

Minicourse on “Level-k Thinking”

Monday, 5 October 2008, National Taiwan University

Tuesday, 6 October 2008, Academia Sinica

Vincent P. Crawford, University of California, San Diego

based mostly on joint work with

**Miguel A. Costa-Gomes, University of Aberdeen
and**

Nagore Iriberry, Universitat Pompeu Fabra

**(Our papers cited below are available at
[http://dss.ucsd.edu/~vcrawfor/.](http://dss.ucsd.edu/~vcrawfor/))**

Introduction

In many applications of game theory, there are ample opportunities for learning from experience in previous play of analogous games.

In such settings, the cognitive requirements for learning to converge to equilibrium are mild, and experiments suggest that people do in fact converge to equilibrium (with some qualifications for mixed-strategy equilibria or extensive-form games).

If only long-run outcomes matter and equilibrium selection does not depend on the details of how people learn from experience, such applications can rely entirely on equilibrium.

There is then no need for deeper understanding of strategic thinking.

Many other applications of game theory involve games played without clear precedents in which initial outcomes matter.

Such applications, which include questions involving comparative statics or mechanism design, depend on predicting initial responses to games, even if eventual convergence to equilibrium is assured.

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

Such applications also depend on predicting initial responses, and may also depend on understanding the structure of learning rules.

Consider for example a “Continental Divide” coordination game from Van Huyck, Cook, and Battalio’s 1997 *JEBO* experiment.

Seven subjects choose simultaneously and anonymously among “efforts” from 1 to 14, with each subject’s payoff determined by his own effort and a summary statistic, the median, of all players’ efforts.

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

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

In the table as displayed here (but not as displayed to subjects), the payoffs of a player’s best responses to each possible median are highlighted in bold; and the payoffs of the (symmetric, pure-strategy) equilibria “all–3” and “all–12” are highlighted in large bold.

Continental divide game payoffs

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

There were ten sessions, each with its own separate group. Half the groups had an initial median of eight or above, and half had an initial median of seven or below. (I suspect that the experimenters chose the design to make this happen, but it's not uncommon.)

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

The median-seven-or-below groups converged almost perfectly to the all-3 equilibrium.

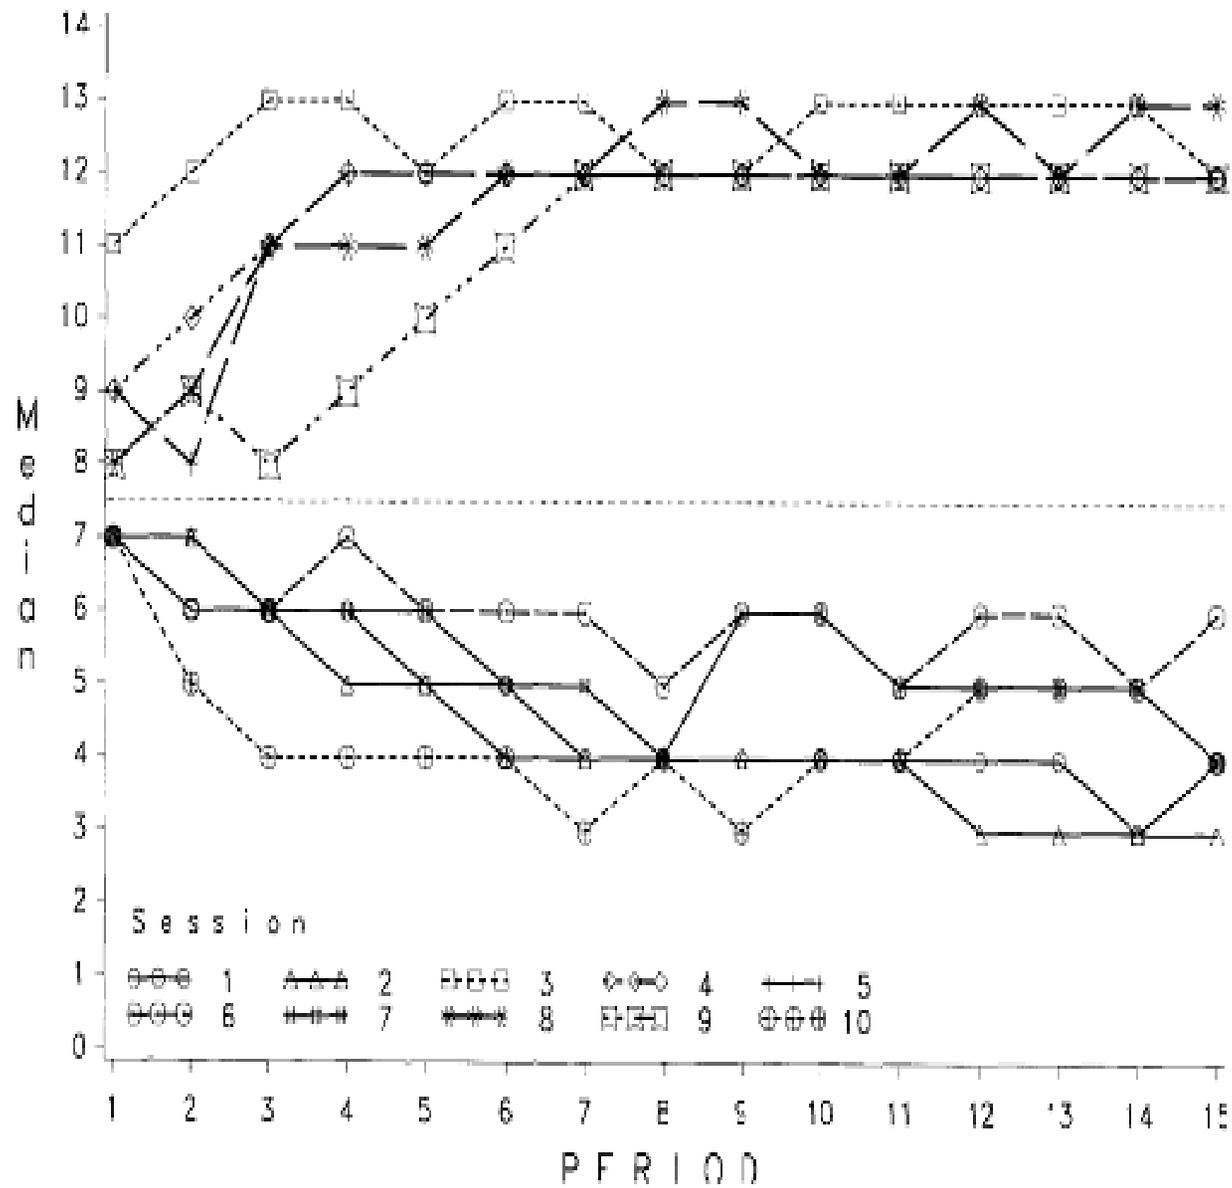


Fig. 3. Median choice in sessions 1 to 10 by period

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

Here we also need to know a subject group's median initial response.

That and a simple view of learning, in which equilibrium selection is determined by which basin of attraction—defined by myopic best responses—subjects' initial responses fell into, seem to determine final outcomes in Van Huyck, Cook, and Battalio's experiment.

Thus, because Van Huyck, Cook, and Battalio's subjects had no prior experience, their initial responses were almost entirely the product of strategic thinking, the focus of this minicourse.

In other applications we need to know more about both initial responses and the structure of subjects' learning rules: See for example Crawford, "Adaptive Dynamics in Coordination Games," 1995 *Econometrica*, and Crawford and Broseta, "What Price Coordination? The Efficiency-enhancing Effect of Auctioning the Right to Play," 1998 *AER*, which discuss Van Huyck, Battalio, and Beil's 1990 *AER*, 1991 *QJE*, and 1993 *GEB* experiments.

What we learn about strategic thinking by studying initial responses will also affect our views about how best to model learning, although the cognitive underpinnings of learning are less well understood.

For more on learning, see chapter 6 in Camerer's 2003 book, *Behavioral Game Theory*, but the literature is evolving rapidly.

Modeling Strategic Thinking in Initial Responses to Games

The cognitive requirements for initial responses to be in equilibrium are more stringent than those for learning to converge to equilibrium:

- Players must have perfectly coordinated beliefs, which without precedents requires perfect models of each other's decisions.

It is easy to imagine strategic thinking being this accurate in simple games, but the thinking required for equilibrium initial responses in complex games is behaviorally far-fetched (see for example Harsanyi and Selten's *A General Theory of Equilibrium Selection in Games*).

Even players who are capable of such thinking may doubt that others are capable of it, or doubt that others believe that others are capable.

Moreover, there is a growing body of laboratory evidence that initial responses often deviate systematically from equilibrium, especially when it requires thinking that is not straightforward.

Modeling initial responses more accurately promises several benefits:

- It can establish the robustness of conclusions based on equilibrium in games where boundedly rational rules mimic equilibrium (for example, Costa-Gomes, Crawford, and Broseta 2001 *Econometrica* or Crawford and Iriberri 2007 *Econometrica*).
- It can challenge the conclusions of applications to games where equilibrium is implausible without learning.
- It can resolve empirical puzzles by explaining the deviations from equilibrium some games evoke (Crawford 2003 *AER*; Camerer, Ho, and Chong 2004 *QJE*; Crawford and Iriberri 2007 *AER*).
- It can yield insights that elucidate the structure of learning, where assumptions about cognition determine which analogies between current and previous games players recognize and distinguish reinforcement from beliefs-based and more sophisticated rules.

