Strategic Thinking

Vincent P. Crawford

**All Souls College, University of Oxford,
and University of California, San Diego**

UNIVERSITY OF OXFORD

UC San Diego

erc European Research Council
Established by the European Commission

19 October 2024

Introduction

Strategic thinking—people's attempts to predict others' decisions in games—pervades human interaction.

Traditional game theory has a built-in model of strategic thinking in its notion of Nash equilibrium, henceforth shortened to "equilibrium".

Equilibrium is defined as a combination of decisions, one for each player, such that each player's decision is rational in the decision-theoretic sense of maximizing her/his expected utility ("payoff"), given other players' decisions.

Although this definition identifies equilibrium decisions without reference to players' beliefs, equilibrium can usefully be viewed as "equilibrium in beliefs," in which rational players have common beliefs about each other's decisions that are self-confirming, given the (decision-theoretically) rational decisions they imply.

In other words, players have "rational expectations" about each other's decisions.

This is a far stronger assumption than rationality of individual decisions.

Even so, equilibrium remains the method of choice in strategic applications.

But I will argue that it is possible in some applications to identify nonequilibrium models of strategic thinking that systematically out-predict equilibrium.

This claim raises two questions:

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

- And if such models exist, how could anyone identify one within the enormous set of logically possible nonequilibrium models?

These questions are answered, to some extent, by recent experimental work that studies strategic thinking by eliciting subjects' initial responses to games.

Equilibrium can be justified in two main ways:

- In learning models, rational players learn to predict others' decisions from their previous responses to analogous games:

  Their decisions then often converge to some equilibrium in the "stage game" that is repeated; they need not think at all, beyond recognizing the analogies

- In thinking models, rational players predict others' decisions in their initial responses to games played without close precedents epistemically:

  If players are rational and it is common knowledge that their thinking follows the logic of some equilibrium, their beliefs will yield decisions in that equilibrium (Brandenburger 1992, Proposition 4)

When equilibrium has a plausible learning justification, it often yields a coherent and empirically reliable account of behavior, with no strategic thinking required.

But when it does not, equilibrium must be justified by strategic thinking.

However, thinking justifications are much more fragile than learning justifications.

In all but the simplest games, equilibrium logic requires fixed-point, contingent, or extensive iterated-dominance reasoning.

These are all subtle enough that even people who can follow them may doubt that others can.

Many applications in which equilibrium lacks a plausible learning justification do not support a blanket assumption that people's decisions will be in equilibrium.

Despite these concerns, most researchers reflexively assume equilibrium:

- Perhaps because they hope equilibrium predictions will still be correct on average, or fear that without equilibrium there will be no basis for analysis

- Or because they overestimate the scope of learning justifications, when few applications have games with precedents as close as learning models assume

When equilibrium is not well justified, it seems more useful to address the above concerns by using theory and evidence to identify nonequilibrium models that are better justified and better predict people's decisions.

The goal is not to supplant equilibrium analysis, but to extend it in ways that make it more useful.

The *potential* benefits are clear:

- In some applications nonequilibrium models may predict that people's strategic thinking will lead them to mimic their equilibrium decisions, thus establishing the robustness of equilibrium-based predictions

- In others they may predict the forms and frequencies of deviations from equilibrium, resolving empirical puzzles equilibrium analysis cannot address

It is not obvious that those potential benefits can be realized.

But in recent experimental work that studies strategic thinking by eliciting people's initial responses to games, few if any subjects use fixed-point or extensively iterated dominance reasoning when equilibrium requires it.

To the extent that there are decision rules that reliably describe what such subjects do (which is considerable), they must be nonequilibrium rules.

Well-designed experiments often precisely identify subjects' decision rules, even within the enormous set of logically possible nonequilibrium rules.

Identifying subjects' decision rules often yields tractable models that can predict, in advance, what deviations from equilibrium will occur and their forms and likelihoods: out-predicting equilibrium despite its rational expectations character.

I next review the leading models of strategic thinking, then review some representative experimental evidence, and finally discuss some applications.

Models of Strategic Thinking

I now review the leading models of strategic thinking, focusing on simultaneous-move ("normal-form") games with symmetric information (I call players' choices "decisions" rather than "strategies"), played in settings that do not allow learning.

(Such models of strategic thinking yield insights that also help in modeling asymmetric information, learning from imperfect analogies, and equilibrium selection via learning, but that's another story.)

I assume the structure of each game is common knowledge, in that each player knows the structure, knows that other players know it, and so on ad infinitum.

To focus as sharply as possible on strategic thinking, I also maintain standard assumptions about individual preferences and judgment.

Apology: In principle, any aspect of behavioral decision theory—such as present-biased, reference-dependent, or social preferences, or heuristics and biases in probabilistic judgment—is equally relevant in behavioral game theory.

But here, as in most of behavioral game theory, I follow a “divide-and-conquer” strategy, taking individual decisions as self-interested and decision-theoretically rational while considering the issue that is unique to games, strategic thinking.

The implicit hope is that in behavioral game theory done this way, behavioral decision theory will be “plug and play”, so that combining the two will yield the full understanding of behavioral game theory needed for useful applications.

A “decision rule” describes a player’s decisions in all the games s/he plays.

A “model” of strategic thinking specifies a frequency distribution over players’ possible decision rules and an error structure.

The decision rules I consider here include:

- Equilibrium

- Rationalizability and *k*-rationalizability

- Quantal response equilibrium (“QRE”), logit QRE (“LQRE”), regular QRE (“RQRE”), and rank-dependent choice equilibrium (“RDCE”)

- Dominance-*k* models (“*Dk*”)

- Level-*k* models (“*Lk*”)

- Cognitive hierarchy models (“CH”)

These are all general decision rules, applicable to any game, which determine players’ beliefs as well as their decisions. They all allow decision errors, except for the set-valued rules rationalizability, *k*-rationalizability, RQRE, and RDCE.

Recall that I am focusing on simultaneous-move ("normal-form") games with symmetric information, whose structures are common knowledge.

One issue that has not been discussed enough in models of strategic thinking is the information that a player needs to implement her/his rule's decisions.

For rules that specify precise decisions, a rule's decision can be thought of as the sum of its mean decision and a mean-zero error term. (This includes rules like QRE and LQRE, which distinguish errors from mean decisions only implicitly.)

I take a rule's mean decisions to represent the *intentional* part of its decisions.

I call a decision rule "self-informed" if a player can identify its mean decisions using only information available to the player her/himself.

This notion can easily be extended to rules like rationalizability, *k*-rationalizability, RQRE, and RDCE, which are set-valued and don't have explicit error structures.

A decision rule that is *not* self-informed entails no *logical* inconsistency.

But assuming that a player's mean decisions depend on information about others' decision rules—the only other information that might matter in symmetric-information games—implicitly ascribes more than usual sophistication to her/him.

Self-informedness is not a standard notion, but it is an important desideratum for decision rules meant to describe strategic thinking, or cognition more generally.

Leading models of strategic thinking differ substantially in self-informedness.

Equilibrium

Recall that equilibrium is defined as a combination of decisions, one for each player, such that each player's decision is rational in the decision-theoretic sense of maximizing her/his expected utility ("payoff"), given other players' decisions.

In the "equilibrium in beliefs" view, rational players have common beliefs about each other's decisions that are self-confirming, given the rational decisions they imply: Players have "rational expectations" about each other's decisions.

Analysts are driven to this rational-expectations assumption, which is far stronger than decision-theoretic rationality, because in games rationality, even if common knowledge, seldom restricts behavior enough to be useful.

(Moreover, many games have multiple equilibria, so even equilibrium does not always make unambiguous predictions, and it is often further augmented by equilibrium-selection refinements such as risk- or payoff-dominance.)

The simplest way to allow for errors in an equilibrium analysis is to add errors to its predicted decisions, yielding a model I will call "equilibrium plus noise".

An equilibrium plus noise player's decision densities are usually assumed to depend on her/his decisions' expected payoffs, evaluated via equilibrium beliefs, that is assuming that all others make their equilibrium decisions with certainty.

Each player's decision densities are usually assumed to have a known form such as logit, independent across players with zero mean and estimated precision.

If an equilibrium or equilibrium plus noise decision rule is augmented by a refinement that uniquely selects an equilibrium, then that rule is self-informed.

Even though the refinement involves others' decisions, it rests on an equilibrium prediction of their decisions, so it can be implemented without further information.

Rationalizability and *k*-rationalizability

In two-person games (*n*-person games differ in ways that are unimportant here), rationalizability or finitely iterated strict dominance and *k*-rationalizability capture the implications of common or finitely iterated knowledge of rationality.

k-rationalizability reflects the implications of *k* levels of mutual knowledge of rationality. Rationalizability is equivalent to *k*-rationalizability for all *k*.

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

A 2-rationalizable strategy (the sets R2) is one for which there exists a profile of others' 1-rationalizable strategies that make it a best response. And so on….

Each notion generally yields set-valued restrictions on individual players' strategy choices (unlike equilibrium, which restricts the relationship of their strategies).

Decision errors are not usually allowed, because their predictions are set-valued.

