frequencies and restricting attention to row players (Player 1s), when A is more salient ($3j - a < 1$) only *L4* chooses A, and that with probability at most 1/3 (Table A3); and 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.)

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%.

Speculating that O’Neill’s subjects’ *initial* responses must not have had an Ace effect, we got the data and found that there was in fact no Ace effect for initial responses.

Instead there was a Joker effect, a full order of magnitude stronger, in which row players played J 56% of the time and column players played it 64% of the time.

Unlike the putative Ace effect, the Joker effect and the other observed frequencies *can* be gracefully explained by a level-*k* model with an *L0* that probabilistically favors the salient A and J cards, in the spirit of Crawford and Iriberri’s analysis of RTH’s data.

Equilibrium or LQRE with payoff perturbations are well-defined for O’Neill’s game, but they both fit significantly worse than our favored level-*k* model.

Our analysis traces the superior portability of the level-*k* model 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 seldom could one extrapolate a specification from one game to another.

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

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	$rj+(1-\varepsilon)[u+v]+(1-r)\,\varepsilon/4$	

We speculated, based on the level-*k* model's success in RTH's and other games, that O'Neill's subjects' *initial* responses must not have had an Ace effect.

There was in fact no Ace effect for initial responses.

Instead there was a Joker effect, a full order of magnitude stronger (but to our knowledge never before mentioned in the literature):

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

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

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

Moreover, unlike the putative Ace effect, the Joker effect and the other observed frequencies *can* be gracefully explained by a level-*k* model with an Obama-McCain *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	Observed or predicted choice frequencies					MSE
		Player	A	2	3	J	
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 $j > 1/4$	1	0.0824	0.1772	0.1772	0.5631	0.0018
	$3j - a < 1, a + 2j < 1$	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 $j > 1/4$	1	0.0000	0.2541	0.2541	0.4919	0.0073
	$3j - a < 1, a + 2j > 1$	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	$a_1 < 1/4, j_1 > 1/4$; $a_2 > 1/4, j_2 < 1/4$	1	0.1804	0.2729	0.2729	0.2739	0.0291
	$3j_1 - a_1 < 1$, $a_1 + 2j_1 < 1, 3a_2 + j_2 > 1$	2	0.1804	0.1804	0.1804	0.4589	0.0117

Crawford and Iriberri's Table 5

Equilibrium or LQRE with perturbations are well-defined for O'Neill's game, but they both fit significantly worse than our favored level-*k* model.

As explained in the paper, equilibrium or LQRE with perturbations are not even well-defined for Rapoport and Boebel's game.

A level-*k* model is well-defined, and explains some but by no means all of the patterns in Rapoport and Boebel's data.

Importantly, Crawford and Iriberri's analysis traces the superior portability of the level-*k* model 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.

End of aside

10. Mr. Schelling Goes to Chicago: Coordination via Payoff Asymmetries and Non-neutral Framing

Perhaps the most famous examples of framing effects in economics are Schelling’s (1960) classic “meeting in New York City” experiments.

Crawford, Gneezy, and Rottenstreich (2008) paired subjects to play games with non-neutral decision labels like Schelling’s, but except for a game with the payoff symmetry of Schelling’s games, they used payoff-asymmetric games like Battle of the Sexes.

In unpaid Chicago pilots, they used naturally occurring labels, pitting the world-famous Sears (now Willis”!) Tower against the little-known AT&T Building across the street.

Framing the game: the Sears Tower, with the AT&T Building on its left

The salience of Sears Tower makes it easy and obvious for subjects to coordinate on the "both-Sears" equilibrium; and almost all do this in the symmetric version of the game.

Since Schelling's experiments with symmetric games, people have assumed that slight payoff asymmetry would not interfere with this.

As we suspected, although the pilot results replicated Schelling's in the symmetric game, there was a substantial decline in coordination with even slight payoff asymmetry.

		P2 (90% Sears)	
		Sears	AT&T
P1 (90% Sears)	Sears	100,100	0,0
	AT&T	0,0	100,100

Symmetric

		P2 (58% Sears)	
		Sears	AT&T
P1 (61% Sears)	Sears	100,101	0,0
	AT&T	0,0	101,100

Slight Asymmetry

		P2 (47% Sears)	
		Sears	AT&T
P1 (50% Sears)	Sears	100,110	0,0
	AT&T	0,0	110,100

Moderate Asymmetry

Chicago Skyscrapers

What’s going on?

Even with slight payoff asymmetry, the game poses a new strategic problem because both-Sears is one player’s favorite way to coordinate but not the other player’s.

Just as in a society of men and women playing Battle of the Sexes, in which Ballet is more salient than Fights, there is a tension between the “label salience” of Sears and the “payoff-salience” of a player’s favorite way to coordinate:

Payoff salience reinforces label salience in one player role (P2s) but opposes it for players in the other (P1s).

This tension may lead players to respond asymmetrically, which in this game is bad for coordination.

To investigate the reasons for the decline in coordination, we conducted more formal, paid treatments using abstract decision labels, pitting X against Y, with X presumed (and shown) to be more salient than Y.

Like the salience of Sears Tower, the salience of the X label makes it obvious for subjects to coordinate on the "both-X" equilibrium; and they again do this in the symmetric version of the game.

But with payoff asymmetry there is again a tension between the "label salience" of X and the "payoff-salience" of a player's favorite way to coordinate:

Payoff salience again reinforces label salience for P2s but opposes it for P1s.

This tension again had a large and surprising effect:

		P2 (76% X)	
		X	Y
P1 (76% X)	X	5,5	0,0
	Y	0,0	5,5

Symmetric

		P2 (28% X)	
		X	Y
P1 (78% X)	X	5,5.1	0,0
	Y	0,0	5.1,5

Slight Asymmetry

		P2 (61% X)	
		X	Y
P1 (33% X)	X	5,6	0,0
	Y	0,0	6,5

Moderate Asymmetry

		P2 (60% X)	
		X	Y
P1 (36% X)	X	5,10	0,0
	Y	0,0	10,5

Large Asymmetry

X-Y Treatments

Even tiny payoff asymmetries caused a large drop in the expected coordination rate, from 64% ($0.64 = 0.76 \times 0.76 + 0.24 \times 0.24$) in the symmetric game to 38%, 46%, and 47% in the asymmetric games.