Overview of Alternative Models of Initial Responses

A variety of models have been proposed to describe experimental subjects' initial responses to games, which normally allow players' responses to be in equilibrium, but do not assume it:

- Adding noise to equilibrium predictions (“equilibrium plus noise”).
- McKelvey and Palfrey’s 1995 *GEB* notion of quantal response equilibrium (“QRE”) and a leading case, logit QRE (“LQRE”).
- The level- k models of Nagel 1995 *AER*; Stahl and Wilson 1995 *GEB*; Ho, Camerer, and Weigelt 1998 *AER*; Costa-Gomes, Crawford, and Broseta 2001 *Econometrica* (“CGCB”); and Costa-Gomes and Crawford 2006 *AER* (“CGC”).
- Camerer, Ho, and Chong’s 2004 *QJE* (“CHC”) closely related cognitive hierarchy (“CH”) model.
- Goeree and Holt’s 2004 *GEB* model of noisy introspection (“NI”).

I now briefly discuss those models, as background for the more detailed discussion of level- k and CH models that follows. Costa-Gomes, Crawford, and Iriberri 2009 *JEEA* gives more detail.

1. Equilibrium plus noise

Equilibrium plus noise adds noise with a specified distribution (usually logit) and an estimated precision parameter to equilibrium predictions.

It sometimes does well, but it often misses systematic patterns in subjects' deviations from equilibrium, which tend to be sensitive to out-of-equilibrium payoffs in patterns that it cannot account for.

And in games with multiple equilibria, where most or all feasible decisions may be part of some symmetric pure-strategy equilibrium, equilibrium plus noise is incomplete in that it does not specify a unique (though possibly probabilistic) prediction conditional on the value of its behavioral parameters (in this case, the precision).

Such multiplicity of predictions has previously been dealt with by estimating an unrestricted probability distribution over equilibria (for example, Bresnahan and Reiss 1991 *Journal of Econometrics*), but the freedom this allows can make the model badly overfit.

To put equilibrium plus noise on a more equal footing with the other models, which are complete in the above sense, and to guard against overfitting, it seems appropriate to add a coordination refinement such as Harsanyi and Selten's risk-dominance or payoff-dominance.

2. Quantal response equilibrium (QRE) and Logit QRE (LQRE)

To capture the payoff-sensitivity of deviations from equilibrium, McKelvey and Palfrey 1995 *GEB* proposed the notion of QRE.

In a QRE players' decisions are noisy, with the probability density of each decision increasing in its expected payoff, evaluated taking the noisiness of others' decisions into account.

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

As the distributions' precision increases, QRE approaches equilibrium; and as their precision approaches zero, QRE approaches uniform randomization over players' feasible decisions.

A QRE model is closed by specifying a response distribution, which is logit in almost all applications.

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

The distributional assumptions are crucial: with an unrestricted distribution QRE can “explain” any given dataset (Haile, Hortacsu, and Kovenock 2008 *AER*). The use of the logit distribution has been guided more by fit and custom than by independent evidence.

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

With estimated precision, LQRE often fits subjects’ initial responses better than an equilibrium plus noise model.

From the point of view of describing strategic thinking, LQRE's fit comes at a cost:

- Players must not only respond to a probability distribution of other players' responses, but also find a generalized equilibrium that is a fixed point in a large space of response distributions: If equilibrium reasoning is cognitively taxing, LQRE reasoning is doubly taxing.
- The mathematical complexity of LQRE means it must almost always be solved for computationally and is not easily adapted to analysis.
- In some settings LQRE fits worse than equilibrium (Camerer, Palfrey, and Rogers 2007; Chong, Camerer, and Ho 2005; Crawford and Iriberri 2007 *Econometrica*), sometimes making systematic *qualitative* errors in predicting deviations from equilibrium (Crawford and Iriberri 2007 *AER*; Östling, Wang, Chou, and Camerer 2008).

3. Level- k and cognitive hierarchy (CH) models

Motivated by these considerations and experimental evidence, some recent work treats 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 structures seems daunting, much of the experimental evidence supports a particular class of models called level- k or cognitive hierarchy (“CH”) models.

Costa-Gomes and Crawford (2006 *AER*, Introduction and Section II.D) summarize the experimental evidence for level- k models, which now includes experiments on games with a wide variety of structures:

- Stahl and Wilson 1994 *JEBO*, 1995 *GEB*
- Nagel 1995 *AER*
- Ho, Camerer, and Weigelt 1998 *AER* (HCW)
- Costa-Gomes, Crawford, and Broseta 2001 *Econometrica* (“CGCB”)
- Costa-Gomes and Crawford 2006 *AER* (“CGC”)
- Cai and Wang 2006 *GEB*
- Wang, Spezio, and Camerer 2007
- Costa-Gomes and Weizsäcker 2008 *Review of Economic Studies*
- Kawagoe and Takizawa 2008 *GEB*

The closely related CH model was introduced and used to analyze a partly overlapping body of experimental data by:

- Camerer, Ho, and Chong 2004 *QJE* (“CHC”)

Level- k models

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

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

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

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

Lk 's only departure from equilibrium is in replacing its assumed perfect model of others' decisions with simplified models that avoid the complexity of equilibrium analysis. Compare Selten (1998 *EER*):

“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.”

$L1$ and higher types make undominated decisions, and in many games Lk complies with k rounds of iterated dominance, so that its decisions are k -rationalizable.

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

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

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

Thus, $L0$ “exists” mainly as $L1$'s model of others, $L2$'s model of $L1$'s model of others, and so on.

Even so, the specification of $L0$ is the main issue in defining a level- k model and the key to its explanatory power.

$L0$ needs to be adapted to the setting as illustrated below, but the definition of higher types via iterated best responses allows a simple, reliable explanation of behavior across different settings.

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

But unlike LQRE, a level- k model requires neither that players respond to nondegenerate distributions of others' responses (except for $L1$'s response to $L0$, whose uniform randomness is simple to respond to) nor that they find fixed points.

Further, a level- k model's point predictions do not depend on estimated precisions, only on the estimated type frequencies.

Their simple recursive structure 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.

Cognitive hierarchy (CH) models

In Camerer et al.'s closely related CH model, type L_k best responds not to L_{k-1} alone but to a mixture of lower-level types, and the type frequencies are treated as a parameterized Poisson distribution.

Although this specification seems more natural than the simpler level- k specification for an outside observer modeling subjects' behavior econometrically, which specification better describes people's behavior is an empirical question (on which the jury is still out).

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

Estimating an unconstrained type distribution 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—the explanation is not credible.

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

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

In a CH model, as in a level- k model, players need not respond to the noisiness of others' decisions (except $L0$'s) or find fixed points, but they do respond to a nondegenerate distribution of lower types' responses, in proportions determined by an estimated parameter.

Like a level- k model, a CH model makes point predictions that do not depend on estimated precisions, only on the Poisson parameter.

A CH model also has a recursive structure, albeit somewhat more complex one than a level- k model's structure.

Level- k and CH models

Like equilibrium plus noise and LQRE, level- k and CH models are applicable to “any” game and have small numbers of parameters.

Like equilibrium plus noise and LQRE, level- k and CH models are general models of strategic behavior, but they are complements to equilibrium plus noise, not competitors.

In a sense, level- k and CH models are competitors to LQRE, which replace LQRE’s payoff-sensitive errors, and players’ responses to them, with a structural model of players’ deviations from equilibrium.

In many games L_k complies with k rounds of iterated dominance.

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

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

These features allow level- k (and CH) models to establish the robustness of conclusions based on equilibrium in games where boundedly rational rules mimic equilibrium, and to capture the sensitivity of deviations from equilibrium to out-of-equilibrium payoffs.

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

4. 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 *GEB* suggest using LQRE for both initial responses and limiting outcomes, with increasing precision as a reduced-form model of learning.

But Goeree and Holt 2004 *GEB* (“GH”) suggest using LQRE for limiting outcomes, instead proposing a Noisy Introspection (“NI”) model to describe initial responses.

GH’s NI model relaxes LQRE’s equilibrium assumption while maintaining its assumption that players respond to a nondegenerate probability distribution of other players’ responses: Instead players form beliefs by iterating best responses as in a level- k model, but higher-order beliefs reflect increasing amounts of noise.

For a given noise distribution, an NI model 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 extremes mimic level- k types: If the noise jumps immediately to ∞ , NI beliefs are $L0$; if it is zero for one iteration and then jumps to ∞ , NI beliefs are $L1$, and so on.
- In applications GH assume that the noisiness of higher-order beliefs grows geometrically with iterations, which yields beliefs similar but not identical to Lk 's; slower noise growth is like higher k .

The resulting NI model is more flexible than LQRE, and cognitively less taxing because it does not require fixed-point reasoning.

But such an NI model is cognitively 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 in applications GH truncate the iteration to ten rounds).

NI's structure, like LQRE's, is not directly grounded in evidence.

In fact the evidence from Nagel's 1995 *AER* and subsequent experiments suggests that the indefinite iteration of best responses and the assumed homogeneity of strategic thinking are unrealistic.

Experimental Evidence for level- k /CH models

There have been a number of econometric horse races among subsets of the equilibrium plus noise, QRE and LQRE, level- k and CH, and NI models just described, which generally tend to favor level- k /CH models but are not conclusive.