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

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

Strictly dominance-solvable

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

Unique equilibrium without dominance

Any equilibrium decision is *k*-rationalizable for all *k*.

In games that are strictly dominance-solvable in *k* rounds, *k*-rationalizability implies that players have the same beliefs, so that any combination of *k*-rationalizable decisions is in equilibrium, as in the first game.

But in general, not all combinations of rationalizable decisions are in equilibrium.

k-rationalizability and rationalizability can allow deviations from equilibrium, as in the second game, where there is a "helix" of beliefs, consistent with common knowledge of rationality, to support any decision combination.

But except for the equilibrium beliefs (M, C), the beliefs in the helix differ across players and rounds, and many are behaviorally implausible.

Rationalizability and *k*-rationalizability decision rules are self-informed, because their sets of decisions are completely determined by iterated knowledge of players' rationality and the structure of the game.

Quantal response equilibrium, logit QRE ("LQRE"), regular QRE ("RQRE"), and rank-dependent choice equilibrium ("RDCE")

In a QRE, players' decisions are noisy, with given error distribution and precision.

A player's decision density is sensitive to her/his decisions' expected payoffs, but—unlike in equilibrium plus noise—is evaluated taking the noisiness of others' decisions into account.

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

As players' error precisions approach ∞, QRE converges to equilibrium.

As players' error precisions approach 0, QRE converges to independent uniform randomization over each player's decisions.

Allowing players' decisions to be noisy is often seen as enhancing QRE's behavioral plausibility. After all, people's decisions *are* noisy.

Anyone who accepts that players' decisions are noisy but insists on equilibrium/rational expectations will be led to define a notion like QRE.

However, QRE loses the behavioral plausibility it gains by allowing decision noise, by replacing equilibrium's fixed-point reasoning in the space of beliefs about decisions (deterministic or mixed with limited support), with fixed-point reasoning in the much larger space of full-support noisy decision distributions.

If the rationale for QRE is that equilibrium reasoning is cognitively taxing, QRE reasoning is doubly taxing.

QRE with a specified error distribution shares equilibrium plus noise’s precision: Given an error distribution it precisely predicts a distribution of players’ decisions.

In applications QRE’s error distribution is often assumed to be logit (“LQRE”).

LQRE (or QRE more generally) allows players to respond to out-of-equilibrium payoffs (which, by definition, equilibrium plus noise cannot do).

Such responses often allow QRE to track subjects’ initial responses to games significantly better than equilibrium plus noise does.

However, this strength is linked to a weakness:

QRE's ability to track subjects' responses is highly sensitive to the *form* of its assumed error distribution—far more sensitive than in quantal response models of individual decisions, or in other nonequilibrium models.

Haile et al. (2008) show that if there is only one observation per game-player pair and there are no restrictions on the error distributions, QRE, by varying the error distributions, can be made to fit any given dataset perfectly.

In other settings, QRE varying the error distributions stops short of allowing a perfect fit, but QRE is still unusually sensitive to assumptions about their forms.

Put another way, the source of QRE's ability to fit the data better is an untested assumption about which theory offers no guidance.

Remedying this may require estimating QRE decision distributions nonparametrically, with enough data from a sufficiently rich experimental design.

Existing analyses haven't done that, with partial exceptions discussed below.

QRE or LQRE decision rules are not self-informed, because a player's mean decision depends on other players' decision distributions.

As I have said, there is no *logical* problem with a non-self-informed model.

But a QRE or LQRE decision rule implicitly assumes that a player who, by hypothesis, cannot precisely identify her/his own equilibrium decision, can identify the precision of the distribution of others' *deviations* from equilibrium.

Spurred by Haile et al.'s (2008) and other criticisms, Goeree et al. (2005) and Goeree et al. (2019) consider the possible QRE decision distributions, replacing conventional functional form assumptions regarding error distributions, such as logit, with intuitive nonparametric qualitative restrictions.

Goeree et al.'s (2005) "regular QRE" or "RQRE" is a QRE with i.i.d. error distributions that are otherwise unrestricted except that their decision densities are monotonically increasing in a decision's expected payoffs (in two complementary ways) and satisfy continuity and a technical condition.

Goeree et al. (2005) characterize RQRE's set-valued restrictions on players decision distributions in a simple perturbed matching pennies example, with one observation per game-player pair, and some cross-game restrictions.

Thus, QRE with i.i.d. error distributions and plausible monotonicity restrictions on error densities does have some testable distribution-free implications.

Goeree et al. (2019) take this line of work a step further, introducing a notion they call “rank-dependent choice equilibrium” or “RDCE”.

Like RQRE, RDCE replaces QRE’s functional form assumptions regarding error distributions with plausible monotonicity conditions, now requiring that the rankings of a player’s decisions’ choice probabilities be the same as the rankings of their expected payoffs given other players’ decision distributions.

Goeree et al. characterize RDCE’s set-valued restrictions on players’ decision distributions, showing that RDCE’s restrictions are equivalent to the set of QRE’s restrictions with decision error distributions unrestricted except for monotonicity.

RDCE’s sets of possible decision probabilities, while possibly large, shrink quickly as the number of feasible decisions increases.

RDCE is cognitively less taxing than equilibrium or QRE, in that a player can find its sets of possible decision probabilities without fixed-point reasoning; but finding the QRE for particular error distributions still requires fixed-point reasoning.

Unlike QRE or LQRE decision rules, RQRE and RDCE rules are self-informed, because the means of their set-valued decisions are completely determined by iterated knowledge of players’ rationality and the structure of the game.

Dominance-*k* models ("*Dk*")

Costa-Gomes and Crawford (2006) introduce a family of decision rules called "Dominance-*k*" or "*Dk*", which makes epistemic analyses of finitely iterated dominance more precise and plays an important role in their analysis of Nagel's (1995) and Ho et al.'s (1998) *n*-person guessing games discussed below.

D1 does one round of deletion of both players' strictly dominated decisions and best responds to a uniform prior over its partner's remaining decisions.

D2 does two rounds of iterated deletion of strictly dominated decisions and best responds to a uniform prior over its partner's remaining decisions, and so on.

If there are no dominated decisions, a *Dk* rule is equivalent to a level-*k* *L1* rule.

In general, *Dk*-1's decisions survive *k* rounds of iterated elimination of dominated decisions and so in two-person games are *k*-rationalizable: *Dk*-1 yields a precise selection from the set of *k*-rationalizable decisions: a heterogeneity-tolerant refinement of *k*-rationalizability.

Dk decision rules are self-informed because their mean decisions are determined by iterated knowledge of players' rationality and the structure of the game.

Level-*k* models ("*Lk*")

Level-*k* decision rules anchor beliefs in a strategically naïve initial assessment of others' likely responses, called *L0*; and then adjust them via thought-experiments with iterated best responses, so that *L1* best responds to *L0*, *L2* to *L1*, and so on.

Level-*k* models bring in strategic considerations—while avoiding fixed-point reasoning—by assuming that *L1* best responds to *L0*, *L2* to *L1*, and so on.

A player's *L0* is usually assumed to be independent and uniform random over other players' feasible decisions.

In *n*-person games it matters whether *L0* is independent across other players or correlated. The limited available evidence suggests that most people have highly correlated, "representative agent"-like models of others. In the *n*-person games discussed below I take *L0* still to be uniform random but to pertain to the payoff-relevant summary of others' decisions, which there is their average decision.

Either way, *L0* players need not exist in the population: They may be only the starting point for higher levels' models of others.

Level-*k* models seem first to have been suggested by Keynes (1936, Chapter 12) at the end of his famous newspaper “beauty contest” analogy:

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

Keynes’ forms of strategic thinking are heterogeneous, discrete, and based on finitely iterated best responses to pure lower levels—like level-*k* decision rules.

Note that this definition of *Lk* best responds to a pure *Lk-1*, rather than:

- An *Lk-1* with decision noise (in the spirit of QRE)
- A mixture of lower levels (as in the CH models discussed next), or
- A mixture of levels with a positive frequency of its own level (which—like equilibrium—would require fixed-point reasoning)

This definition does make *Lk* cognitively more accessible than QRE, *Dk*, or CH.

The right *Lk* definition is an empirical question, and there is evidence for this one.

Lk decision rules (for $k > 0$) are rational; they depart from equilibrium only in that their beliefs are based on simplified, nonequilibrium models of others.

Lk decision rules, like *Dk*-1 rules (but unlike the CH *Lk* rules discussed below), respect *k*-rationalizability, making decisions that in two-person games survive *k* rounds of iterated deletion of strictly dominated strategies.

Dk-1 respects *k*-rationalizability by construction, while *Lk*'s respect is a by-product of iterating best responses.

Thus an *Lk* model, like a *Dk* model, can be viewed as a heterogeneity-tolerant refinement of *k*-rationalizability.

By basing beliefs on iterated best responses to a common *L0* anchor, an *Lk* model avoids many of the unrealistic belief helices that sometimes support extreme decision combinations as rationalizable or *k*-rationalizable.