Perhaps more surprisingly (and unlike in the unpaid Chicago Skyscrapers treatment), the pattern of miscoordination reversed as asymmetric games progressed from small to large payoff differences:

With slightly asymmetric payoffs, most subjects in both roles favored their partners' payoff-salient decisions.

But with moderate or large asymmetries, most subjects in both roles switched to favoring their own payoff-salient decisions.

There are two things to explain here:

Why didn't subjects in the asymmetric games ignore the payoff asymmetry, which cannot be used to break the symmetry as required for coordination, and use the salience of Sears Tower to coordinate?

Why did the pattern of miscoordination reverse as the asymmetric games progressed from small to large payoff differences?

Standard notions such as equilibrium plus noise with refinements and QRE ignore labeling, and so cannot help.

A level-k model can gracefully explain the patterns in the data, but again it's important to have an *L0* that realistically describes people's beliefs about others' instinctive reactions to the tension between label- and payoff- salience that seems to drive the results.

Assume that *L0* is the same in both player roles, and that it responds instinctively to both label and payoff salience; but with a "payoffs bias" that favors payoff over label salience, other things equal:

In symmetric games *L0* chooses X with some probability greater than ½.

And in any asymmetric game, (for simplicity only) whether or not label-salience opposes payoff-salience, *L0* chooses its payoff-salient decision with probability $p > ½$.

(These assumptions are consistent with Crawford and Iriberri's (2007b) assumptions, because their games had no payoff-salience. However, there remain some unresolved issues about how to generalize these assumptions.)

Under these assumptions about *L0*, *L1*'s and *L2*'s choices in roles P1 and P2 are completely determined by *p*, the extent of *L0*'s payoff bias.

Except in symmetric games, even though *L0*'s choice probabilities are the same for P1s and P2s, they imply *L1* and *L2* choice probabilities that differ across player roles due to the asymmetric relationships between label and payoff salience for P1s and P2s.

Simple calculations show that a level-*k* model can track the reversal of the pattern of miscoordination between the slightly asymmetric game and the games with moderate or large payoff asymmetries if (and only if) $0.505 (= 5.1/[5.1+5]) < p < 0.545 (= 6/[6+5])$, so that *L0* has only a modest payoff bias.

If *p* falls into this range and the population frequency of *L1* is 0.7 and that of *L2* is 0.3, close to most previous estimates, the model's predicted choice frequencies differ from the observed frequencies by more than 10% only in the symmetric game, where the model somewhat overstates the homogeneity of our subject pool.

11. Huarangdao and D-Day: Communication of Intentions in Outguessing Games

General Kongming: "Have you forgotten the tactic of 'letting weak points look weak and strong points look strong'?"

General Cao Cao: "Don't you know what the military texts say? 'A show of force is best where you are weak. Where strong, feign weakness.'"

—Luo Guanzhong's historical novel, *Three Kingdoms* (thanks to Duozhe Li)

The quotations refer to a two-person outguessing game with complete information and one-sided preplay communication of intentions via cheap talk.

In the story, set around 200 A.D., fleeing general Cao Cao chose between two escape routes, the easier Main Road and the awful Huarong Road, trying to avoid capture by pursuing General Kongming (http://en.wikipedia.org/wiki/Battle_of_Red_Cliffs).

Kongming (the sender in this example) waited in ambush along the Huarong Road and set campfires there, thus sending a deceptively truthful message.

Cao Cao (the receiver), misjudging Kongming's communication strategy, inverted the truthful message and was caught by Kongming (but was later released).

Consider a simple model of the underlying game (without communication):

		Kongming	
		Main	**Huarong**
Cao Cao	**Main**	**3** **-1**	**0** **1**
	Huarong	**1** **0**	**2** **-2**

Huarongdao

- Cao Cao loses 2 and Kongming gains 2 if Cao Cao is captured.
- But both Cao Cao and Kongming gain 1 by taking the Main Road, whether or not Cao Cao is captured: It’s important to be comfortable, even if (especially if?) if you think you’re about to die.

The key issues here are how Kongming should choose his message and how Cao Cao—knowing Kongming is choosing strategically, trying to anticipate Cao Cao's interpretation—should interpret Kongming's message.

In real settings like this, a receiver's thinking often assigns a prominent role to the literal meanings of messages, without necessarily taking them at face value; and a sender's thinking takes this into account.

But a standard equilibrium analysis precludes a role for the literal meanings of payoff-irrelevant messages (Crawford and Sobel 1982; see however Farrell's 1993 neologism-proofness refinement, which depends on meanings).

Moreover, there is no equilibrium (refined or not) in a zero-sum (or this) two-person game in which cheap talk conveys information or the receiver responds to the message.

In such an equilibrium, if there was information in the sender's message that the receiver found it optimal to respond to, the receiver's response would help him and so hurt the sender, who would then prefer to make his message uninformative.

Thus communication can have no effect in any equilibrium, and as a result the underlying game must be played according to its unique mixed-strategy equilibrium, as if there were no communication phase.

Yet intuition suggests that in many such situations:

- The sender's message and action are part of a single, integrated strategy.
- The sender tries to anticipate which message will fool the receiver and chooses it nonrandomly.
- The sender's action differs from what he would have chosen with no opportunity to send a message.

Huarongdao is only one datapoint (if that).

But there's another example in which essentially the same thing happened.

Consider the Allies' choice of where to invade Europe on D-Day (6 June 1944), the motivating example of Crawford (2003).

The underlying game can also be modeled as an outguessing game:

		Germans	
		Defend Calais	**Defend Normandy**
Allies	**Attack Calais**	**1** **-1**	**-2** **2**
	Attack Normandy	**-1** **1**	**1** **-1**

- Attacking an undefended Calais is better for the Allies than attacking an undefended Normandy, so better for them on average.
- Defending an unattacked Normandy is worse for the Germans than defending an unattacked Calais and so worse for them on average.

Now imagine that D-Day is preceded by a message from the Allies to the Germans regarding their intentions about where to attack, as in Operation Fortitude South (http://en.wikipedia.org/wiki/Operation_Fortitude).