See for example Costa-Gomes, Crawford, and Iriberri 2009 *JEEA* and the references cited there.

However, as noted above, level- k /CH models have an advantage over the alternatives in that they are *directly* linked to experimental evidence (rather than just doing well in data-fitting exercises).

I now give the flavor of some of the experimental evidence.

The flavor of the evidence on which level- k /CH models are based is illustrated by Nagel's 1995 *AER* experimental results for n -person guessing games:

- 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 who haven't already done so, 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 guess 0.

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

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

- It's dominated to guess more than 50 ($1/2 \times 100 \leq 50$).
- Unless you think other people will make dominated guesses, it's also dominated to guess more than 25 ($1/2 \times 50 \leq 25$).
- 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 the arguments for equilibrium in the other examples, because it depends “only” on iterated knowledge of rationality, not on players having the same beliefs.

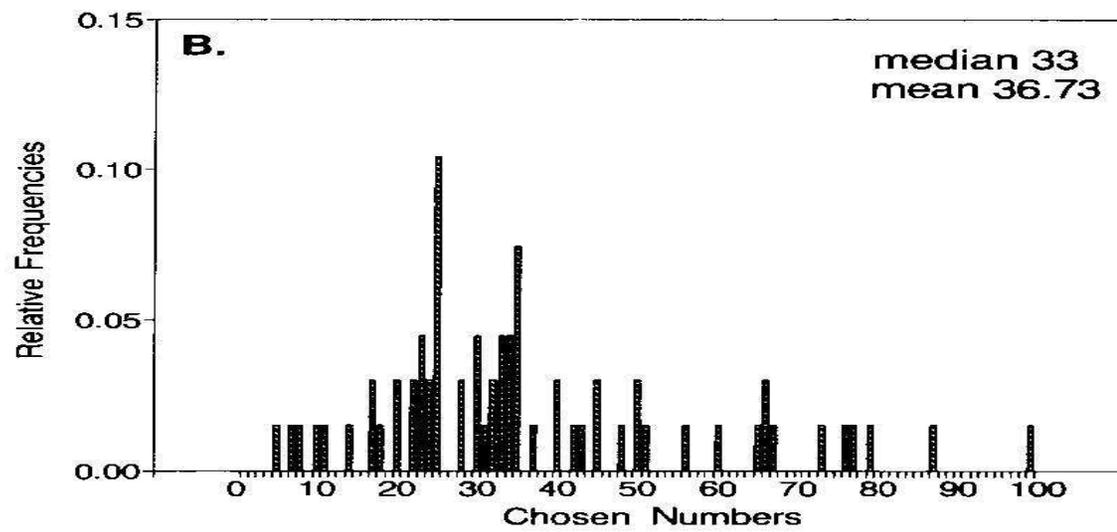
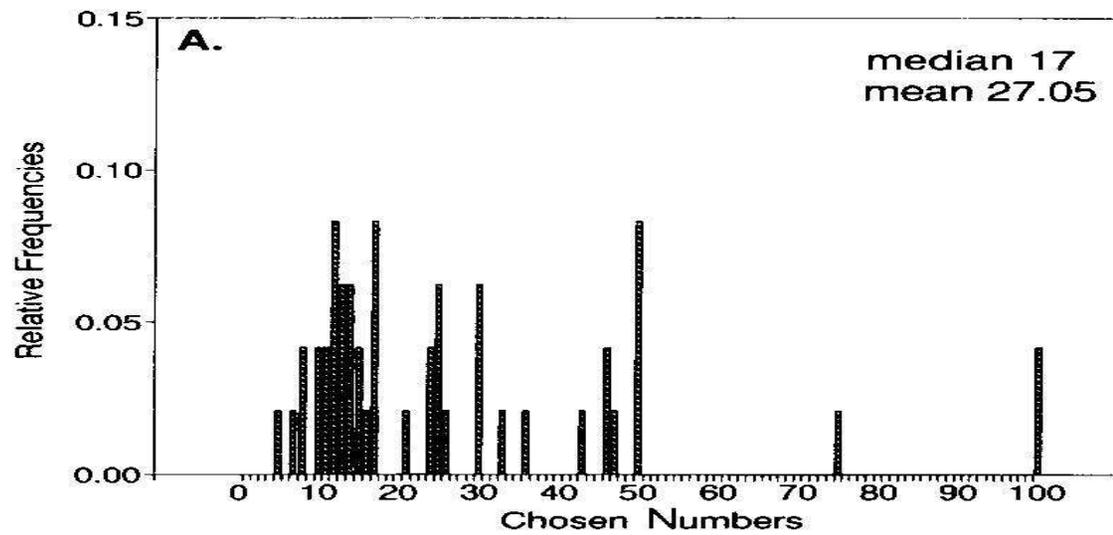
However, even people who are rational themselves are seldom certain that others are rational, or that others believe that they themselves are rational, and so on.

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

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

Her subjects never played their equilibrium strategies initially, and their responses resembled neither equilibrium plus noise nor QRE (for any reasonable distribution; but recall Haile et al. 2008 *AER*).

Instead there were spikes that suggest a discrete, heterogeneous distribution of strategic thinking "types," respecting 0 to 3 rounds of iterated dominance (first picture $p = 1/2$; second $p = 2/3$):



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

- In one, which we 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 others' remaining strategies.
- In another, which we call "level- k " or " Lk ," a player starts with a naïve prior, $L0$, over others' possible guesses and then iterates the best response mapping k times, with $k = 1, 2$, or perhaps 3 .

In games like Nagel's, $L0$ is usually taken to be uniform random, and for simplicity I will take this to refer to the average of others' guesses.

(In other n -person games, whether $L0$ is taken to be correlated or independent is important, and the limited evidence that is available suggests that people act as if they had a single, perfectly correlated model of others. But here we mostly consider two-person games.)

Dk starts with iterated knowledge of rationality and then invokes a naïve prior; by standard measures its cognitive requirements are close to *Lk+1*'s, and both respond similarly to dominance.

In Nagel's $[0, 100]$ games with $p < 1$, *Dk*'s and *Lk+1*'s guesses are perfectly confounded:

- *Dk* guesses $([0+100p^k]/2)p$ and *Lk+1* guesses $[(0+100)/2]p^{k+1}$; both track the spikes in her data.

Despite the lack of separation, many theorists interpret Nagel's results as evidence that subjects explicitly performed iterated dominance, the way we teach students to solve such games.

In HCW's 1998 *AER* and CGCB's 2001 *Econometrica* experiments, *Dk*'s and *Lk+1*'s guesses are weakly separated, and the results are inconclusive. But in CGC's 2006 *AER* experiments their guesses are strongly separated, and the results clearly favor *Lk* over *Dk* rules.

Experimental evidence for level- k /CH models continued

Camerer (*Behavioral Game Theory*, Chapter 5), CHC (Section IV), and CGC 2006 *AER* (Introduction, Section II.D) summarize the experimental evidence for level- k /CH models in games with a variety of structures.

Here I give the flavor of the evidence by summarizing CGC's results, which are fully consistent with the results of previous experiments that elicit initial responses to games, but more precise.

CGC's 2006 *AER* Design

CGC's experiments randomly and anonymously paired subjects to play series of 2-person guessing games, with no feedback.

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

The goal was to focus on how players model others' decisions by studying strategic thinking “uncontaminated” by learning.

(“Eureka!” learning was possible, but it can be tested for and is rare.)

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

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

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 previous guessing experiments, the targets and limits were always the same for both players, and varied at most across treatments.)

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, so the equilibrium is determined by players' upper limits. (When the product is < 1 , the equilibrium is determined by players' lower limits in a similar way.)

In equilibrium the first player guesses his upper limit of 500, but the second player guesses 750, below his upper limit of 900.

More generally, the games have essentially unique equilibria 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 stress-tests equilibrium, which responds much more strongly to the product of the targets than alternative rules do, and enhances the separation of equilibrium from alternative rules.

(It also reveals other interesting patterns, not discussed here; see Crawford, "Look-ups as the Windows of the Strategic Soul: Studying Cognition via Information Search in Game Experiments" at <http://dss.ucsd.edu/~vcrawfor/#Search>.)

Most of the evidence from normal-form games suggests defining $L0$ as uniformly random over the feasible range of decisions.