In estimation, an *Lk* rule is usually assumed to have possibly payoff-sensitive decision errors, often taken to be logit as in equilibrium plus noise or LQRE.

Estimates suggest that there are positive frequencies only of *L1*, *L2*, and *L3.*

When the error structure is neutral regarding level frequencies (as in *Lk* models but not CH models; see below) the frequency of *L0* is usually estimated to be 0.

An estimated level-*k* model shares equilibrium's generality and some of its simplicity and tractability.

For given population level frequencies, it generically uniquely predicts the distribution of outcomes even in games with multiple equilibria.

Lk decision rules are self-informed because the means of their decisions are completely determined by iterated best responses and the structure of the game.

In a model with errors assumed to be i.i.d. across players, an *Lk* player may need knowledge of other players' decisions to estimate the distribution of her/his own errors, but not to estimate what is intentional, her/his own mean decisions.

Cognitive hierarchy models ("CH")

Camerer et al.'s (2004) CH model is a close relative of *Lk* models.

Like a level-*k* *Lk*, a CH *Lk* (for $k > 0$) ignores the possibility of others at the same level; but unlike a level-*k* *Lk* it best responds not only to *Lk-1* but to a mixture of all lower-level rules, including *L0*, again assumed to be uniform random.

The population distribution of levels is assumed to be Poisson over all levels including *L0*, with parameter τ; a CH *Lk*'s beliefs about level frequencies are derived by conditional Bayesian updating.

A CH *Lk* is assumed not to make decision errors. Instead the Poisson distribution of levels doubles as an error structure. Thus τ is the model's only parameter.

Estimates of τ, which equals the Poisson's average *k*, are often close to 1.5, which for the Poisson over all levels requires a population frequency of *L0* of 22%, far higher than unconstrained estimates with a conventional error structure.

(Some later work assumes that the Poisson's support includes only levels *L1* and higher, freeing up the frequency of *L0*, which seems more realistic.)

A CH *L1* is the same as a level-*k* *L1* and respects simple dominance.

Higher-level CH rules may differ from *Lk* rules and may violate *k*–rationalizability.

However, CH *Lk* beliefs (unlike level-*k* *Lk* beliefs) always become more accurate as *k* increases; but I don't assign much significance to limiting results as $k \to \infty$.

Like an *Lk* model, a CH model avoids some of the unrealistic belief helices that can support extreme decision combinations as rationalizable or *k*-rationalizable.

CH decision rules need not find fixed points or best respond to others' decision noise, and so share most of the cognitive ease of level-*k* rules.

However, CH rules are less accessible than level-*k* rules in that they best respond to an estimated mixture of lower-level rules.

A CH *Lk*'s best responding to a mixture of all lower levels may *seem* more realistic than a level-*k* *Lk*'s best responding to *Lk*-1, but which better describes behavior is an empirical question, and most subjects are not econometricians.

Like an estimated *Lk* model, an estimated CH model shares equilibrium's generality and some of its simplicity and tractability.

And for given population level frequencies, it generically uniquely predicts the distribution of outcomes even in games with multiple equilibria.

However, unlike an *Lk* decision rule, a CH rule is *not* self-informed, because its mean decisions for $k > 1$ depend on τ, the average of all players' levels, whose estimates involve other players' decisions.

As with QRE, a CH rule is not logically inconsistent, but it implicitly assumes players are sophisticated enough to predict the mean of other players' levels.

Summary of the leading rules' desiderata

	Precise	Set-valued
Self-Informed	**Equilibrium** ***Lk*** ***Dk***	**Rationalizability** ***k*-Rationalizability** **RQRE, RDCE**
Not Self-Informed	**QRE** **LQRE** **CH**	

On balance, theory and experimental results favor level-*k* models over the others; looking ahead to the experiments discussed next, regarding fit:

● *k*-rationalizability > rationalizability > equilibrium (as they must, logically); *k*-rationalizability's fit for low *k* seems enough better to justify its greater flexibility

● *k*-rationalizability > *Lk* (logically) >> *Dk*-1; *Lk*'s fit seems close enough to *k*-rationalizability's fit to justify *Lk*'s greater precision

● *Lk* >> equilibrium for a large majority of subjects

● *Lk* ≈ CH, LQRE; but *Lk* is self-informed while CH and LQRE are not; we have very little information about how *Lk*'s fit compares to RQRE's or RDCE's

Experimental Evidence

I now review some representative experimental evidence on strategic thinking.

The most informative such experiments have subjects play series of different games, with varying partners and without feedback to suppress learning and repeated-game effects and motivate subjects to respond to each game as if played in isolation: eliciting their “initial responses”.

The structures are publicly announced, to justify using the results to test theories of players’ decisions in complete-information versions of the games.

Subjects’ decisions are analyzed within subjects, assuming that a single decision rule determines a subject’s decision in each of the games s/he plays.

The games are chosen to separate the decision profiles implied by equilibrium and other rules as strongly as possible. This makes a subject’s profile a strategic “fingerprint,” which often identifies her/his rule with great precision.

n-person guessing games

Nagel (1995), Ho, Camerer, and Weigelt (1998; "HCW"), and Bosch-Domènech et al. (2002) experimentally studied *n*-person guessing games directly inspired by Keynes' beauty contest analogy in the passage quoted above.

In Nagel's and HCW's games, *n* subjects (15-18 in Nagel, 3 or 7 in HCW) made simultaneous guesses between lower and upper limits (0 to 100 in Nagel, 0 to 100 or 100 to 200 in HCW). In Bosch-Domènech et al. the same games were played by more than 7500 volunteers recruited from subscribers of the newspapers *Financial Times*, *Spektrum der Wissenchaft*, or *Expansión*.

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

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

(Nagel's and HCW's subjects each played a game repeatedly, but their first-round guesses can be viewed as initial responses to a game if they treated their own influences on future guesses as negligible, which is plausible for all but HCW's 3-subject groups. Bosch-Domènech et al.'s subjects played only once.)

In Nagel's leading treatments:

- 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

Nagel's games are dominance-solvable, with a unique equilibrium (ignoring discreteness) that can be found by iteratively eliminating dominated guesses.

If $p = 1/2$:

- It's dominated to guess more than 50 (because $1/2 \times 100 \leq 50$).
- Unless you think that other people will make dominated guesses, it's also dominated to guess more than 25 (because $1/2 \times 50 \leq 25$).
- And so on, down to 12.5, 6.25, 3.125, and eventually to 0.

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

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

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

Nagel's subjects never made equilibrium guesses initially; HCW's rarely did, and Bosch-Domènech et al.'s (who had more time and could consult) rarely did.

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

Nagel's Figure 1 (top p = 1/2, bottom p = 2/3)

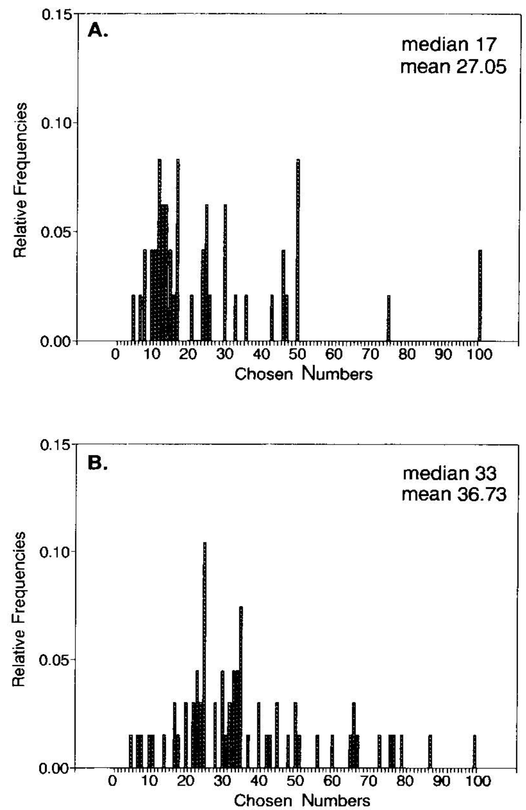

Bosch-Domènech et al.'s Figure 1.