Imagine further that the message is (approximately!) cheap talk.

A "Tank" from Operation Fortitude

In what sense did essentially the same thing happen in both Huarongdao and D-Day?

In each case the deception succeeded, but the sender won in the less beneficial of the two possible ways to win.

Kongming's message was literally truthful—he lit fires on the Huarong Road and ambushed Cao Cao there—but Cao Cao was fooled because he misread Kongming’s message strategy and inverted the message.

In D-Day the message was literally deceptive but the Germans were fooled because they “believed” it (either because they were credulous or because they inverted the message one too many times).

The sender’s and receiver’s message strategies and beliefs were different, but the outcome—what happened in the underlying game—was the same: The sender won, but in the less beneficial of the two possible ways.

The quotations tell us why Cao Cao was fooled by Kongming’s message: with a truthful *L0*, his rationale resembles *L1* thinking, while Kongming’s resembles *L2* thinking.

But the quotations don’t tell us why in each case the sender won only in the less beneficial way.

To restate the puzzle more concretely, for both D-Day and Huarongdao:

- Why did the receiver allow himself to be fooled by a costless (hence easily faked) message from an *enemy*?

- If the sender expected his message to fool the receiver, why didn't he reverse it and fool the receiver in the way that would have allowed him to win in the *more* beneficial way?

  (Why didn't the Allies feint at Normandy and attack at Calais? Why didn't Kongming light fires and ambush Cao Cao on the Main Road?)

A level-*k* analysis suggests that it was more than a coincidence that the same thing happened in both cases.

Although *Sophisticated* subjects are rare in laboratory experiments, one hopes they are more common in field settings; and it is interesting to see whether a plausible model allows deception between *Sophisticated* players.

Accordingly, let Allies' and Germans' types be drawn from separate distributions, each including both level-*k* or *Mortal* types (as Crawford 2003 called them, in honor of Puck, in *A Midsummer Night's Dream*, Act 3: "Lord, what fools these mortals be!") and a fully strategically rational or *Sophisticated*, type.

Mortal types use step-by-step procedures that generically determine unique pure strategies, and avoid simultaneous determination of the kind used to define equilibrium; recall the Selten (1998) quote above.

Sophisticated types know everything about the game, including the distribution of *Mortal* types; and so play an equilibrium in a "reduced game" between possible *Sophisticated* players, taking *Mortals*' choices as given.

How should *L0* be adapted to an extensive-form game with communication?

Here a uniform random *L0* seems quite unnatural. For sender or receiver, the instinctive reaction to a message in a language one understands is surely to focus on its literal meaning, even if one ends up either lying or not taking the message at face value.

The level-*k* model therefore anchors *Mortal* types' messages and responses on *L0*s based on truthfulness for senders and credulity for receivers, just as in the informal literature on deception.

(The literature has not yet converged on whether *L0* receivers should be defined as credulous or uniform random—compare Ellingsen and Östling (2010)—but the distinction is partly semantic because truthful *L0* senders imply that *L1* receivers are also credulous.)

Mortal Allied types' simplified models of other players make *L1* or higher *Mortal* Allied types always expect to fool the Germans, either by lying (like the Allies) or by telling the truth (like Kongming).

Given this, all *L1* or higher *Mortal* Allied types send a message they expect to make the Germans think they will attack Normandy, and then attack Calais.

If we knew the Allies and Germans were *Mortal*, we could now derive the model's implications from an estimate of the type frequencies of *Mortal* Allies who tell the truth or lie, and of *Mortal* Germans who believe or invert the Allies' message.

But the analysis must also take into account the possibility of *Sophisticated* Allies and Germans, who know everything about the game, including the distribution of *Mortal* types, and play an equilibrium in the resulting game.

To take into account the possibility of *Sophisticated* Allies and Germans, note that *Mortals* players' strategies are determined independently of each other's and *Sophisticated* players' strategies, and so can be treated as exogenous (even though they affect other players' payoffs).

Plug in the distributions of *Mortal* Allies' and Germans' independently determined behaviors to obtain a "reduced game" between *Sophisticated* Allies and *Sophisticated* Germans.

Because *Sophisticated* players' payoffs are influenced by *Mortal* players' decisions, the reduced game is no longer zero-sum, its messages are not cheap talk, and it has incomplete information.

The sender's message, ostensibly about his intentions, is in fact read by a *Sophisticated* receiver as a signal of the sender's type.

Thus, the possibility of *Mortal* players completely changes the character of the game between *Sophisticated* players, which is what gives the model the ability to explain the effectiveness of communication in a zero-sum game and the possibility of deception between *Sophisticated* players.

The equilibria of the reduced game are determined by the population frequencies of *Mortal* and *Sophisticated* senders and receivers.

There are two leading cases, with different implications:

- When *Sophisticated* Allies and Germans are common—not behaviorally plausible—the reduced game has a mixed-strategy equilibrium whose outcome is virtually equivalent to D-Day's without communication.

- When *Sophisticated* Allies and Germans are rare, the game has an essentially unique pure-strategy equilibrium, in which *Sophisticated* Allies can predict *Sophisticated* Germans' decisions, and vice versa.

In the latter, pure-strategy equilibrium, *Sophisticated* Germans always defend Calais (because they know that *Mortal* Allies, who predominate when *Sophisticated* Allies are rare, will always attack Calais).

Sophisticated Allies send the message that fools the most likely kind of *Mortal* German (feinting at Calais or Normandy depending on whether more *Mortal* Germans believe than invert messages), and then attack Normandy.

There is never a pure-strategy equilibrium in which *Sophisticated* Germans defend Normandy while *Sophisticated* Allies attack Calais.

In such an equilibrium any deviation from *Sophisticated* Allies' equilibrium message would "prove" to *Sophisticated* Germans that the Allies were *Mortal*, making it optimal for *Sophisticated* Germans to defend Calais.

If in the equilibrium *Sophisticated* Allies attacked Normandy while *Sophisticated* Germans defended Normandy, then the Allies' message would fool only the most likely kind of *Mortal* German—*Sophisticated* Germans are never fooled in a pure-strategy equilibrium, and a message cannot fool both *Mortal* Germans who believe and those who invert messages—with expected payoff gain equal to the frequency of the most likely kind of *Mortal* German times the payoff of attacking undefended Normandy.