In addition to *Equilibrium* and the level- k types $L1$, $L2$, and $L3$ defined for a random $L0$, CGC's data analysis considered two "iterated dominance" types:

- $D1$, which does one round of dominance and then best responds to a uniform prior over its partner's remaining decisions
- $D2$, which does two rounds and then best responds to a uniform prior over its partner's remaining decisions

CGC also considered:

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

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

	<i>L1</i>	<i>L2</i>	<i>L3</i>	<i>D1</i>	<i>D2</i>	<i>Eq.</i>	<i>Soph.</i>
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

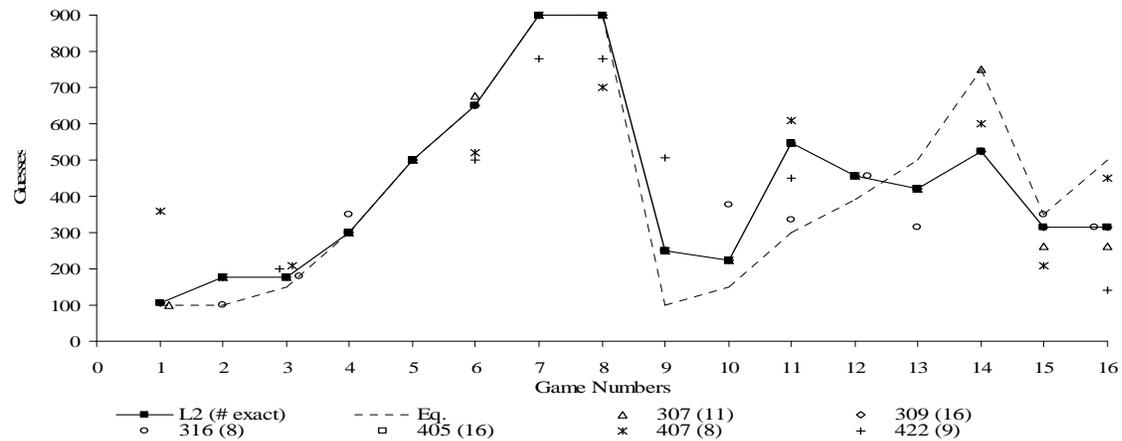
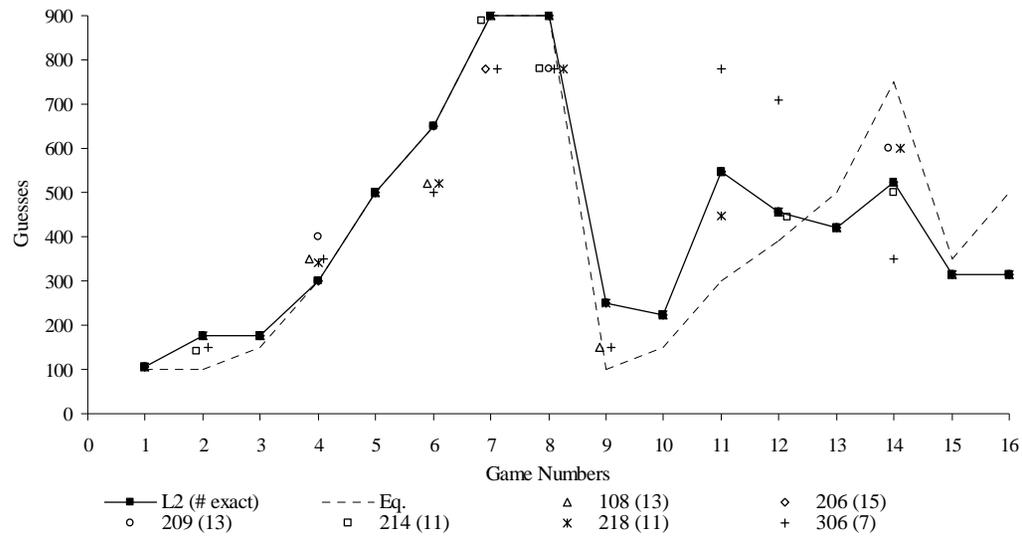
Of the 88 subjects in CGC's main treatments, 43 made guesses that complied *exactly* (within 0.5) with one type's guesses in from 7 to 16 of the games (20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*).

For example, CGC's Figure 2 shows the "fingerprints" of the 12 subjects whose guesses conformed most closely to *L2*'s; 72% of these guesses were exact; only the deviations 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 very powerful evidence for the type whose guesses are tracked.

If a subject chooses 525, 650, 900 in games 1-3, we "know" he's *L2*.

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



CGC's Figure 2. "Fingerprints" of 12 Apparent L2 Subjects

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

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

Type	Apparent from guesses	Econometric from guesses	Econometric from guesses, excluding random	Econometric from guesses, with specification test	Econometric from guesses and search, with specification test
<i>L1</i>	20	43	37	27	29
<i>L2</i>	12	20	20	17	14
<i>L3</i>	3	3	3	1	1
<i>D1</i>	0	5	3	1	0
<i>D2</i>	0	0	0	0	0
<i>Eq.</i>	8	14	13	11	10
<i>Soph.</i>	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 Table 1

CGC's analysis suggests several conclusions:

- There are no *Dk* subjects. (Subjects respect iterated dominance to the extent that *Lk* types do, not because they explicitly perform it.)
- There are no *Sophisticated* subjects.
- Thus there are only *L1*, *L2*, *L3*, and *Equilibrium* subjects (however the subjects who guessed closest to *Equilibrium* appear to be following hybrid types that only mimic equilibrium in some games).
- By a quirk of CGC's design (p. 1763), the data are neutral on the level-*k* versus CH definitions of *L2* and *L3*.
- But the data appear inconsistent with CH's assumed Poisson type distribution, which given the estimated frequencies of *L1*, *L2*, and *L3* subjects would imply many more *L0* subjects than CGC found.
- Equilibrium plus noise and LQRE miss clear patterns in the data: "errors" are structural or cognitive, with little payoff-sensitivity.

CGC's Specification test

CGC's data analysis is based on an a priori list of types, which were chosen for behavioral plausibility and consistency with previous work.

For the 43 of 88 subjects with high rates of exact compliance with some type's guesses, the data are pretty conclusive anyway: Fits this good can only happen to the extent that the types are well-specified.

(Even so, doubts remain about the subjects with high exact compliance with *Equilibrium*; see Crawford, "Look-ups as the Windows of the Strategic Soul: Studying Cognition via Information Search in Game Experiments".)

But for the 45 of 48 subjects whose types must be econometrically estimated, there is 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.

With regard to overfitting, 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 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 for concreteness *L2*.

The pseudotypes of subjects now estimated to be *L2* would then outperform the non-*L2* types estimated for them, and would also make approximately the same (*L2*) 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 small clusters involving 11 of the 88 subjects, and the subjects in these clusters were also left unclassified in Table 1.

The paper and web appendix discuss what the subjects in each cluster seemed to be doing; most of it appears idiosyncratic, hence reasonable to treat as part of the error term in a simple model.

Because a cluster must contain at least two subjects, it is reasonable to anticipate finding more than the five CGC found in a larger sample.

But because any such clusters did not reach the two-subject threshold in CGC's sample of 88, they probably make up at most about 2% of any larger sample.

Clusters that small are reasonably be treated as errors, on the grounds that extending the theory to encompass them isn't worth it.

On this basis, CGC concluded that:

- A level- k model explains a large fraction of the part of subjects' deviations from equilibrium that can be explained by a model.
- Although the model explains only half or a bit more of subjects' deviations from equilibrium, it may still be optimal for a modeler to treat the rest of the deviations as errors.

Illustration from M. M. Kaye's *The Far Pavilions*

I now give an example that illustrates how one might use a level- k model in applications.

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

Both have a single, *strategically simultaneous* choice between North and South—that is, their choices are time-sequenced, but the Pursuers must make their choice irrevocably before they learn Ash's choice.

If the pursuers catch Ash, they gain 2 and he loses 2.

But South is warm, and North is the Himalayas with winter coming.

Thus both Ash and the Pursuers gain an extra 1 for choosing South, whether or not Ash is caught:

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

Escape!

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

Escape!

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

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

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

As in other perturbed matching pennies games, this equilibrium is intuitive for the Pursuers, but not for Ash.

Equilibrium does not reflect this intuition, but experimental data from such games suggest that people's decisions do reflect it, with average p s above the analog of 1/4 and sometimes average q s above 3/4.

Back in the novel, Ash overcomes his intuition and goes North. The Pursuers unimaginatively go South, so Ash escapes...and the novel can continue...romantically...for 900 more pages.

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

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

Types	Ash	Pursuers
<i>L0</i>	uniformly random	uniformly random
<i>L1</i>	South	South
<i>L2</i>	North	South
<i>L3</i>	North	North
<i>L4</i>	South	North
<i>L5</i>	South	South

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

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

How do we know Ash's type?

One advantage of using fiction as data is that the narrative sometimes reveals cognition as well as decisions:

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

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

Koda Dad's advice reflects the belief that the Pursuers think Ash is $L1$, so that Ash will go south because it's “kinder” and that (assuming the Pursuers are $L0$) the Pursuers are no more likely to pursue him there.

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

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

$L3$ ties my record-high k for a clearly explained level- k type in fiction.

Poe's *The Purloined Letter*

(<http://xroads.virginia.edu/%7EHYPER/POE/purloine.html>) has another $L3$, but Conan Doyle doesn't even have an $L1$!

I offer a US\$100 reward for the first *clearly explained* $L4$ or higher in fiction or non-game-theoretic nonfiction.

I suspect even postmodern fiction may have no L s higher than $L3$, because they wouldn't be credible.

But I would be delighted to pay off for a true $L4$, postmodern or not.

In applications we don't usually have an author identifying characters' types for us.

But if the game is clearly defined, and we have data on people's decisions, we can specify a level- k model and derive its implications, and then use them to estimate the frequency distribution of types.

Alternatively, we can calibrate the model using previous estimates.

If the model includes errors, we can also estimate or calibrate the types' precision parameters.

In games like Escape!, even though Lk types best responding to nonequilibrium beliefs typically don't randomize, the estimated type distribution implies a mixture of decisions for each role that reflects players' strategic uncertainty.