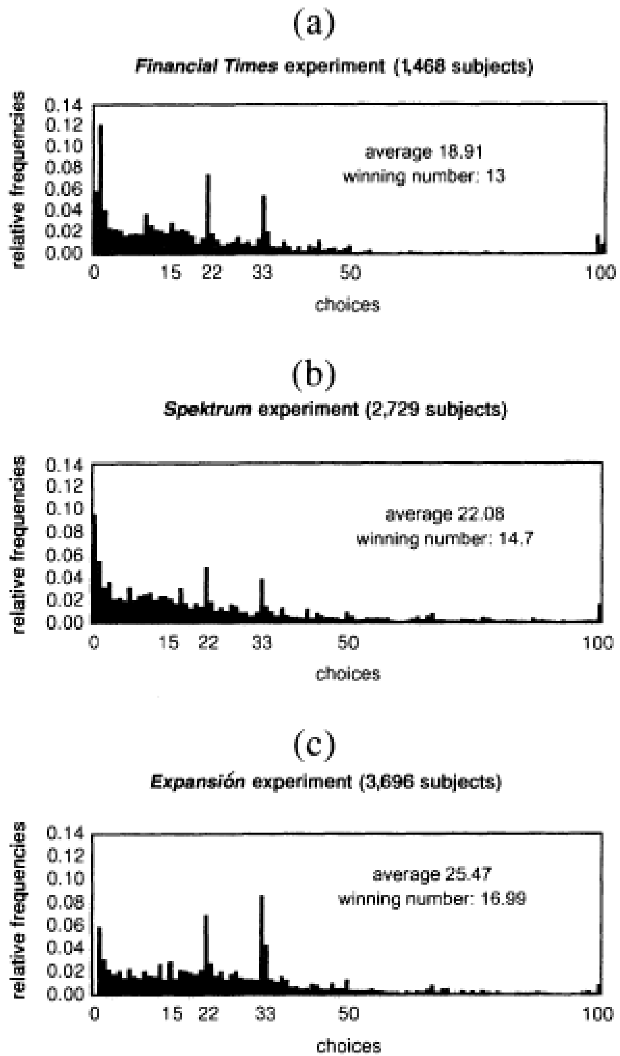

FIGURE 1. RELATIVE FREQUENCIES OF CHOICES IN THREE NEWSPAPER EXPERIMENTS

These data resemble neither equilibrium plus noise nor "equilibrium taking noise into account", as in a QRE—that is, not for any reasonable error distribution, though by Haile et al's (2008) result we could make the data an exact QRE for a carefully contrived but *un*reasonable distribution.

The data do suggest that deviations from equilibrium have a coherent structure.

The guess distributions have spikes that track $50p^k$ for $k = 1, 2, 3$ across treatments with different *p*s, respecting up to 3 rounds of iterated dominance.

Like spectrograph peaks that reflect chemical elements, the spikes are evidence of a partly deterministic, discrete, and individually heterogeneous structure.

Compare Keynes's "We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees."

Nagel's, HCW's, and Bosch-Domènech et al.'s results show clearly that, even in such simple games, subjects' initial responses do *not* reflect fixed-point reasoning or extensively iterated dominance or even best responses.

Nagel and HCW interpreted their spikes as evidence of level-*k* decision rules:

Lk guesses $[(0+100)/2]p^k \equiv 50p^k$.

But many theorists interpret the spikes as subjects performing finitely iterated deletion of dominated decisions.

Costa-Gomes and Crawford's (2006) *Dk* rules, in which a player does *k* rounds of iterated dominance and then best responds to a uniform prior over other players' remaining decisions (or here, their average), are a natural way of making finitely iterated dominance yield a unique decision:

Dk-1 guesses $([0+100p^{k-1}]/2)p \equiv 50p^k$.

(By an unimportant quirk of my notation, it is *Dk*-1 that is *Lk*'s cousin, not *Dk*.)

Which is it? *Dk*-1 and *Lk* are weakly separated in some other experiments, but we need to know more.

Two-person guessing games

Nagel's and HCW's large strategy spaces greatly increase the informativeness of subjects' decisions, but it is a weakness that each subject played only one game.

Although each subject played her/his game repeatedly, later choices confound strategic thinking with learning, so there's only one real observation per subject.

Even so, if subjects treat their own influences on others' future decisions as negligible, first-round choices can be viewed as initial stage-game responses.

But Nagel's and HCW's between-subjects variation across treatments is much less informative about thinking than within-subjects variation across games.

Costa-Gomes and Crawford's (2006; "CGC") design increases power by combining large strategy spaces with Stahl and Wilson's (1994, 1995) and Costa-Gomes, Crawford, and Broseta's (2001) series of games.

(CGC also supplemented their analysis of subjects' decisions by monitoring and analyzing their searches for hidden but freely accessible payoff information; Crawford 2008 gives the details.)

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

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

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

A subject's guesses in the 16 games form a strategic "fingerprint" that comes close to revealing a subject's decision rule without an econometric "middleman".

Another advantage is that in two-person guessing games, subjects know that other subjects don't view them as a negligible part of the interaction; thus these games more fully engage and identify subjects' strategic thinking.

In CGC's games, each player has his own lower and upper limit, both strictly positive (which makes the games finitely dominance-solvable).

(Guesses outside a subject's limits are automatically adjusted up to the lower or down to the upper limit as necessary: This is a trick to enhance the separation of information search implications, which is not important for this discussion.)

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

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

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

CGC’s games have essentially unique equilibria (“essentially” due solely to the automatic adjustment), determined 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 is one enhances the separation of equilibrium from other decision rules.

This also stress-tests equilibrium in that CGC’s design includes games that differ in whether the product of targets is slightly greater or slightly less than one.

Equilibrium responds to such differences much more strongly than behaviorally plausible nonequilibrium decision rules do.

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]. (What would you guess?)

The product of targets 1.05 > 1 so equilibrium is determined by the upper limits.

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

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

Any guess outside [315, 500] is then iteratively dominated for the first player.

Any guess outside [472.5, 750] is then iteratively dominated for the second.

And so on until the equilibrium at (500, 750) is reached after 22 iterations.

(Cognitively and behaviorally, 22 ≈ ∞.)

CGC's data analysis assumed (with testing) that each subject's guesses were determined, up to logit errors, by a single decision rule in all 16 games.

The main analysis restricted attention to an a priori list of plausible rules:

- *L0*, *L1*, *L2*, and *L3* as defined above, with *L0* uniform random
- *D1* and *D2* as defined above
- *Equilibrium*, which makes its equilibrium decisions
- *Sophisticated*, which best responds to the probability distributions of others' decisions, estimated in CGC's data analysis from the observed frequencies

CGC tested the restriction to this list of decision rules and found it to be a reasonable approximation to the population support of subjects' rules.

CGC's large strategy spaces and independent variation of targets and limits across games very strongly separate decision rules' implications.

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

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

Of the 88 subjects in CGC's main treatments, 43 made guesses that complied *exactly* (within 0.5) with one type's guesses in from 7 to 16 of the games: 20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*; with 200 to 800 possible exact guesses in 16 different games, exact compliance is very powerful evidence for a decision rule.

(Crawford 2008 questions the subjects with high *Equilibrium* compliance.)

CGC's results clearly identify subjects' rules, and strongly favor *Lk* over *Dk*-1: Subjects respect *k*-rationalizability for small *k* not because they explicitly perform iterated dominance, but because they follow rules that implicitly respect it.

By contrast, in a small matrix game there can be many possible reasons for choosing a strategy; and even in Nagel's and HCW's large-strategy-space games rules as cognitively disparate as *Lk* and *Dk*-1 yield identical decisions.

Further, because CGC's rules build in risk-neutral, self-interested rationality, *L1*, *L2*, and *L3* subjects' deviations from equilibrium can be traced to their simplified models of others, not to irrationality, risk aversion, altruism, spite, or confusion.

The analysis supports a uniform random *L0*, which may reflect sampling payoffs initially without considering others' incentives or principle of insufficient reason.

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

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

And econometric estimates of their rules are concentrated on *L1*, *L2*, *L3*, and *Equilibrium* in roughly the same proportions. (Analysis of subjects' searches for hidden but freely accessible payoff information confirms and sharpens this.)

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

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

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

CGC's Figure 1.

Lessons for modeling strategic thinking

- Strategic thinking is *heterogeneous*: There are significant numbers of *L1*, *L2*, *L3*, and *possibly* (as explained in Crawford 2008) *Equilibrium* subjects

- There are few if any *Dk* or *Sophisticated* subjects and few if any *L4* or higher level *Lk* subjects; thinking seldom directly relies on extensively iterated dominance or even best responses, or relies on fixed-point reasoning

- In non-dominance solvable games or *k*-dominance-solvable games for high values of *k*, people's decisions deviate systematically from equilibrium; a simple class of models with an estimated or calibrated distribution of *L1*, *L2*, *L3*, and *Equilibrium* players, then shares the generality and much of the tractability of equilibrium analysis and can systematically out-predict equilibrium

Outguessing in games with symmetric information: *Far Pavilions* Escape

The following example illustrates how a level-*k* model works in the simplest possible outguessing game, and the interpretation of *L0*.

(It's interesting to compare this example with the outguessing game between Holmes and Moriarty in Conan Doyle's story "The Final Problem", which von Neumann and Morgenstern (1944) used to illustrate mixed-strategy equilibrium.)

Early in M.M. Kaye's historical novel *The Far Pavilions*, the main male character, Ash/Ashok, is trying to escape from his Pursuers along a north-south road.

Ash and his pursuers must choose between North and South. Although Ash moves first in time, his Pursuers must make their choice before they learn Ash's choice, so the two players' choices are strategically simultaneous.

South is warm, while North are the Himalayas, with winter coming.

If the Pursuers catch Ash, they gain two units of payoff and Ash loses two.

In addition, both the Pursuers and Ash gain one extra unit for choosing South.

This yields the payoff matrix:

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

***Far Pavilions* Escape**

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

Ash's mentor Koda Dad (played by Omar Sharif in the HBO miniseries, but unfortunately this scene is only in the book, p. 97) advises Ash to ride north:

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

Ash follows his mentor's advice, the Pursuers unimaginatively go south, and Ash escapes—to live for another 900 pages.