But such *Sophisticated* Allies could reverse both their message and attack location, again fooling the most likely kind of *Mortal* German, but now with expected payoff gain equal to the frequency of that kind of German times the higher payoff of attacking undefended Calais, a contradiction. Thus in any pure-strategy equilibrium, *Sophisticated* Germans must defend Calais no matter what message they hear.

Thus, in a sense (resting on a preference for pure-strategy equilibria), the model explains why *Sophisticated* Germans don't defend Normandy while *Sophisticated* Allies attack Calais, even though this would have been more profitable for the Allies if it succeeded.

In the pure-strategy equilibrium that exists when *Sophisticated* Allies and Germans are rare, the Allies' message and action are part of a single, integrated strategy; and the probability of attacking Normandy is much higher than if no communication was possible.

The Allies choose their message nonrandomly, the deception succeeds most of the time, but it allows the Allies to win in the less beneficial way.

Thus for plausible parameter values, with no unexplained difference in the sophistication of Allies and Germans, the model explains why *Sophisticated* Germans might allow themselves to be "fooled" by a costless message from a *Sophisticated* enemy: It is an unavoidable cost of exploiting mistakes by *Mortal* enemies, who are much more common.

Nonetheless, *Sophisticated* players in either role do strictly better than their *Mortal* counterparts; their advantage comes from the ability to avoid being fooled and/or to choose which *Mortal* type(s) to fool.

In the mixed-strategy equilibrium that prevails when *Sophisticated* Allies and Germans are common, *Sophisticated* players' equilibrium mixed strategies offset each other's gains from fooling *Mortal* Receivers, and in each role *Sophisticated* and *Mortal* players have equal expected payoffs.

This suggests that in an adaptive analysis of the dynamics of the type distribution, as in Conlisk 2001 *AER*, the frequencies of *Sophisticated* types will grow until the population is in or near (depending on costs) the region of mixed-strategy equilibria in which types' expected payoffs are equal.

Thus *Sophisticated* and *Mortal* players can coexist in long-run equilibrium.

12. Alphonse and Gaston: Communication of Intentions in Coordination Games

"After you, Alphonse." "No, you first, my dear Gaston!"

—Frederick B. Opper's comic strip, *Alphonse and Gaston* (http://en.wikipedia.org/wiki/Alphonse_and_Gaston)

"What we got here…is a failure to communicate."

—Paul Newman as the title character in Cool Hand Luke

(http://www.imdb.com/title/tt0061512/)

If level-*k* models allow preplay communication of intentions to affect the outcomes of zero-sum games, it should come as no surprise that they also allow effective communication in coordination games.

Ellingsen and Östling (2010) and Crawford (2007), not discussed in detail here, adapt Crawford's (2003) approach to study different aspects of preplay communication of intentions in coordination and other games.

Ellingsen and Östling use a level-*k* model to study the effectiveness of a single round of one- or two-sided preplay communication in games where communication of intentions plays various roles.

Crawford uses a level-*k* model to study the effectiveness of one- or multi-round two-sided communication in games like Battle of the Sexes, building on Farrell's 1987 *RAND J* and Rabin's 1994 *JET* analyses.

In each case the power of the analysis stems from the use of a model that does not assume equilibrium, which is question-begging in this context; but which imposes a realistic structure less agnostic than rationalizability.

13. October Surprise: Communication of Private Information in Outguessing Games

"…The news that day was the so-called 'October Surprise' broadcast by bin Laden. He hadn't shown himself in nearly a year, but now, four days before the [2004 presidential] election, his spectral presence echoed into every American home. It was a surprisingly complete statement by the al Qaeda leader about his motivations, his actions, and his view of the current American landscape. He praised Allah and, through most of the eighteen minutes, attacked Bush,…. At the end, he managed to be dismissive of Kerry, but it was an afterthought in his 'anyone but Bush' treatise….

Inside the CIA…the analysis moved on a different [than the presidential candidates' public] track. They had spent years, as had a similar bin Laden unit at FBI, parsing each expressed word of the al Qaeda leader…. What they'd learned over nearly a decade is that bin Laden speaks only for strategic reasons…. Today's conclusion: bin Laden's message was clearly designed to help the President's reelection."

—Suskind, *The One Percent Doctrine*, 2006, pp. 335-6 (quoted in Jazayerli 2008 http://www.fivethirtyeight.com/2008/10/guest-column-will-bin-laden-strike.html).

A. October Surprise

The quotation refers to a zero-sum two-person game with incomplete information and one-sided preplay communication of private information via cheap talk.

Only bin Laden knows which candidate he wants; and, talk being cheap, he will say what it takes to help his candidate win. A representative American voter knows only that he wants whichever candidate bin Laden doesn’t want.

A representative American voter knows only that he wants the opposite of what bin Laden wants.

The key issues are how bin Laden should relate his statement to what he really wants and how the American should interpret bin Laden’s statement, knowing that bin Laden is choosing the message strategically.

Once again, the literal meanings of messages are likely to play a prominent role in applications, but equilibrium analysis precludes such a role.

There is again no equilibrium in which cheap talk conveys information, or in which the receiver responds to the sender’s message.

Consider, however, a level-*k* model in which *L0* is anchored on truthfulness for the sender (bin Laden) and credulity for the receiver (American).

(Or one could derive credulity for an *L1* receiver and start from there.)

An *L0* or *L1* American believes bin Laden's message, and therefore votes for whichever candidate bin Laden attacks.

An *L0* bin Laden who wants Bush to win attacks Kerry, but an *L1* (*L2*) bin Laden who wants Bush to win attacks Bush to induce *L0* (*L1*) Americans to vote for Bush.

Given bin Laden's attack on Bush, an *L0* or *L1* American ends up voting for Bush, and an *L2* American ends up voting for Kerry.

Bin Laden's message is always influential, but he needs to choose it to fool the most prevalent kind of American—believer or inverter—as in Crawford's (2003) analysis.

An *L2* bin Laden believes Americans are *L1*, so "reverse psychology" will be effective.

B. Experimental Evidence: Wang, Spezio, and Camerer (2010)

I now discuss some experimental evidence on communication of private information in discretized versions of Crawford and Sobel's (1982) sender-receiver games.