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

But the model, like the data, tends to deviate from equilibrium in the direction intuition suggests:

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

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

Applications

I now give several applications to illustrate the use of level- k models to resolve empirical/experimental puzzles about strategic behavior that do not seem to be gracefully resolved by equilibrium analysis:

- Camerer, Ho, and Chong's 2004 *QJE* analysis of “magical” coordination in market-entry games, simplified here to Battle of the Sexes games as in Crawford 2007, “Let's Talk it Over...”.
- Crawford and Iriberri's 2007 *AER* analysis of systematic deviations from unique mixed-strategy equilibrium in zero-sum two-person hide-and-seek games with non-neutrally framed locations.
- Crawford, Gneezy, and Rottenstreich's 2008 *AER* analysis of coordination via Schelling-style focal points.
- Crawford's 2003 *AER* analysis of preplay communication of intentions in zero-sum two-person games.
- Crawford and Iriberri's 2007 *Econometrica* analysis of overbidding in independent-private-value and common-value auctions.

Other interesting applications not discussed here include:

- Camerer, Ho, and Chong's 2004 *QJE* CH analyses of speculation and zero-sum betting and of money illusion.
- Cai and Wang's 2006 *GEB*; Wang, Spezio, and Camerer's 2007; and Kawagoe and Takizawa's 2008 *GEB* level-*k* experimental analyses of communicating private information.
- Ellingsen and Östling's 2008 level-*k* analysis of one-round preplay communication of intentions in coordination and other games.
- Crawford's 2007 (<http://dss.ucsd.edu/~vcrawfor/#Talk>) level-*k* analysis of one- and multi-round preplay communication of intentions in coordination games.
- Goldfarb and Yang's 2008 *J. Marketing Research* CH analysis of field data on technology adoption decisions (like entry games).
- Östling, Wang, Chou, and Camerer's 2008 CH analysis of field and lab data on Poisson LUPI (lowest unique positive integer games).
- Crawford, Kugler, Neeman, and Pauzner's 2009 *JEEA* level-*k* analysis of optimal auction design.

The applications suggest that the definition of Lk in terms of iterated best responses yields models that “work” in a variety of settings.

But they also suggest that LO must often be adapted to the setting.

The type frequencies must also sometimes be adapted to the setting.

The flexibility of LO may raise particular doubts because its specification seems to give the modeler a lot of freedom.

But because LO reflects players’ models of others’ instinctive reactions to the game, it is a decision-theoretic or “psychological” concept.

As the applications will illustrate, there is often enough intuition and evidence to specify a “psychological” LO in new applications.

If LO were an even partly strategic concept, there would be much less intuition and evidence, and the models would be less portable.

“Magical” coordination in Battle of the Sexes and market-entry Games (CHC 2004 *QJE* and Crawford 2007)

In market-entry experiments, a number of 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 the market capacity; but a given negative profit if too many enter.

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

The natural equilibrium benchmark prediction is the symmetric mixed-strategy equilibrium, in which each player enters with a given probability that makes all 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 regularly have better ex post coordination (number of entrants closer to market capacity) than in the symmetric equilibrium.

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

(But no one would be at all surprised by this unless he believed in equilibrium, so it would only look like magic to a game theorist.)

Camerer, Ho, and Chong's 2004 *QJE*, Section III.C, analysis shows that Kahneman's magic can be explained by a level- k model. I now do a similar level- k analysis in a simplified two-person market-entry game with capacity one, which is like Battle of the Sexes.

		In	Out
In	0	0	1
Out	1	a	0

$a > 1$ Market Entry

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

The expected coordination rate is $2p(1 - p) = 2a/(1+a)^2$.

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

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

I assume for simplicity that the frequency of $L0$ is 0, and that $L0$ chooses its action randomly, with $\Pr\{\text{In}\} = \Pr\{\text{Out}\} = 1/2$.

$L1$ s mentally simulate $L0$ s' random decisions and best respond, choosing In; $L2$ s choose Out; $L3$ s choose In; and $L4$ s choose Out.

		In	Out
In	0	0	1
Out	1	a	0

$a > 1$ **Market Entry**

Types	$L1$	$L2$	$L3$	$L4$
$L1$	In, In	In, Out	In, In	In, Out
$L2$	Out, In	Out, Out	Out, In	Out, Out
$L3$	In, In	In, Out	In, In	In, Out
$L4$	Out, In	Out, Out	Out, In	Out, Out

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

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

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

This is maximized at $v = \frac{1}{2}$ where it takes the value $\frac{1}{2}$.

Thus for v near $\frac{1}{2}$, which is plausible, the coordination rate is close to $\frac{1}{2}$. (For more extreme values the rate is worse, $\rightarrow 0$ as $v \rightarrow 0$ or 1 .)

By contrast, the mixed-strategy equilibrium coordination rate, $\frac{2a}{(1 + a)^2}$, is maximized when $a = 1$, where it takes the value $\frac{1}{2}$.

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

The level- k model yields a completely different view of coordination than a traditional refined-equilibrium model:

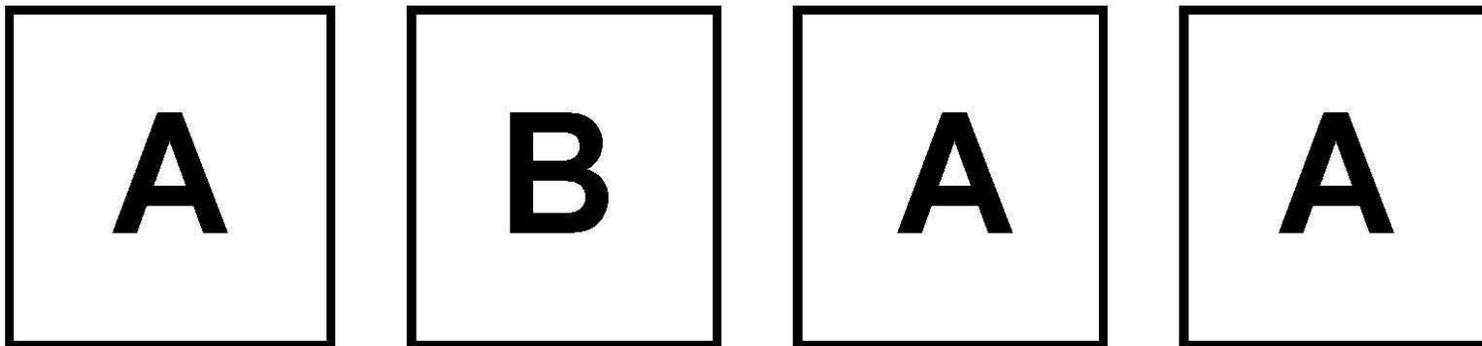
- Equilibrium and selection principles like risk- or payoff-dominance play no role whatsoever in players' strategic thinking.
- Coordination, when it occurs, is an accidental (though statistically predictable) by-product of players' non-equilibrium decision rules.
- Even though decisions are simultaneous and there is no communication or observation of the other's decision, the predictable heterogeneity of strategic thinking allows more sophisticated players such as $L2$ s to mentally simulate the decisions of less sophisticated players such as $L1$ s and accommodate them, just as Stackelberg followers would.
- This mental simulation doesn't work perfectly, because an $L2$ doesn't know the other's type. Neither would it work if strategic thinking were homogeneous. But it is surprising that it works at all.

Role-Asymmetric Deviations from Mixed-Strategy Equilibrium in Hide-and-Seek Games with Non-neutral Framing of Locations (Crawford and Iriberri 2007 *AER*)

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.

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 uniqueness is analogous to the uniqueness of “*B*”.

However, the analysis will show that its psychological effects are quite different.

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

Similar landscapes are common in "folk game theory":

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

(comment on the poisoning of Ukrainian presidential candidate—now president—Viktor Yushchenko)

(The meeting with government officials is analogous to RTH's B, but there's nothing in this example analogous to the end locations.)

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

(Garrison Keillor (1997) on multiple-choice tests)

(With four possible choices arrayed left to right, this example is very close to RTH's design.)

Hide-and-seek has a clear equilibrium prediction, which leaves no room for framing to systematically influence the outcome.

Because it's a zero-sum two-person game, the arguments for playing equilibrium strategies are stronger than usual.

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

(Although in this game any strategy, pure or mixed, is a best response to equilibrium beliefs, systematic deviations of aggregate choice frequencies from equilibrium probabilities must (with high probability) have a cause that is partly common across players, and are therefore indicative of systematic deviations from equilibrium.)

TABLE 1—AGGREGATE CHOICE FREQUENCIES IN RTH'S TREATMENTS

RTH-4	<i>A</i>	<i>B</i>	<i>A</i>	<i>A</i>
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	<i>A</i>	<i>A</i>	<i>B</i>	<i>A</i>
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	<i>A</i>	<i>A</i>	<i>B</i>	<i>A</i>
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	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
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	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
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	<i>A</i>	<i>B</i>	<i>A</i>	<i>A</i>
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 Iriberry's Table 1

Folk game theory deviates from equilibrium logic in ways that are reminiscent of RTH's results.

Any game theorist worth his salt would respond to the Yushchenko quote:

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

with

“If investigators thought that way, a meeting with government officials is precisely where a government *would* try to kill an opponent.”

As we will see, the quote reflects the reasoning of an *L1* poisoner, or equivalently of an *L2* investigator reasoning about an *L1* poisoner.

RTH's data 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 does not help to explain the asymmetry of choice distributions.

Resolution:

The role asymmetry in behavior and how it is linked to the game's payoff asymmetry points strongly in the direction of level- k thinking or CH, and is a mystery from the viewpoint of other theories we know of.

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 $p > 1/2$ and $q > 1/4$.