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

(*L3* ties my personal best *k* for a clearly explained level-*k* rule in fiction. I suspect that no higher level would be credible to a general audience.

And I dare you to try to explain the fixed-point reasoning that underlies this game's mixed-strategy Nash equilibrium to a general audience.)

Pursuers who expect Ash to go south simply because it's "kinder" must be modeling Ash as an *L1* responding to a uniform random *L0*.

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

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

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

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

In a more complex game, a uniform random *L0* could approximate random sampling of an *L1* player's decisions' payoffs, unstratified by the other player's incentives to make the decisions.

Level-k versus equilibrium analysis in Escape

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

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

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

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

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

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

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

Rule	**Ash**	**Pursuers**
L0	Uniform random	Uniform random
L1	South	South
L2	North	South
L3	North	North
L4	South	North
L5	South	South

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

The level-*k* model correctly predicts the outcome provided that Ash was either *L2* or (as we know) *L3* and Pursuers were either *L1* or (as Koda Dad believed) *L2*.

If we don't have an omniscient narrator telling us how players think, but the game is well-defined and we have data, we can specify a level-*k* model, derive its implications, and estimate the rule frequency distribution.

Or we can calibrate the model using estimates from analogous settings.

Suppose, for example, we assume that each player role in Escape is filled from an (empirically plausible) 50-30-20 mixture of *L1*s, *L2*s, and *L3*s, with no errors.

Then the level-*k* model predicts that Ash goes North with probability 1/2 (= 0.30 + 0.20) and Pursuers go South with probability 4/5 (= 0.50 + 0.30)

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

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

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

Yet experimental subjects' aggregate choices in initial responses to such games follow decision-theoretic intuition in both roles.

(Even Ash's counterintuitive choice would not contradict this pattern if he were a subject, because his revealed level is in the minority.)

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

Market-entry games with symmetric information

As suggested by Camerer et al.'s (2004, Section III.C) CH analysis of market-entry games, a level-*k* analysis suggests a view of tacit coordination in games like Battle of the Sexes that is radically different than an equilibrium view.

Subjects in market-entry experiments come close, on average, to exactly filling market capacity, which led Kahneman (1988) to remark, "…to a psychologist, it looks like magic".

But as Camerer et al. note, there is also something that would look like magic even to a game theorist: Ssubjects regularly achieve better ex post coordination (number of entrants stochastically closer to market capacity) than in the natural equilibrium benchmark, the symmetric mixed-strategy equilibrium.

Camerer et al. give a unified CH explanation of both kinds of magic: The predictable heterogeneity of players' strategic thinking allows some of them to mentally simulate others' entry decisions and accommodate them, breaking the symmetry as required for coordination.

Higher-level players in effect become like Stackelberg followers, with coordination benefits for all.

I adapt Camerer et al.'s CH analysis to a level-*k* analysis of two-person Battle of the Sexes, which illustrates the key points. Assume that $a > 1$.

		Column	
		H	**D**
Row	**H**	**0** **0**	**1** ***a***
	D	***a*** **1**	**0** **0**

Battle of the Sexes

When players cannot distinguish each other or communicate, it can be shown that the asymmetric equilibria are effectively unplayable.

The unique symmetric equilibrium is in mixed strategies, with $p \equiv \Pr\{H\} = a/(1+a) > ½$ (when $a > 1$) for both players.

The expected coordination rate is $2p(1–p) = 2a/(1+a)^2$; and players' 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 levels, *L1*, *L2*, *L3*, or *L4*, with each player role filled by a draw from the same distribution, which restricts attention to symmetric outcome distributions.

As in most previous analyses, I assume that *L0* chooses its decision randomly, with Pr{H} = Pr{D} = ½.

Higher levels' best responses are easily calculated:

*L1*s mentally simulate *L0*s' random decisions and best respond to them, choosing H; similarly, *L2*s choose D, *L3*s choose H, and *L4*s choose D.

Although *L3* behaves like *L1* here, and *L4* like *L2*, I retain all four for comparability with the analysis below; but I set the frequency of *L0* to 0.

The model's predicted outcome distribution is determined by the outcomes of the possible type pairings in Table 1 and the level frequencies:

Rules	***L1***	***L2***	***L3***	***L4***
L1	H, H	H, D	H, H	H, D
L2	D, H	D, D	D, H	D, D
L3	H, H	H, D	H, H	H, D
L4	D, H	D, D	D, H	D, D

Table 1. Level-*k* Outcomes

The rule frequencies are assumed independent of payoffs, in keeping with the fact that they are intended as general models of strategic behavior; as a result, the model's predicted outcome distribution (without errors) is independent of *a*.

With symmetry across roles, players have equal ex ante payoffs, which are proportional to the expected coordination rate.

Lumping *L1* and *L3* together and letting v denote their total probability, and lumping *L2* and *L4* together and letting $(1-v)$ denote their probability, the coordination rate is $2v(1-v)$, maximized at $v = ½$ where it takes the value ½.

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

The mixed-strategy equilibrium coordination rate, $2a/(1+a)^2$, is maximized when $a = 1$ where it takes the value ½, but it converges to 0 like $1/a$ as $a \to \infty$.

Even for moderate values of a, the level-*k* coordination rate is quite likely to be higher than the symmetric equilibrium rate.

The level-*k* model improves upon the symmetric equilibrium by "relaxing" the incentive constraints requiring players' responses to be in equilibrium.

Because level-*k* levels best respond to nonequilibrium beliefs, it is natural to compare the level-*k* outcome to the best symmetric rationalizable outcome, in which each player plays a nonequilibrium mixed strategy with $v \equiv \Pr\{H\} = ½$.

When $v = ½$, the level-*k* model can be viewed as using the heterogeneity of strategic thinking to purify this best symmetric rationalizable outcome. With levels exogenous, level-*k* thinking may not make this outcome attainable.

The level-*k* analysis suggests a view of coordination that is profoundly different than the equilibrium view:

Although players' decisions are simultaneous and there is no communication, the predictable heterogeneity of strategic thinking allows some players to mentally simulate others' entry decisions and accommodate them, just as (noisy) Stackelberg followers would, with coordination benefits for all players.

Equilibrium and selection principles like risk- or payoff-dominance play no role in level-*k* players' thinking; and coordination, when it occurs, is an accidental (but statistically predictable) by-product of their non-equilibrium decision rules.

Outguessing with symmetric information and communicating intentions

(Crawford 2003 and Crawford, Costa-Gomes, and Iriberri 2013, Section 9.1; I focus here on deceptive communication with opposing interests; on coordination see Ellingsen and Östling 2010 and Crawford 2017).

D-Day

The Allies chose where to invade Europe on D-Day: Calais or Normandy.

- Invading an undefended Calais is better for the Allies than invading an undefended Normandy (but the Germans know that too)

- Similarly, defending an uninvaded Normandy is worse for the Germans than defending an uninvaded Calais

Before the attack, the Allies placed a fake invasion army in the Thames Estuary: an approximately cheap talk message regarding their intentions, with an obvious literal meaning, Calais (http://en.wikipedia.org/wiki/Operation_Fortitude).

The Allies invaded at Normandy, while the Germans believed the message (or perhaps inverted it one too many times) and persistently overdefended Calais.

A “Tank” from Operation Fortitude South in the Thames Estuary, 1944

Huarongdao (from Luo Guanzhong's historical novel *Three Kingdoms*)

Defeated, fleeing General Cao Cao chose between two escape routes, the easy Main Road and the awful Huarong Road, trying to evade pursuing General Kongming (http://chinesepuzzles.org/huarong-pass-sliding-block-puzzle/).

Other things equal, both generals preferred the Main Road.

Kongming waited in ambush along Huarong Road and set campfires there, sending an approximately cheap talk message with an obvious literal meaning.

Cao Cao expected a lie, inverted the message, and was caught on Huarong Road.

Puzzles

In both D-Day and Huarongdao, the message-sender deceived the message-receiver and won, but in the *less* beneficial of the two possible ways.

- Why did the receiver allow himself to be fooled by an easily faked, approximately cheap talk message from an *enemy*?

- And if the sender expected his deception to succeed, why didn't he reverse the message and 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?)

An analysis should also reconcile the answers to these puzzles with the surface differences in the senders' messaging strategies and the receivers' responses:

- The Allies' messages was literally a lie, but it was sent with the belief that the Germans would believe it or—perhaps expecting a double bluff—invert it one too many times

- Kongming's message was literally true, but Cao Cao, expecting a lie, was fooled by Kongming's double bluff into inverting the message

In this case Luo Guanzhong tells us what his fictional generals were thinking:

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

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

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

Equilibrium in D-Day and Huarongdao

- In each example, one player, the Sender, sends a nonbinding, cheap talk message to the receiver regarding his planned action; lying has no direct cost

- The other player, the Receiver, observes the message

- The Sender and Receiver make decisions that jointly determine their payoffs

- The Sender's and Receiver's payoffs are (at least approximately) opposed

In games like these, in any equilibrium the Sender's cheap talk message must be uninformative, and the Receiver must ignore it.