Sender observes state S = 1, 2, 3, 4, or 5, sends message M = 1, 2, 3, 4, or 5. Receiver observes message, chooses action A = 1, 2, 3, 4, or 5.

The Receiver's choice of A determines the welfare of both:

- The Receiver's ideal outcome is A = S.

- The Sender's ideal outcome is A = S + b.

The Receiver's von Neumann-Morgenstern utility function is $110 - 20|S - A|^{1.4}$, and the Sender's is $110 - 20|S + b - A|^{1.4}$.

The difference in preferences varied across treatments: b = 0, 1, or 2.

Crawford and Sobel's theoretical analysis characterized the possible equilibrium relationships between Sender's observed S and Receiver's choice of A, which determines the informativeness of communication.

They showed, for a class of models with continuous state and action spaces that generalizes Wang et al.'s examples (except for discreteness), that all equilibria are "partition equilibria", in which as illustrated below, the Sender partitions the set of states into contiguous groups and tells the Receiver, in effect, only which group his observation lies in.

For any given difference in Sender's and Receiver's preferences (b), there is a range of equilibria, from a "babbling" equilibrium with one partition element to more informative equilibria that exist when b is small enough.

Under reasonable assumptions there is a "most informative" equilibrium, which has the most partition elements and gives the Receiver the highest ex ante (before the Sender observes the state) expected payoff.

As the preference difference decreases, the amount of information transmitted in the most informative equilibrium increases (measured either by the correlation between S and A or the Receiver's expected payoff).

The unambiguous part of Crawford and Sobel's characterization of equilibrium concerns the possible relationships between S and A.

Because messages have no direct effect on payoffs ("cheap talk"), there is nothing to tie down their meanings in equilibrium.

As a result, any equilibrium relationship between S and A can be supported by any sufficiently rich language, with the meanings of messages determined by players' equilibrium beliefs.

Behaviorally, however, in experiments like Wang et al.'s with a clear correspondence between state and message—S = 1, 2, 3, 4, or 5 and M = 1, 2, 3, 4, or 5—or where communication is in a common natural language, the interpretations of messages are dictated by their literal meanings.

Thus messages are always understood—even if not always believed.

Wang et al.'s data analysis therefore fixes the meanings of Sender subjects' messages at their literal values.

Even with this restriction, when b = 0 or 1 in their design (Sender's and Receiver's preferences are close enough) there are multiple equilibria.

Wang et al.'s analysis then focuses on the "most informative" equilibrium.

When b = 0, the most informative equilibrium has M = S and A = S: perfect truth-telling, credulity, and information transmission, as is intuitively plausible when Sender and Receiver have identical preferences.

When b = 2, the most informative equilibrium has Senders sending a completely uninformative message M = {1, 2, 3, 4, 5} for any value of S; and Receivers ignoring it, hence choosing A = 3, which is optimal given their prior beliefs, for any value of M.

(A babbling equilibrium also exists when b = 0 or 1, but then it is not the most informative equilibrium.)

When b = 1, the most informative equilibrium has Senders sending M = 1 when S = 1 but M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and Receivers choosing A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}.

(The Sender's message M = {2, 3, 4, 5} is the simplest way to implement the intentional vagueness of this partition equilibrium. Another way would be for the Sender to randomize M uniformly on {2, 3, 4, 5} when S = 1.)

Thus, when b = 1 the difference in preferences causes noisy information transmission even in the most informative equilibrium.

Importantly, however, the Receiver's beliefs on hearing the Sender's message M are necessarily an unbiased—though noisy—estimate of S:

In equilibrium there is no lying or deception, only intentional vagueness.

(When b = 1, there's another, more informative equilibrium, found by David Eil of UCSD, in which Senders send M = {1, 2} when S = 1 or 2 but M = {3, 4, 5} when S = 3, 4, or 5; and Receivers choose A = 2 when M = {1, 2} and A = 4 when M = {3, 4, 5}. But this equilibrium is not "robust", in that Senders who observe S = 2 are indifferent between M = {1, 2} and M = {3, 4, 5}.)

Turning to Wang et al.'s results, when b = 0 Senders almost always set M = S and Receivers almost always set A = M: The result is near the perfect information transmission predicted by the most informative equilibrium.

Figure 1 shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S: A circle's size shows the Sender's message frequencies. A circle's darkness and the poorly visible numbers inside show the Receiver's action frequencies.

Figure 1: Raw Data Pie Charts (b=0)
(Hidden Bias-Stranger)