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

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

- Given $L0$'s attraction to salient locations, $L1$ hidiers choose *central A* to avoid $L0$ seekers and $L1$ seekers avoid *central A* searching for $L0$ hidiers (the data suggest that end locations are more salient than B).
- For similar reasons, $L2$ hidiers choose *central A* with probability between 0 and 1 (breaking payoff ties) and $L2$ seekers choose it with probability 1.
- $L3$ hidiers avoid *central A* and $L3$ seekers choose it with probability between zero and one (breaking payoff ties).
- $L4$ hidiers 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	Seeker	Expected payoff	Choice probability	Expected payoff	Choice probability
	$p < 2q$	$p < 2q$	$p > 2q$	$p > 2q$		$p < 2q$	$p < 2q$	$p > 2q$	$p > 2q$
<i>L0 (Pr, r)</i>					<i>L0 (Pr, r)</i>				
A	—	$p/2$	—	$p/2$	A	—	$p/2$	—	$p/2$
B	—	q	—	q	B	—	q	—	q
A	—	$1-p-q$	—	$1-p-q$	A	—	$1-p-q$	—	$1-p-q$
A	—	$p/2$	—	$p/2$	A	—	$p/2$	—	$p/2$
<i>L1 (Pr, s)</i>					<i>L1 (Pr, s)</i>				
A	$1-p/2 < 3/4$	0	$1-p/2 < 3/4$	0	A	$p/2 > 1/4$	0	$p/2 > 1/4$	1/2
B	$1-q < 3/4$	0	$1-q < 3/4$	0	B	$q > 1/4$	1	$q > 1/4$	0
A	$p+q > 3/4$	1	$p+q > 3/4$	1	A	$1-p-q < 1/4$	0	$1-p-q < 1/4$	0
A	$1-p/2 < 3/4$	0	$1-p/2 < 3/4$	0	A	$p/2 > 1/4$	0	$p/2 > 1/4$	1/2
<i>L2 (Pr, r)</i>					<i>L2 (Pr, r)</i>				
A	1	1/3	1/2	0	A	0	0	0	0
B	0	0	1	1/2	B	0	0	0	0
A	1	1/3	1	1/2	A	1	1	1	1
A	1	1/3	1/2	0	A	0	0	0	0
<i>L3 (Pr, u)</i>					<i>L3 (Pr, u)</i>				
A	1	1/3	1	1/3	A	1/3	1/3	0	0
B	1	1/3	1	1/3	B	0	0	1/2	1/2
A	0	0	0	0	A	1/3	1/3	1/2	1/2
A	1	1/3	1	1/3	A	1/3	1/3	0	0
<i>L4 (Pr, v)</i>					<i>L4 (Pr, v)</i>				
A	2/3	0	1	1/2	A	1/3	1/3	1/3	1/3
B	1	1	1/2	0	B	1/3	1/3	1/3	1/3
A	2/3	0	1/2	0	A	0	0	0	0
A	2/3	0	1	1/2	A	1/3	1/3	1/3	1/3
Total	$p < 2q$		$p > 2q$		Total	$p < 2q$		$p > 2q$	
A	$rp/2+(1-\varepsilon)[t/3+u/3]+(1-r)\varepsilon/4$		$rp/2+(1-\varepsilon)[u/3+v/2]+(1-r)\varepsilon/4$		A	$rp/2+(1-\varepsilon)[u/3+v/3]+(1-r)\varepsilon/4$		$rp/2+(1-\varepsilon)[s/2+v/3]+(1-r)\varepsilon/4$	
B	$rq+(1-\varepsilon)[u/3+v]+(1-r)\varepsilon/4$		$rq+(1-\varepsilon)[t/2+u/3]+(1-r)\varepsilon/4$		B	$rq+(1-\varepsilon)[s+v/3]+(1-r)\varepsilon/4$		$rq+(1-\varepsilon)[u/2+v/3]+(1-r)\varepsilon/4$	
A	$r(1-p-q)+(1-\varepsilon)[s+t/3]+(1-r)\varepsilon/4$		$r(1-p-q)+(1-\varepsilon)[s+t/2]+(1-r)\varepsilon/4$		A	$r(1-p-q)+(1-\varepsilon)[t+u/3]+(1-r)\varepsilon/4$		$r(1-p-q)+(1-\varepsilon)[t+u/2]+(1-r)\varepsilon/4$	
A	$rp/2+(1-\varepsilon)[t/3+u/3]+(1-r)\varepsilon/4$		$rp/2+(1-\varepsilon)[u/3+v/2]+(1-r)\varepsilon/4$		A	$rp/2+(1-\varepsilon)[u/3+v/3]+(1-r)\varepsilon/4$		$rp/2+(1-\varepsilon)[s/2+v/3]+(1-r)\varepsilon/4$	

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

Model	Ln L	Parameter estimates	Observed or predicted choice frequencies				MSE	
			Player	A	B	A		A
Observed frequencies (624 hidiers, 560 seekers)			H	0,2163	0,2115	0,3654	0,2067	—
			S	0,1821	0,2054	0,4589	0,1536	
Equilibrium without perturbations	-1641,4		H	0,2500	0,2500	0,2500	0,2500	0,00970
			S	0,2500	0,2500	0,2500	0,2500	
Equilibrium with restricted perturbations	-1568,5	$e_H \equiv e_S = 0,2187$ $f_H \equiv f_S = 0,2010$	H	0,1897	0,2085	0,4122	0,1897	0,00084
			S	0,1897	0,2085	0,4122	0,1897	
Equilibrium with unrestricted perturbations	-1562,4	$e_H = 0,2910, f_H = 0,2535$ $e_S = 0,1539, f_S = 0,1539$	H	0,2115	0,2115	0,3654	0,2115	0,00006
			S	0,1679	0,2054	0,4590	0,1679	
Level- <i>k</i> with a role-symmetric LO that favors salience	-1564,4	$p > 1/2$ and $q > 1/4, p > 2q,$ $r = 0, s = 0,1896, t = 0,3185,$ $u = 0,2446, v = 0,2473, \epsilon = 0$	H	0,2052	0,2408	0,3488	0,2052	0,00027
			S	0,1772	0,2047	0,4408	0,1772	
Level- <i>k</i> with a role-asymmetric LO that favors salience for seekers and avoids it for hidiers	-1563,8	$p_H < 1/2$ and $q_H < 1/4,$ $p_S > 1/2$ and $q_S > 1/4,$ $r = 0, s = 0,66, t = 0,34,$ $\epsilon = 0,72; u \equiv v \equiv 0$ imposed	H	0,2117	0,2117	0,3648	0,2117	0,00017
			S	0,1800	0,1800	0,4600	0,1800	
Level- <i>k</i> with a role-symmetric LO that avoids salience	-1562,5	$p < 1/2$ and $q < 1/4, p < 2q,$ $r = 0, s = 0,3636, t = 0,0944,$ $u = 0,3594, v = 0,1826, \epsilon = 0$	H	0,2133	0,2112	0,3623	0,2133	0,00006
			S	0,1670	0,2111	0,4549	0,1670	

Crawford and Iriberry's Table 3

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

($L4$ s don't matter here because they never choose *central A* (Table 2), hence they are not implicated in the major aggregate patterns.)

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

Given this, a 50-50 mix of $L1$ s and $L2$ s in both player roles would imply (Table 2, 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*.

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 our analysis suggests that RTH’s subjects were actually quite strategic and in fact more than usually sophisticated (with many *L3s* and even some *L4s*)—they just didn’t follow equilibrium logic.

In Crawford and Iriberry'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, although it yields a small improvement in fit (Table 3), would beg the question of why subjects' responses were so role-asymmetric and risk overfitting.

("Beg" means "refuse to address," not "emphasize the importance of".)

As noted above, in the analogous analysis of the Yushchenko quote ("Any government wanting to kill an opponent...would not try it at a meeting with government officials") the quote reflects the reasoning of an *L1* poisoner or an *L2* investigator reasoning about an *L1* poisoner.

Model evaluation

Although our empirically based prior about the hump shape and location of the type distribution imposes some discipline, the freedom to specify $L0$ leaves room for doubts about overfitting and portability.

To see if the proposed level- k explanation of RTH's results is more than a “just-so” story, Crawford and Iriberri compare 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 similar payoff perturbations, and
- Alternative level- k specifications (for example, with role-asymmetric $L0$ or an $L0$ that avoids salience, as in Table 3).

Crawford and Iriberry test 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.

A more challenging test regards portability, the extent to which a model estimated from subjects' responses to one game can be extended to predict or explain other subjects' responses to different games.

Crawford and Iriberry considered 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.

They tested for portability by using the leading alternative models, estimated from RTH's data, to "predict" subjects' initial responses in O'Neill's and Rapoport and Boebel's games.

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

One player, say the row player (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 (s)	2 (s)	3 (s)	J (h)
A (h)	0 1	1 0	1 0	0 1
2 (h)	1 0	0 1	1 0	0 1
3 (h)	1 0	1 0	0 1	0 1
J (s)	0 1	0 1	0 1	1 0

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 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 tests directly address 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.

Even Obama and McCain could 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 our RTH estimates suggested that there, 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, and such evidence is more likely to yield convergence than if it were partly a strategic issue.

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

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

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

Crawford and Iriberri did not then have O'Neill's data, the most natural choice.

But they did know that discussions of it had been dominated by an "Ace effect," whereby row and column players, aggregated over all 105 rounds, played A with frequencies 22.0% and 22.6%, significantly above the equilibrium 20%.