If the Sender made his message informative, the Receiver's optimal response to it would increase the Receiver's own payoff and thus reduce the Sender's, who would therefore do better by making his message uninformative.

Equilibrium behavior therefore makes the communication phase irrelevant.

Yet in the world, deception is common and succeeds, even in zero-sum games.

Level-*k* thinking in D-Day and Huarongdao

Although the experimental evidence from games without communication is broadly consistent with a uniform random *L0*, intuition and the limited evidence suggest that in games with communication, the first thing we do when hearing a message, even from an enemy, is to try to understand its literal meaning.

This suggests anchoring *L0* in truthfulness for Senders or credulity for Receivers.

Higher levels can be defined by iterated best responses as before.

(The level-*k* analysis assumes no direct ad hoc lying costs.)

The Luo Guanzhong quotations suggest that Kongming was *L3* and Cao Cao *L2* with a truthful or respectively credulous *L0*.

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

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

“Behavioral equilibrium” in D-Day and Huarongdao

Assume the Sender and Receiver are each drawn from a population including both level-*k* (called “mortal” in Crawford 2003) and *Sophisticated* players.

- Level-*k* players avoid fixed-point reasoning, anchor beliefs on truthfulness or credulity, and determine beliefs and decisions by iterated best responses

- *Sophisticated* players choose equilibrium decisions in a reduced game that reflects the possibility and frequencies of level-*k* and *Sophisticated* players

Analyze the reduced game between Sophisticated Senders and Receivers, with level-*k* players’ mechanical and uninteractively predictable responses plugged in.

The main goal is to learn whether and when the possibility of level-*k* players in each role allows *Sophisticated* Senders to “deceive” *Sophisticated* Receivers.

The underlying game is approximately zero-sum, with symmetric information and cheap talk messages.

The fact that some players might be level-*k* makes the character of the reduced game between *Sophisticated* Senders and Receivers completely different:

- As *Sophisticated* players' payoffs are influenced by level-*k* players' decisions, the reduced game is not zero-sum and its messages are not cheap talk

- The reduced game has asymmetric information about the Sender's decision rule, and a *Sophisticated* Receiver reads the Sender's message about his intentions as an informative signal about his rule, not (directly) his intentions

If *Sophisticated* Senders and Receivers have high frequencies, even well below 1, the reduced game has a unique mixed-strategy equilibrium, which is outcome-equivalent to the equilibrium of the underlying game without communication. *Sophisticated* players then protect their level-*k* counterparts from exploitation.

If *Sophisticated* Senders and Receivers have low frequencies, the reduced game has an essentially unique pure-strategy equilibrium, in which *Sophisticated* Senders send the message that deceives the most frequent kind of level-*k* Receiver and then try for the *less* beneficial way to win: D-Day and Huarangdao!

There is never a sensible equilibrium in which *Sophisticated* Senders try for the *more* beneficial way to win.

For, in such an equilibrium any deviation from the *Sophisticated* Sender's equilibrium message would "prove" to a *Sophisticated* Receiver that the Sender is level-*k*, leading a *Sophisticated* Receiver to try for the less beneficial way to win, and thus leading a *Sophisticated* Sender to try for the less beneficial way.

Thus, with no unexplained difference in the sophistication of Senders and Receivers, and for plausible parameter values, the level-*k* model explains why *Sophisticated* Receivers might allow themselves to be "deceived", and why *Sophisticated* Senders don't try for the more beneficial way to win.

Communicating private information

(Wang, Spezio, and Camerer 2010 ("WSC"); based on Crawford and Sobel 1982 and Crawford 2003.) WSC's "Blodget" frame:

During the tech-stock bubble, Wall Street security analysts were alleged to inflate recommendations about the future earnings prospects of firms in order to win investment banking relationships with those firms. Specifically, analysts of Merrill Lynch used a five-point rating system (1 = Buy to 5 = Sell) to predict how the stock would perform. They usually gave two 1–5 ratings for short run (0–12 months) and long run (more than 12 months) performance separately.

Henry Blodget, Merrill Lynch's famously optimistic analyst, "did not rate any Internet stock a 4 or 5" during the bubble period (1999 to 2001). In one case, the online direct marketing firm LifeMinders, Inc. (LFMN), Blodget first reported a rating of 2-1 (short run "accumulate"—long run "buy") when Merrill Lynch was pursuing an investment banking relationship with LFMN. Then, the stock price gradually fell from $22.69 to the $3–$5 range. While publicly maintaining his initial 2-1 rating, Blodget privately e-mailed fellow analysts that "LFMN is at $4. I can't believe what a POS [piece of shit] that thing is." He was later banned from the security industry for life and fined millions of dollars.

(Source: Complaint, Order, and Final Judgement in Securities and Exchange Commission v. Henry M. Blodget, (2003) Civ. 2947 (WHP) (S.D.N.Y.).)

Equilibrium (without lying costs) in Blodget

- In each case one player, the Sender, observes a signal relevant to both their payoffs and sends a cheap talk message about it; lying has no direct cost
- The Receiver observes the message and makes a decision that, with the Sender's true signal, determines both Sender's and Receiver's payoffs
- The Sender's and Receiver's preferences about how the Receiver's decision relates to the Sender's signal are qualitatively similar, but differ systematically

Crawford and Sobel showed that in all equilibria, the Sender divides the possible signals into intervals and tells the Receiver only which interval the signal fell in.

There is always a "babbling" equilibrium; and when the Sender's and Receiver's preferences are sufficiently far apart it is unique, except for messaging variations.

When the Sender's and Receiver's preferences are closer together there are also more informative equilibria, but all have some intentional vagueness.

The closer the Sender's and Receiver's preferences, the more information is transmitted in the most informative equilibrium.

Puzzles

- Why would anyone be systematically fooled by a cheap talk message from someone whose interests are different?

- Why do Senders tend to lie in the direction that would push credulous Receivers in their favored direction?

- Why are Senders more truthful and Receivers more credulous than in any equilibrium?

None of these puzzles are resolved by the equilibrium analysis.

WSC's experimental design

WSC's experimental design closely follows their Blodget example.

- A Sender observes the state, S = 1, 2, 3, 4, or 5, and sends a message, M = 1, 2, 3, 4, or 5 (the clear correspondence between state and message labelings ensures that messages are understood and makes lying a meaningful concept)

- A Receiver observes the message M, and chooses an action A = 1, 2, 3, 4, or 5, which together with S determines his and the Sender's welfare

- Senders and Receivers have single-peaked preferences, assuming no lying costs, with the Receiver's ideal outcome A = S and the sender's A = S + b (ignoring boundaries)

- The design varies the difference between Sender's and Receiver's preferences across three treatments: b = 0, 1, or 2

Equilibrium versus level-*k* thinking (without lying costs) in Blodget

WSC's equilibrium and level-*k* analyses assume no direct ad hoc lying costs.

WSC's equilibrium analysis focused on the most informative equilibria in the games created by their three treatments, as benchmarks.

In WSC's level-*k* analysis, subjects' excessive (relative to equilibrium) truthfulness and credulity are explained as a residue of level-*k* anchoring.

(In an equilibrium analysis without lying costs, the effects of such anchoring would be completely massaged away, but *Lk*'s finitely iterated best responses do not completely massage them away.)

In WSC's Figures 1-3, a circle's size shows Senders' message frequencies (columns) in the various states (rows) and a circle's darkness and the numbers inside it show Receivers' action frequencies.

- In Figure 1 the Sender's and Receiver's preferences are identical (b = 0); the most informative equilibrium has truth-telling and credulity: M = S and A = S

  There are no significant deviations from the most informative b = 0 equilibrium.

- In Figure 2 the Sender's and Receiver's preferences differ somewhat (b = 1); the most informative equilibrium has the Sender sending M = 1 when S = 1 and the Receiver responding with A = 1; and otherwise the Sender's message distribution is the same for S = 2, 3, 4, 5, and the Receiver responds A = 3 or 4

  Both Senders and Receivers deviate systematically from the most informative b = 1 equilibrium.

  Senders lie in the direction (above the diagonal) that would make credulous Receivers choose actions Senders would prefer, while making messages more truthful than in the equilibrium (M distributions shift right as S goes from 2 to 5).

  And Receivers are more credulous (A > S, A > best response to Senders).