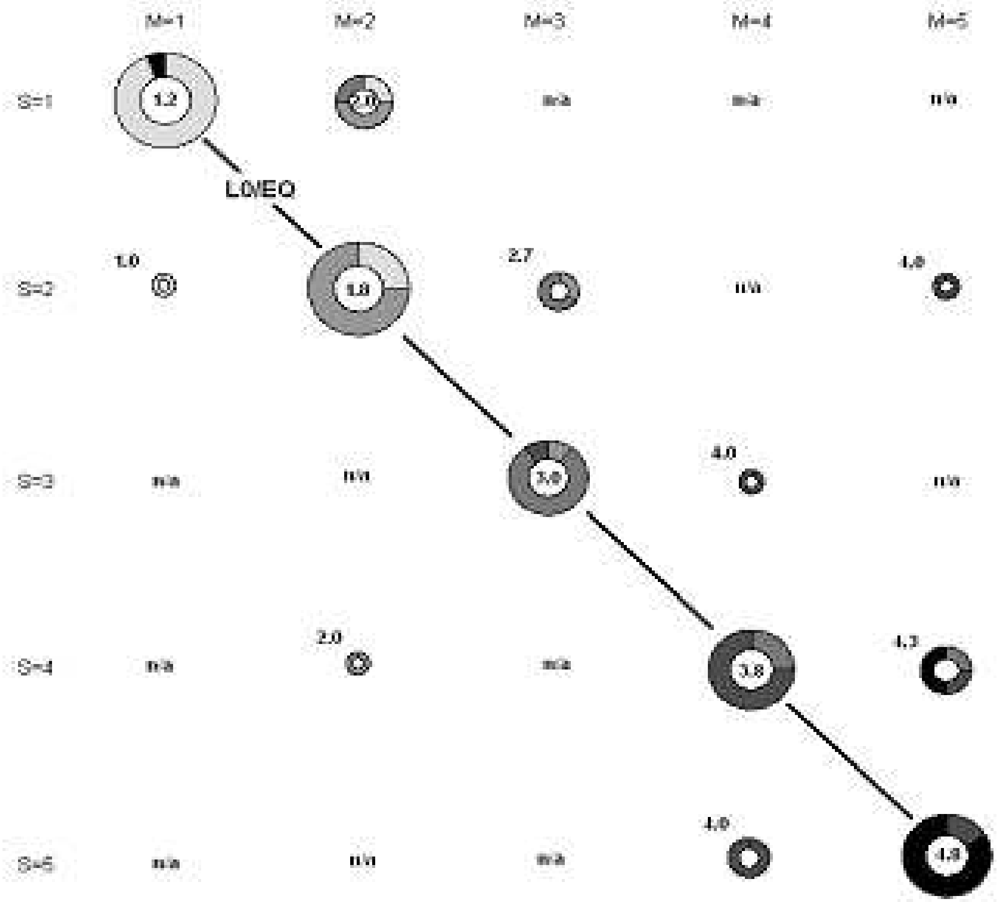

As b increases to b = 1 or b = 2, the amount of information transmitted decreases as predicted by Crawford and Sobel's equilibrium comparative statics, but there are also systematic deviations from the most informative (or any) equilibrium, and lying and successful deception occur.

In Figure 3 (next slide; b = 2 omitted from Wang et al.'s label by accident), in the essentially unique, most informative equilibrium M = {1, 2, 3, 4, 5}, so equilibrium message distributions would look the same for all five rows; and equilibrium actions would be concentrated on A = 3.

However, although the observed actions are fairly close to A = 3, message distributions shift rightward as S increases (going down in the table); thus:

● Most Senders exaggerate the truth (most messages above the diagonal), apparently trying to move Receivers from Receivers' ideal action A = S toward Senders' ideal action A = S + 2 (or 5, whichever is smaller).

● Even so, there is some information in Senders' messages (message distributions shift rightward going down in the table, so messages are positively correlated with the state).

● Receivers are usually deceived to some extent (average A usually > S).

Figure 3: Raw Data Pie Chart, (Hidden Bias-Stranger)

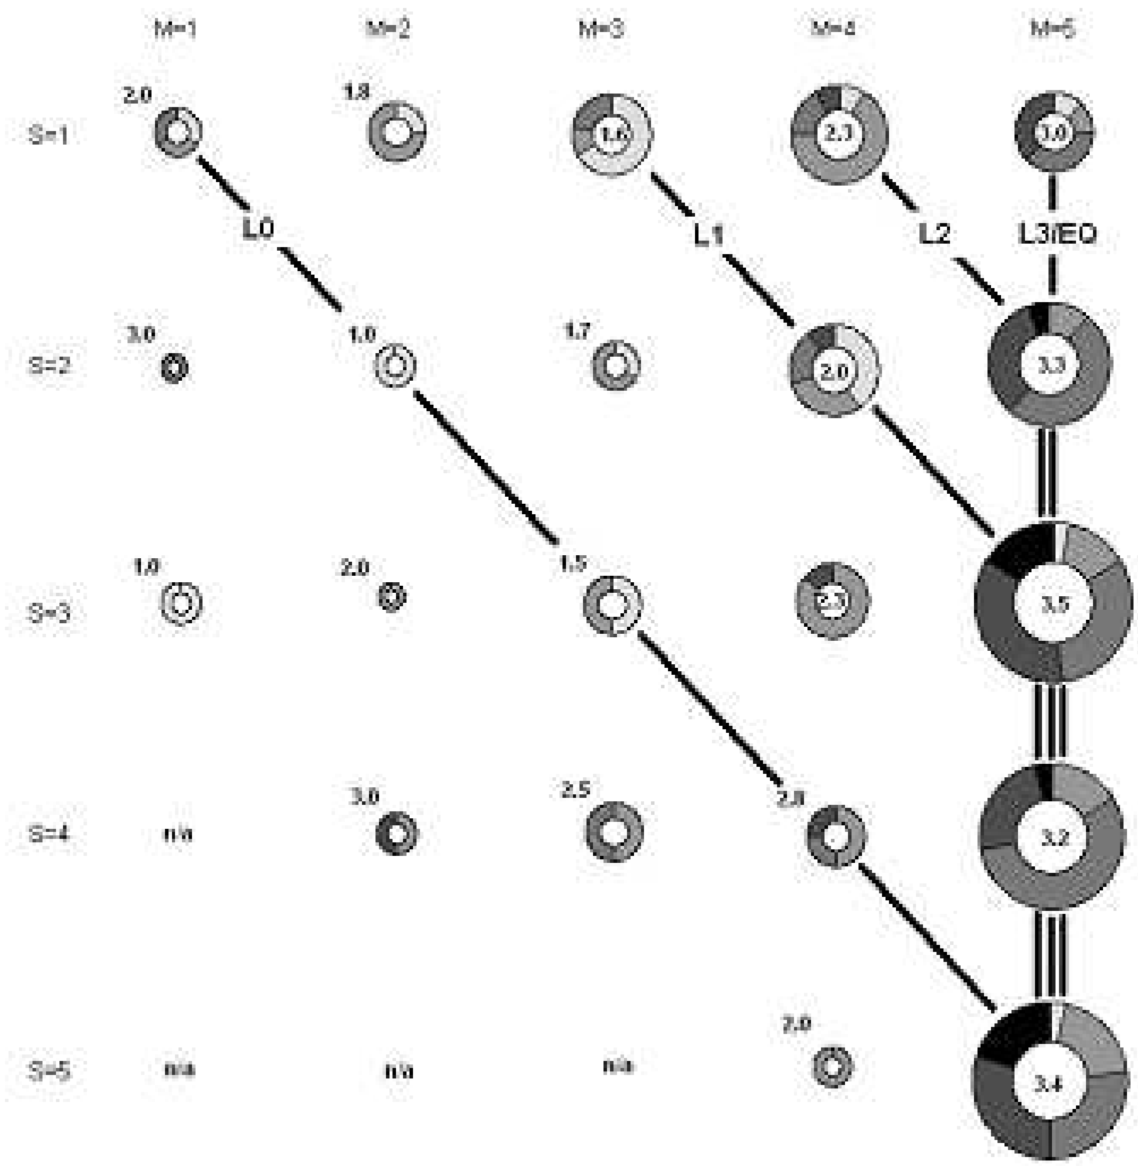

When b = 1, in the most informative robust equilibrium, the Sender's message is M = 1 when S = 1 and M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and the Receiver chooses A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}. Thus, in equilibrium the distributions of messages and actions would be the same for S = 2, 3, 4, or 5.