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

But what about the equally powerful connotations of the Joker and its unique payoff role?

They seem to make it even more salient than A, but in the aggregate data row subjects chose J with frequencies of only 36%, and column subjects with frequencies of only 43%.

Further, with an Obama-McCain specification of $L0$ and the resulting types' decisions in O'Neill's game (Tables A3 and A4 in the paper's web appendix, next slides), no behaviorally plausible level- k model will make row players ("1s") play A more than the equilibrium 20%.

Excluding $L0$ s, depending on whether A or J is more salient this would require a population of almost entirely $L4$ s or, respectively, $L3$ s.

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	1/2	0	1/2	0
3	0	0	1	1/2	3	1/2	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.
<i>L0 (Pr. R)</i>			<i>L0 (Pr. r)</i>		
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
<i>L1 (Pr. S)</i>			<i>L1 (Pr. s)</i>		
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
<i>L2 (Pr. T)</i>			<i>L2 (Pr. t)</i>		
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
<i>L3 (Pr. U)</i>			<i>L3 (Pr. u)</i>		
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
<i>L4 (Pr. V)</i>			<i>L4 (Pr. v)</i>		
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	$Ra+(1-\epsilon)[u/3]+(1-r)\epsilon/4$		A	$ra+(1-\epsilon)[s+t/3]+(1-r)\epsilon/4$	
2	$r(1-a-j)/2+(1-\epsilon)[t/2+u/3]+(1-r)\epsilon/4$		2	$r(1-a-j)/2+(1-\epsilon)[t/3]+(1-r)\epsilon/4$	
3	$R(1-a-j)/2+(1-\epsilon)[t/2+u/3]+(1-r)\epsilon/4$		3	$r(1-a-j)/2+(1-\epsilon)[t/3]+(1-r)\epsilon/4$	
J	$Rj+(1-\epsilon)[s+v]+(1-r)\epsilon/4$		J	$rj+(1-\epsilon)[u+v]+(1-r)\epsilon/4$	

Crawford and Iriberry decided to get the data and test the model on it anyway, speculating (based on the level- k model's success in RTH's and other games) that initial responses must not have an Ace effect.

The hunch was right: There is no Ace effect for initial responses.

Instead there is 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 roughly ten times larger than $(22-20)\%$ and $(22.6-20)\%$.)

Moreover, unlike the Ace effect, the Joker effect and the other frequencies *can* be gracefully explained by a level- k model with an Obama-McCain $L0$ that probabilistically favors salient A and J cards.

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 LO that favors salience	$a > 1/4$ and $j > 1/4$ $3j - a < 1, a + 2j < 1$	1	0,0824	0,1772	0,1772	0,5631	0,0018
		2	0,1640	0,1640	0,1640	0,5081	0,0066
Level- k with a role-symmetric LO that favors salience	$a > 1/4$ and $j > 1/4$ $3j - a < 1, a + 2j > 1$	1	0,0000	0,2541	0,2541	0,4919	0,0073
		2	0,2720	0,0824	0,0824	0,5631	0,0050
Level- k with a role-symmetric LO 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 LO that favors salience for locations for which player is a seeker and avoids it for locations for which player is a hider	$a_1 < 1/4, j_1 > 1/4;$ $a_2 > 1/4, j_2 < 1/4$ $3j_1 - a_1 < 1,$ $a_1 + 2j_1 < 1, 3a_2 + j_2 > 1$	1	0,1804	0,2729	0,2729	0,2739	0,0291
		2	0,1804	0,1804	0,1804	0,4589	0,0117

Crawford and Iriberry's Table 5

Equilibrium or LQRE with perturbations are well-defined for O'Neill's game, but 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. Level- k is well-defined, and explains some but not all patterns in their data.)

Crawford and Iriberry's analysis traces the superior portability of the level- k model directly to the fact that LO is psychological rather than strategic, and is based on simple and universal intuition and evidence.

(If LO 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 they did from RTH's games to O'Neill's.)

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

Miscoordination in Schelling-style coordination games with non-neutral framing of decisions (Crawford Gneezy, and Rottenstreich 2008 *AER* (“CGR”))

CGR randomly paired subjects to play games with non-neutral framing of decisions like those in Schelling’s (1960) classic “meeting in New York City” experiments.

But except for a symmetric game like Schelling’s games, CGR used games with payoff asymmetries like Battle of the Sexes.

As in Schelling’s experiments, there was a commonly observable labeling of decisions:

In unpaid pilots, run in Chicago, CGR used naturally occurring labels, pitting the world-famous Sears Tower versus the little-known AT&T Building across the street.

		P2	
		Sears	AT&T
P1	Sears	100,100	0,0
	AT&T	0,0	100,100

Symmetric

		P2	
		Sears	AT&T
P1	Sears	100,101	0,0
	AT&T	0,0	101,100

Slight Asymmetry

		P2	
		Sears	AT&T
P1	Sears	100,110	0,0
	AT&T	0,0	110,100

Moderate Asymmetry
Chicago Skyscrapers



Sears Tower with the AT&T Building in the background on its left (the AT&T Building is actually almost as tall as Sears Tower)

The Chicago Skyscrapers results replicated Schelling's results in the Symmetric version of the game, but 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

The salience of Sears Tower makes it easy and in principle obvious for subjects to coordinate on the “both-Sears” equilibrium; and they 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.

But 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’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, CGR conducted more formal, paid treatments used abstract labels, pitting X against Y, with X presumed (and shown) to be more salient than Y.

		P2	
		X	Y
P1	X	5,5	0,0
	Y	0,0	5,5
Symmetric			

		P2	
		X	Y
P1	X	5,5.1	0,0
	Y	0,0	5.1,5
Slight Asymmetry			

		P2	
		X	Y
P1	X	5,6	0,0
	Y	0,0	6,5
Moderate Asymmetry			

		P2	
		X	Y
P1	X	5,10	0,0
	Y	0,0	10,5
Large Asymmetry			

Like the salience of Sears Tower, the salience of X 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			

Even tiny payoff asymmetries cause 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.

Puzzles:

- 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?

Resolution:

Standard notions such as 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.

CGR 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 $\frac{1}{2}$.
- In any asymmetric game, (for simplicity only) whether or not label-salience opposes payoff-salience, $L0$ chooses its payoff-salient decision with probability $p > \frac{1}{2}$.

(These assumptions are consistent with Crawford and Iriberri's $L0$ assumptions, because their games had no payoff-salience.)

Under these assumptions about $L0$, $L1$'s and $L2$'s choices for 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 (CGR's Table 3, reproduced next slide) 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 the subject pool (Table 3).

	Symmetric Labeled (SL)	Asymmetric Slight Labeled (ASL)	Asymmetric Moderate Labeled (AML)	Asymmetric Large Labeled (ALL)
Payoffs for coordinating on “X”	\$5, \$5	\$5, \$5.10	\$5, \$6	\$5, \$10
Payoffs for coordinating on “Y”	\$5, \$5	\$5.10, \$5	\$6, \$5	\$10, \$5
Pr{X} for P1 L0	$> \frac{1}{2}$	$1-p$	$1-p$	$1-p$
Pr{X} for P2 L0	$> \frac{1}{2}$	p	p	p
Pr{X} for P1 L1	1	1	0	0
Pr{X} for P1 L2	1	0	1	1
Pr{X} for P2 L1	1	0	1	1
Pr{X} for P2 L2	1	1	0	0
Total P1 predicted Fr{X}	100%	$100q\%$	$100(1-q)\%$	$100(1-q)\%$
Total P1 predicted Fr{X} $q=0.7$	100%	70%	30%	30%
Total P1 observed Fr{X}	76%	78%	33%	36%
Total P2 predicted Fr{X}	100%	$100(1-q)\%$	$100q\%$	$100q\%$
Total P2 predicted Fr{X} $q=0.7$	100%	30%	70%	70%
Total P2 observed Fr{X}	76%	28%	61%	60%
Table 3. L1's and L2's choice probabilities in X-Y treatments when $0.505 < p < 0.545$				

CGR's Table 3

The details are as follows:

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

- In the symmetric game, with no payoff salience, *L0* favors the salience of X.
- *L1* P1s and P2s therefore both choose X.
- *L2* P1s and P2s do the same.

In this case the model predicts that 100% of P1s and P2s will choose X. Thus, here it makes the same prediction as equilibrium selection based on salience as in a Schelling focal point. This is fairly accurate, but it overstates the homogeneity of the subject pool.

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

Slight Asymmetry

- In the slightly asymmetric game, with $p > 0.505$ ($= 5.1/[5.1+5]$), the payoff differences are small enough that *L1* P1s choose P2s' payoff-salient decision, X, because *L1* P1s think it is sufficiently likely that *L0* P2s will choose X that X yields them higher expected payoffs.
- *L2* P2s, who best respond to *L1* P1s, thus choose X as well.
- With $p > 0.505$, *L1* P2s choose P1s' payoff-salient decision, Y, because *L1* P2s think it sufficiently likely that *L0* P1s will choose Y.
- *L2* P1s thus choose Y.

In this case the model predicts that *L1* P1s choose X and *L2* P1s choose Y, while *L1* P2s choose Y and *L2* P2s choose X. Thus, when $q = 0.7$, the model predicts that 70% of P1s will choose X but only 30% of P2s will choose X, reasonably close to the observed 78% and 28%.