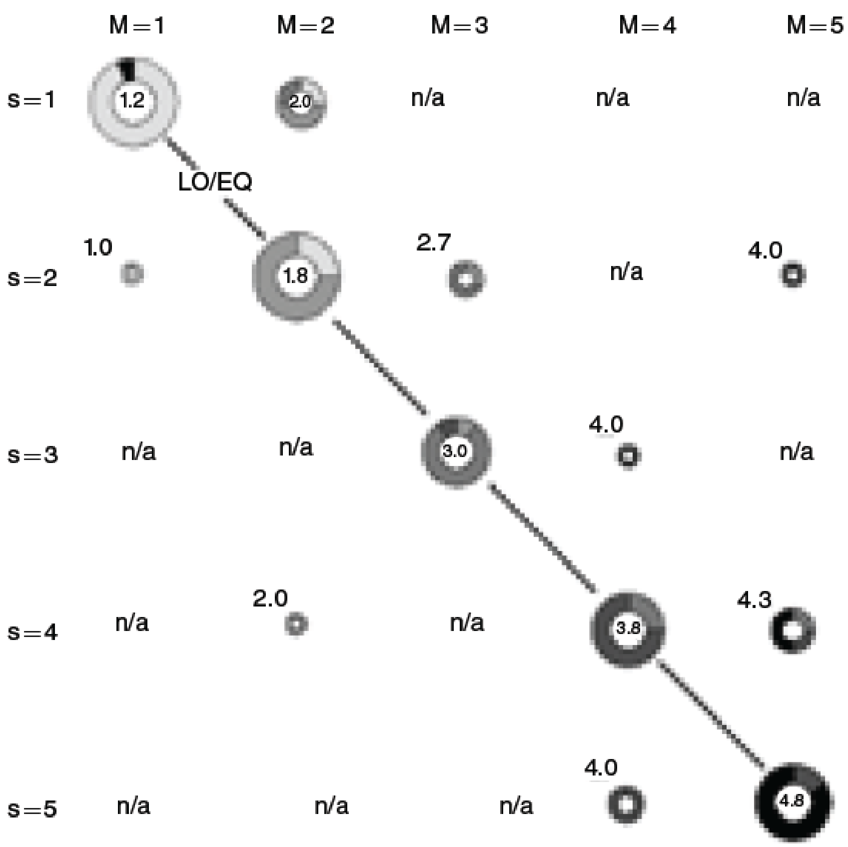

FIGURE 1. RAW DATA PIE CHARTS ($b = 0$)
(HIDDEN BIAS–STRANGER)

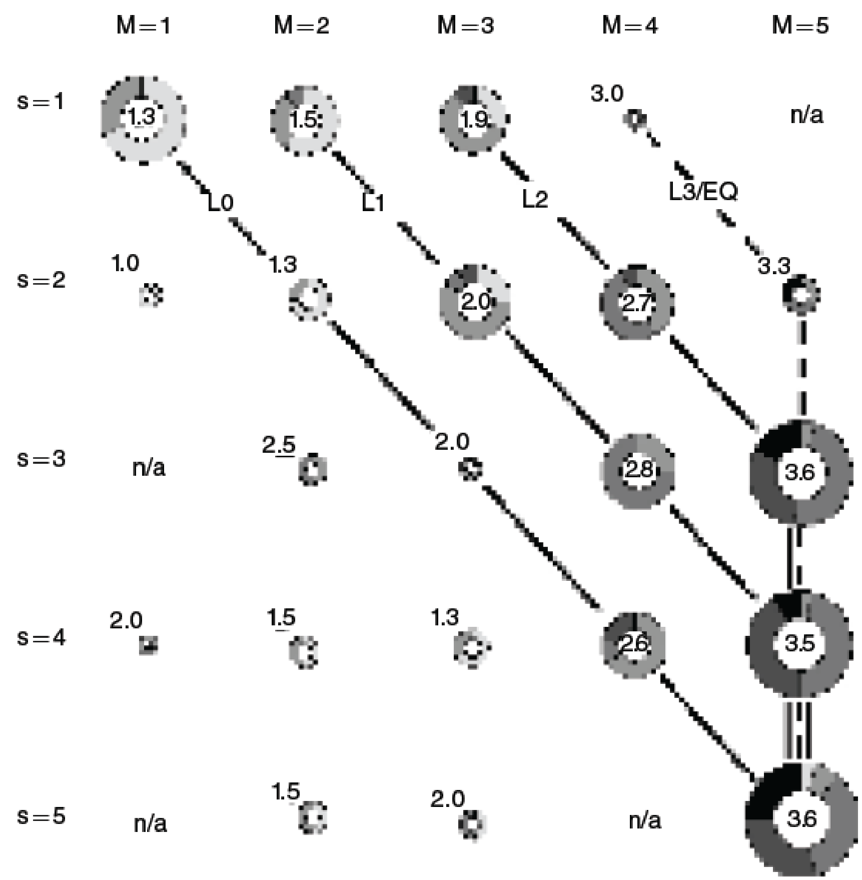

FIGURE 2. RAW DATA PIE CHART ($b = 1$)
(HIDDEN BIAS–STRANGER)

- In Figure 3 the Sender’s and Receiver’s preferences differ a great deal (b = 2); the only equilibria are babbling equilibria, in which the Sender’s message distribution is the same for all S, and the Receiver ignores the Sender’s messages and chooses A = 3, the optimal action given the Receiver’s prior

  Both Senders and Receivers deviate systematically from the most informative b = 2 equilibrium.

  Senders again lie in the direction (above diagonal) that would make credulous Receivers choose actions sender would prefer, while making messages more truthful than in equilibrium (M distributions shift right as S goes from 1 to 5).

  Receivers are again more credulous (A > S: > the best response to Senders).

- Despite the systematic deviations from equilibrium when b = 1 or 2, the amount of information transmitted, measured by the correlation between S and A, declines with the distance between Sender’s and Receiver’s preferences, as suggested by Crawford-Sobel’s equilibrium-based comparative statics result

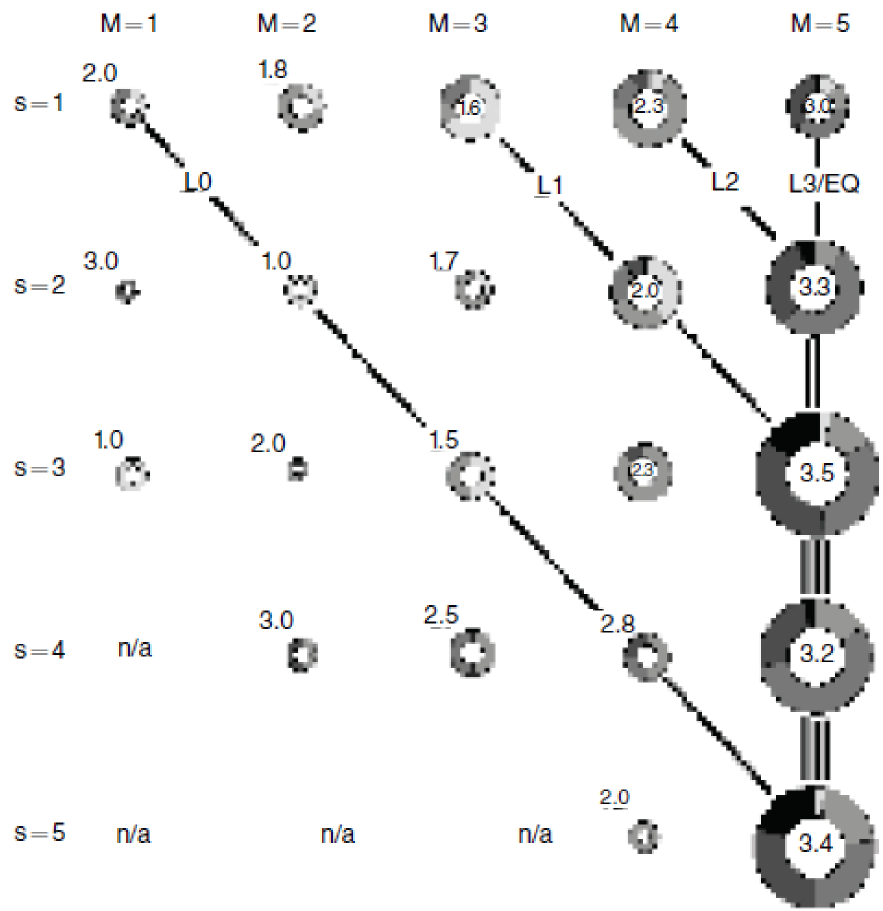

FIGURE 3. RAW DATA PIE CHART ($b = 2$)
(HIDDEN BIAS–STRANGER)

Following Crawford's 2003 model of level-*k* communication of intentions, WSC analyzed their results using a level-*k* model of communication of private information, again assuming no direct lying costs. A Receiver's best outcome is then A = S, ignoring boundaries, here and below; and a Sender's is A = S + b.

- In the level-*k* model, players anchor beliefs in a truthful Sender *L0*, which sets M = S; and a credulous Receiver *L0*, which sets A = M

- *L1* Senders best respond to *L0* Receivers, "puffing" their messages by b: M = S + b, so *L0* Receivers choose S + b, which would yield an *L1* Sender's best action, given a credulous Receiver

- *L1* Receivers best respond to *L1* Senders, de-puffing messages by b: A = M – b, which would yield an *L1* Receiver's ideal action, given her/his belief that *L1* Senders best respond to *L0* Receivers, setting M = S + b

- *L2* Senders best respond to *L1* Receivers, puffing by 2d: M = S + 2b; *L2* Receivers best respond to *L2* Senders, de-puffing by 2b: A = M – 2b; etc.