By contrast, turning to Figure 2 (b = 1; next slide):

- Senders almost always exaggerate the truth (messages above the diagonal), apparently trying to move Receivers from Receivers' ideal action A = S toward Senders' ideal action A = S + 1.

- Even so, there is some information in Senders' messages (message distributions shift rightward going down in the table, so messages are positively correlated with the state).

- Receivers are usually deceived to some extent (average A usually > S).

Figure 2: Raw Data Pie Chart (b=1)
(Hidden Bias-Stranger)

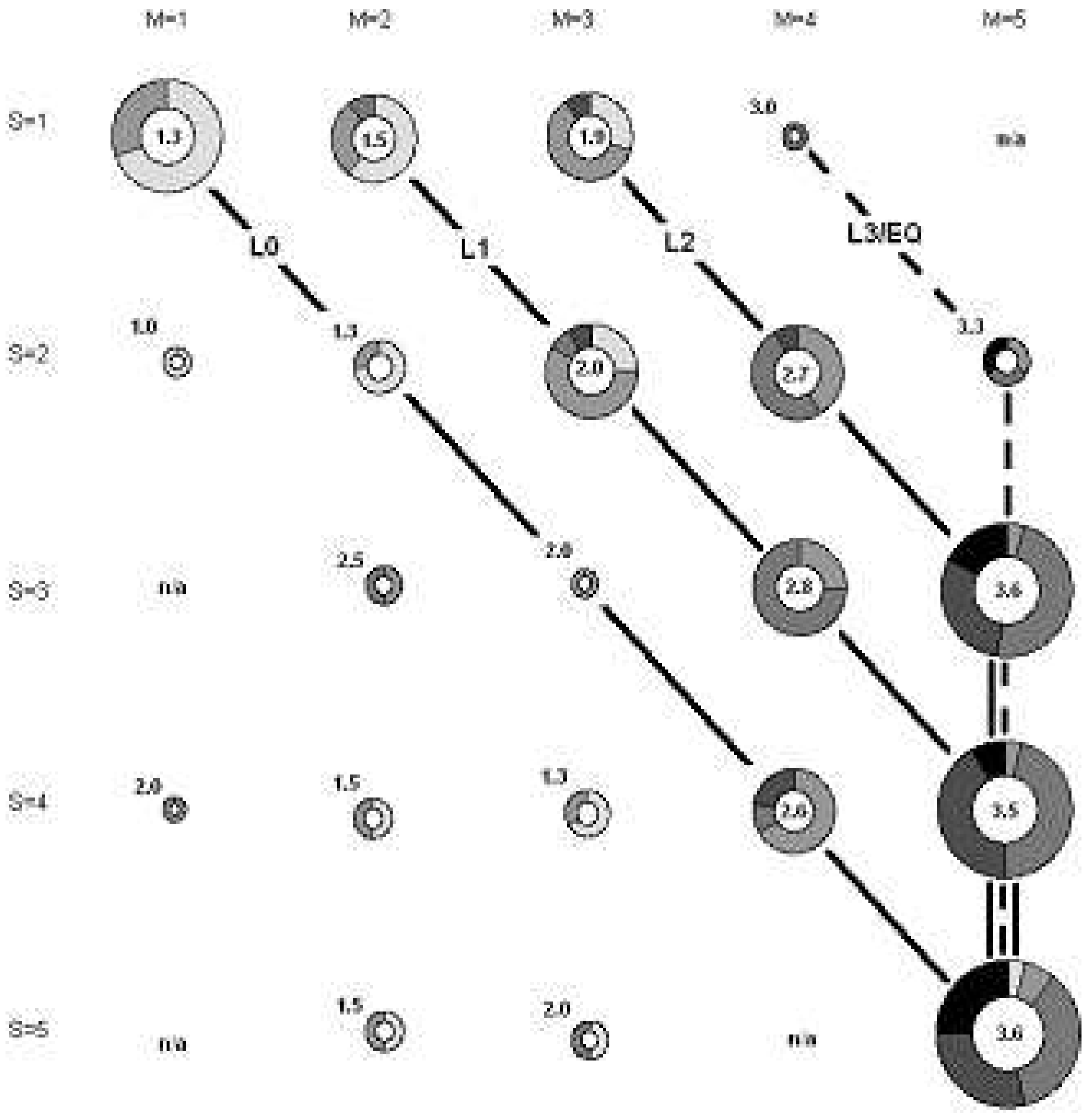

What kind of model can explain results like this? Wang et al., following Cai and Wang (2006), propose a level-*k* explanation based on Crawford's (2003) analysis of preplay communication of intentions:

Anchor beliefs in a truthful Sender *L0*, which sets M = S; and a credulous Receiver *L0* (which also best responds to an *L0* Sender), setting A = M.

L1 Senders best respond to *L0* Receivers by inflating their messages by b: M = S + b (up to M = 5), so that *L0* Receivers will choose S + b, yielding the Sender's ideal action given S.

L1 Receivers (as defined by Wang et al.; the numbering is a convention) best respond to *L1* Senders by discounting the message, normally setting A = M – b, yielding Receivers' ideal action given M = S + b of S.

The qualification "normally" reflects Wang et al.'s assumption that *L1* Receivers take into account that when b = 2, *L1* senders with S = 3, 4, or 5 all send M = 5, with the result that *L1* Receivers, knowing that S is equally likely to be 3, 4, or 5, choose A = 4 instead of A = M – 2b = 3.