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

Moderate Asymmetry

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

Large Asymmetry

- In the games with moderate or large payoff asymmetries, $L0$'s payoff bias is strong enough, but not too strong ($p < 0.545 (= 6/[6+5])$), that $L1$ P1s and P2s both choose their own instead of their partners' payoff- salient decisions, Y for P1s and X for P2s.
- $L2$ P1s choose X and $L2$ P2s choose Y.

In this case the model predicts that $L1$ P1s choose Y and $L2$ P1s choose X, while $L1$ P2s choose X and $L2$ P2s choose Y. Thus, when $q = 0.7$, the model predicts that 30% of P1s will choose X but 70% of P2s will choose X, again close to the observed 33-36% and 61-60%.

Preplay communication of intentions in zero-sum two-person games (Crawford 2003 *AER*)

Consider a simple perturbed matching pennies game, viewed as a model of the Allies' choice of where to invade Europe on D-Day:

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

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



An Inflatable “Tank” from Operation Fortitude

In an equilibrium analysis of a zero-sum game preceded by a cheap-talk message regarding intentions, the sender must make his message uninformative, and the receiver must ignore it.

If in equilibrium the receiver found it optimal to respond to the message, his response would benefit him, and so hurt the sender, who would therefore do better by making the message uninformative.

Thus communication can have no effect in any equilibrium, and the underlying game must be played according to its unique mixed-strategy equilibrium.

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.

Moreover, in my stylized version of D-Day:

- The deception succeeded (the Allies faked preparations for invasion at Calais, the Germans defended Calais and left Normandy lightly defended, and the Allies then invaded Normandy).
- But the sender won in the less beneficial of the two possible ways.

Admittedly, D-Day is only one datapoint (if that)....

But there's an ancient Chinese antecedent of D-Day, Huarongdao, in which General Cao Cao chooses between two roads, the comfortable Main Road and the awful Huarong Road, trying to avoid capture by General Kongming (thanks to Duoze Li of CUHK for the reference to Luo Guanzhong's historical novel, *Three Kingdoms*).

		Kongming	
		Main Road	Huarong
Cao Cao	Main Road	-1 3	1 0
	Huarong	0 1	-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.

In Huarongdao, essentially the same thing happened as in D-Day: Kongming lit campfires on the Huarong road; Cao Cao was fooled by this into thinking Kongming would ambush him on the *Main Road*; and Kongming captured Cao Cao but only by taking Huarong Road.

(The ending however was happy: Kongming later let Cao Cao go.)

In what sense did the “essentially the same thing” happen?

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

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 inverted the message.

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 way.

Why was Cao Cao fooled by Kongming's message?

As we have already seen in *Far Pavilions* Escape!, one advantage of using fiction as data is that it can reveal cognition:

- *Three Kingdoms* gives Kongming's rationale for sending a deceptively truthful message: "Have you forgotten the tactic of 'letting weak points look weak and strong points look strong'?"
- It also gives Cao Cao's rationale for inverting Kongming's message: "Don't you know what the military texts say? 'A show of force is best where you are weak. Where strong, feign weakness.' "

Cao Cao must have bought a used, out-of-date edition....

As we will see, with *L0* suitably adapted to this setting, Cao Cao's rationale resembles *L1* thinking; but Kongming's rationale resembles *L2* thinking.

Puzzle:

We can now 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?)
- Was it a coincidence that the same thing happened in both cases?

Resolution:

A level- k analysis suggests that it was more than a coincidence.

Assume that Allies' and Germans' types are drawn from separate distributions, including both level- k , or *Mortal*, types and a fully strategically rational, or *Sophisticated*, type (interesting but rare).

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 *EER* quote:

“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.”

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

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

Here a uniform random $L0$ does not seem natural, at least for senders.

Instead *Mortal* types' behaviors regarding the message are anchored 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 2008—but the distinction is partly semantic because $L1$ receivers' best responses to truthful $L0$ senders are credulous.)

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 such 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 type frequencies.

But the analysis can usefully be extended to allow the possibility of *Sophisticated* Allies and Germans.

To do this note that *Mortals'* strategies are determined independently of each other's and *Sophisticated* players' strategies, and so can be treated as exogenous (even though they affect others' payoffs).

Next, plug in the distributions of *Mortal Allies'* and *Germans'* independently determined behavior 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.

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 that 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 equilibrium, in which *Sophisticated* Allies can predict *Sophisticated* Germans' decisions, and vice versa.

In the latter kind of equilibrium, *Sophisticated* Allies send the message that fools the most common kind of *Mortal* German (depending on how many believe messages and how many, like Cao Cao, invert them) and attack Normandy; while *Sophisticated* Germans defend Calais (because they know that *Mortal* Allies, who predominate in this case, will attack Calais).

For subtle reasons, there is no pure-strategy equilibrium in which *Sophisticated Allies* feint at Normandy and attack Calais (p. 143):

The asymmetry across actions of Proposition 1's conclusion that a *Sophisticated Receiver's* strategy is R, R in *all* pure-strategy sequential equilibria is an important part of the model's explanation of Operation Fortitude. The conclusion is trivial if the probability of a *Mortal Sender* is high enough to make R a dominant strategy in the underlying game; but, somewhat surprisingly, it remains valid even if the probability is not high enough, as long as the game has a pure-strategy sequential equilibrium. It superficially resembles the babbling equilibrium of the standard analysis, but it actually stems from the assumption that *Mortal Senders* play U.²³ In a pure-strategy sequential equilibrium, a *Sophisticated Sender's* deviation from his equilibrium message "proves" to a *Sophisticated Receiver* that the Sender is *Mortal*, making R the Receiver's best response off the equilibrium path. If a *Sophisticated Sender* plays U on the equilibrium path, the conclusion is immediate. If, instead, a *Sophisticated Sender* plays D on the equilibrium path while a *Sophisticated Receiver* plays L, the *Sophisticated Sender's* message fools only the most frequent type of *Mortal Receiver*, at a payoff gain of 1 per unit of probability. But such a *Sophisticated Sender* could reverse his message and action, again fooling the most frequent type of *Mortal Receiver*, but now at a payoff gain of $a > 1$ per unit, a contradiction.

In the pure-strategy equilibrium, 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 of the possible ways.

Thus for plausible parameter values, without postulating an unexplained difference in the sophistication of Allies and Germans, the model explains why even *Sophisticated* Germans might allow themselves to be "fooled" by a costless message from an enemy.

In a weaker sense (resting on a preference for pure-strategy equilibria and high-probability predictions), the model also explains why *Sophisticated* Allies don't feint at Normandy and attack Calais, even though this would be more profitable if it succeeded.

Overbidding in independent-private-value and common-value auctions (Crawford and Iriberri 2007 *Econometrica*)

Equilibrium predictions		
	First-Price	Second-Price
Independent-Private-Value Auctions	Shaded Bidding	Truthful Bidding
Common-Value Auctions	Value Adjustment + Shaded Bidding	Value Adjustment

Puzzle: Systematic overbidding (relative to equilibrium) has been observed in subjects' initial responses to all kinds of auctions (Goeree, Holt, and Palfrey 2002 *JET*; Kagel and Levin 1986 *AER*, 2000; Avery and Kagel 1997 *JEMS*; Garvin and Kagel 1994 *JEBO*).

(With independent private values, most of the examples that have been studied experimentally do not separate level- k from equilibrium bidding strategies, hence our choice to study GHP's results.)

But the literature has proposed completely different explanations of overbidding for private- and common-value auctions:

- Risk-aversion and/or joy of winning for private-value auctions.
- Winner's curse for common-value auctions.

Resolution:

Crawford and Iriberri propose a level- k analysis that provides a unified explanation of these results, without invoking risk-aversion and/or joy of winning.

Their analysis extends Kagel and Levin's 1986 *AER* and Holt and Sherman's 1994 *AER* analyses of "naïve bidding".

It also builds on Eyster and Rabin's ("ER") 2005 *Econometrica* analysis of "cursed equilibrium" and CHC's (2004, Section VI) CH analysis of zero-sum betting.

The analysis makes it possible to explore how to extend level- k models to an important class of incomplete-information games.

It also links experiments on auctions to experiments on strategic thinking.

It also makes it possible to explore the robustness of equilibrium auction theory to failures of the equilibrium assumption.

The key issue is how to specify L_0 ; there are two natural possibilities:

- *Random L_0* bids uniformly on the interval between the lowest and highest possible values (even if above own realized value).
- *Truthful L_0* bids its expected value conditional on its own signal (meaningful here, though not in all incomplete-information games).

In judging these, bear in mind that L_0 describes only the instinctive starting point of a subject's strategic thinking about others; higher L_k s model the actual strategic thinking.

The model constructs separate type hierarchies on these L_0 s, and allows each subject to be one of the types, from either hierarchy.

Random (Truthful) L_k is L_k defined by iterating best responses from *Random (Truthful) L_0* ; and is not itself random or truthful.

Given a specification of $L0$, the optimal bid must take into account:

- Value adjustment for the information revealed by winning (only in common-value auctions).
- The bidding trade-off between the higher price paid if the bidder wins and the probability of winning (only in first-price auctions).

With regard to value adjustment, Random $L1$ does not condition on winning because Random $L0$ bidders bid randomly, hence independently of their values; Random $L1$ is “fully cursed” (ER).

All 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.

The question, empirically, is whether the distribution of types' bids (for example, a 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 $L2$.