L2 Senders best respond to *L1* Receivers by inflating their messages by 2b: M = S + 2b (up to M = 5), so that *L1* Receivers will set A = M – b = S + b, yielding Senders' ideal action given S.

L2 Receivers best respond to *L2* Senders by discounting the message, normally setting A = M – 2b, yielding Receivers' ideal action given M = S + 2b of S.

The qualification "normally" reflects Wang et al.'s assumption that *L2* Receivers take into account that when b = 1, *L2* senders with S = 3, 4, or 5 all send M = 5, with the result that *L2* Receivers, knowing that S is equally likely to be 3, 4, or 5, choose A = 4 instead of A = M – 2b = 3.

L2 Receivers also take into account that when b = 2, *L2* senders with S = 2, 3, 4, or 5 send M = 5, with the result that *L2* Receivers, knowing that S is equally likely to be 2, 3, 4, or 5, choose A = 4 instead of A = M – 2b = 3.

Figure 2: Raw Data Pie Chart (b=1)
(Hidden Bias-Stranger)

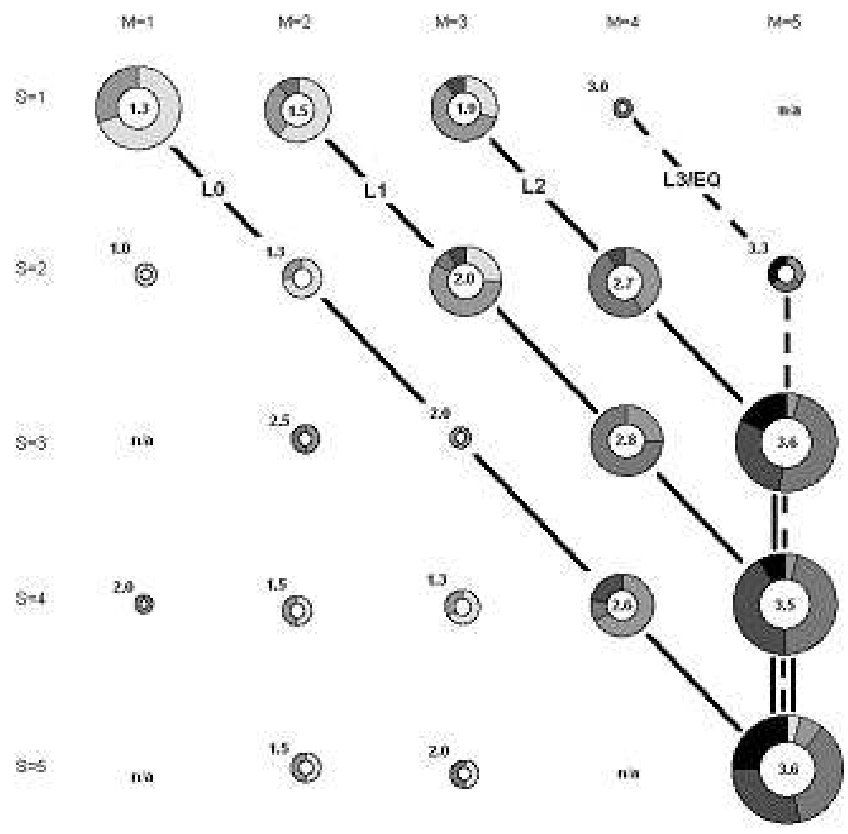

Figure 3: Raw Data Pie Chart, (Hidden Bias-Stranger)

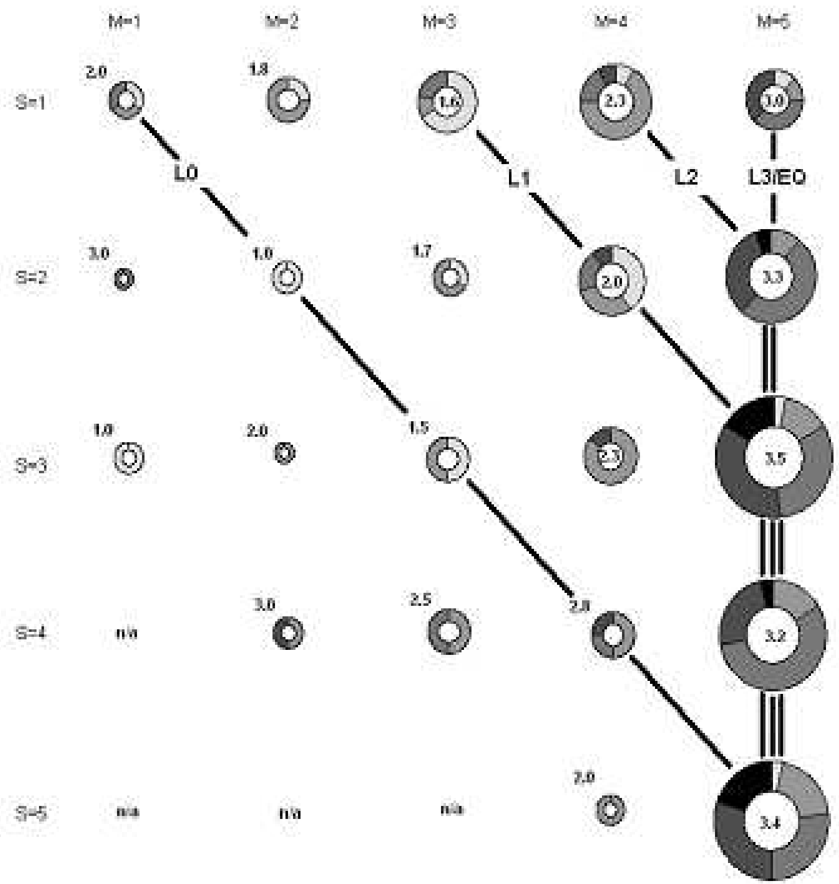

Note that when b = 1, *L1*, *L2*, and *Eq* all predict M = 5 when S = 4 or 5; and when b = 2, *L1*, *L2*, and *Eq* all predict M = 5 when S = 3, 4, or 5.

Econometric estimation classifies 18% of 16 Sender subjects as *L0*, 25% *L1*, 25% *L2*, 14% *Sophisticated*, and 18% *Equilibrium* (note different type definitions).