The labels in Figures 1-3 show a close association between Senders' and Receivers' decisions and *L1*, *L2*, or *L3* behavior. (These are explicitly labeled in Figures 2 and 3; *L1*, *L2*, or *L3* behave the same as equilibrium in Figure 1.)

Overall, the level-*k* model gives a unified explanation of the main fact patterns:

● Senders lie in the direction that would make credulous Receivers choose actions the Sender would prefer, trying to outguess Receivers' discounting

● Senders' messages are nonetheless more truthful than in any equilibrium

● Receivers' responses are more credulous than in any equilibrium

Even though the model assumes lying has no direct costs, *Lk* behavior is anchored in a truthful or credulous *L0*; this gives *Lk* a residue of truthfulness or credulity that only equilibrium reasoning would completely massage away.

The sensitivity of *Lk*'s behavior to the distance between Sender's and Receiver's preferences also explains why Crawford-Sobel's equilibrium comparative statics result is qualitatively robust even to large, systematic deviations from equilibrium.

Asymmetric information and informational naïveté

A large literature documents and analyzes the phenomenon of the winner's curse, whereby bidders who win common-value auctions tend to find that they have bid more than the object is worth.

A related literature studies the phenomenon of zero-sum betting and its analogues in asset markets, whereby people bet or trade more than is consistent with any equilibrium.

Both phenomena have their roots in informational naïveté, whereby people do not fully take into account the dependence of others' decisions on their own private information, which would allow inferences that improve their decisions.

The classic equilibrium-based analyses of informational naïveté are Milgrom and Weber (1982) on auctions and Milgrom and Stokey (1982) on speculative trading.

Milgrom and Weber show (among many other things) how to correct one's bid in common-value auctions for the information about the object that will be contingently revealed (through other bidders' lower value estimates) by winning.

Milgrom and Stokey show that if traders in an asset market start out in a market equilibrium—Pareto-efficient given their information—giving them new private information cannot make new trades mutually beneficial.

For, any such new trades would make it common knowledge that both traders had benefited from it, contradicting the efficiency of the initial equilibrium.

Milgrom and Stokey's no-trade result was dubbed by Milt Harris the "Groucho Marx Theorem":

"I sent the club a wire stating, 'Please accept my resignation. I don't want to belong to any club that will accept people like me as a member'."
—Groucho Marx (1959, p. 321), Telegram to the Beverly Hills Friar's Club

Milgrom and Stokey's result might also be called the "Sky Masterson Theorem":

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

—Obadiah ("The Sky") Masterson, quoting his father in "The Idyll of Miss Sarah Brown" (*Guys and Dolls: The Stories of Damon Runyon*, 1932)

(Unlike Groucho's quip, the Runyon story implicitly questions Savage's "small worlds" assumption, which seems an important part of real zero-sum betting.)

By contrast with Milgrom and Stokey's no-trade theorem, speculative zero-sum trades are common in real financial markets and real life.

Here's an example from the All Souls College betting book (many others are possible). [My scholia are in square brackets.]

[Warden] Davis bets [Fellow whose study overlooked the Gt Quad] Ryan 1 doz oysters that the lawn in the Gt Quad is more circular than it is gibbous or oval.

Adjudicant[s]: [Fellow] Perkins, [Fellow] Häcker

[Signed]: J Davis MJ Ryan
1.X.05 [a Saturday]

[To paraphrase former U.S. Supreme Court Justice Potter Stewart, bettors and adjudicants seem to have thought they would know eccentricity when they saw it.

And so they did.]

Ryan wins
BH
J Perkins

The bettors’ information about the shape of All Souls’s Great Quad:

Warden Davis’s View

Fellow Ryan’s View

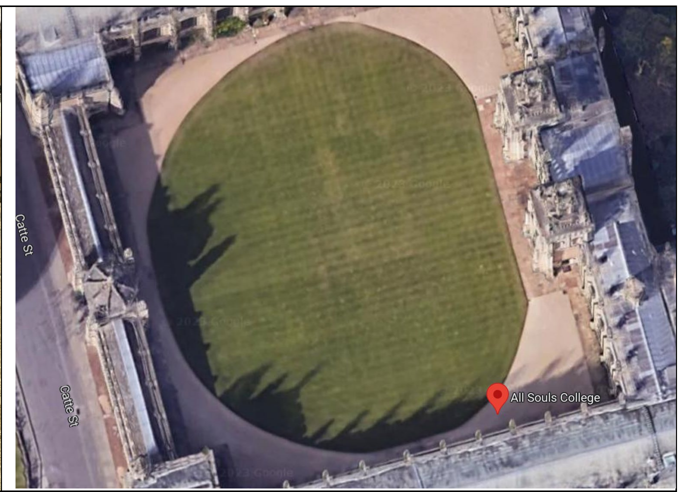

The Lord’s View

Informational naïveté is an empirically important strategic issue, but (Bayesian Nash) equilibrium rules it out by assuming people form rational expectations about others' decisions, including how they depend on their private information.

A generalized level-*k* analysis, as in Camerer et al. (2004), Crawford and Iriberri (2007a), and Brocas, Carrillo, Camerer, and Wang (2014), suggests an explanation (see also Brown et al. 2012):

- Take *L0*'s decisions to be uniform random over the other player's (or players') feasible decisions, and *independent of its own private value*

- Define higher levels by iterated best responses as before

In this "random level-*k*" model, *L0* may seem odd at first. But *L0* is not an actual player, it is a player's naïve model of others whose values he does not observe.

Reasoning contingent on others' possible values is logically possible, but behaviorally far-fetched, often requiring high-dimensional fixed-point reasoning.

This yields a model of informational naivete akin to Eyster and Rabin's (2005) notion of cursed equilibrium, while unbundling naïveté from equilibrium.

A random *L1* believes that others' decisions are independent of their own private information. It is therefore vulnerable to the winner's curse and zero-sum bets:

- Sky Masterson's father was worried that his son would be an *L1*: decision-theoretically rational but insufficiently skeptical of offers "too good to be true"

- Milgrom and Stokey (1982), speculating on why zero-sum trades occur despite their no-trade theorem, conjecture rules Naïve Behavior, which sticks with its prior but otherwise behaves rationally, like a random *L1*; and First-Order Sophistication, which best responds to Naïve Behavior like a random *L2*

- A random *L1* is "fully cursed" as in Eyster and Rabin's (2005) cursed-equilibrium analysis, but it need not be correct on average about others' decisions (as assumed in a cursed equilibrium)

- Unlike cursed equilibrium, which bundles equilibrium and informational naïveté, a random level-*k* model derives naïveté from aversion to contingent reasoning

But the model's usefulness is an empirical question, on which the jury is still out.

With regard to auctions, Crawford and Iriberri (2007a) re-did Milgrom and Weber's (1982) equilibrium analysis of first- and second-price, independent-private-value and common-value auctions with a level-*k* model and used the results to re-analyze data from the classic auction experiments, finding that a random level-*k* model *mostly* outperforms equilibrium or cursed equilibrium.

With regard to zero-sum betting, Brocas et al. (2014) (see also Camerer et al. 2004, Section VI.A) report experiments that have the power to distinguish between a random level-*k* model and explanations based on joy of gambling, etc.

Unlike Milgrom and Stokey (1982), whose theory implicitly assumes a magical Walrasian auctioneer, Brocas et al. study simple, completely-specified betting games—as needed for a laboratory implementation.

Brocas et al.'s experiments were on three-state betting games (close to zero-sum game-theoretic analogues of Milgrom and Stokey's market trading model).

There are three ex ante equally likely states, A, B, C.

Player 1 learns privately either that the state is {A or B} or that it is C.

Player 2 learns privately either that the state is A or that it is {B or C}.

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

Players then choose simultaneously whether to Bet or Pass.

A player who chooses Pass, or who chooses Bet while the other player chooses Pass, earns 10, whatever the state.

If both choose Bet, they get their payoffs in the table for whichever state occurs.

All this is publicly announced (to induce common knowledge).

This game has a unique sensible equilibrium, identifiable via 3 rounds of "iterated weak dominance" (there's a nonsensical equilibrium in which both always Pass).

Round 1 of iterated weak dominance (Bet, Pass)

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

Round 2

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

Round 3

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

In equilibrium, there is no betting in any state (player 1 is *willing* to bet in state C).

Despite this clear equilibrium prediction, in Brocas et al.'s experiments (and in several previous ones), half the subjects expressed a willingness to bet, in patterns that varied systematically with their player role and the state.

In a careful clustering analysis that let the data speak directly, Brocas et al. use a random level-*k* model, again taking *L0*'s decisions to be uniformly distributed and independent of its own value, to interpret their results.

L1 respects only simple dominance:

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

L2 respects two rounds of iterated weak dominance:

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

And *L3* respects three rounds (enough for equilibrium in this game):

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

If all subjects were *L1s*, 100% of player 1s and 67% of player 2s would be willing to bet, with 100% betting in states B and C: each many more than in the data.

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

But Brocas et al. find clusters of subjects whose behavior corresponds to each of *L1*, (arguably) *L2*, and *L3*; and also a cluster of “irrational” players.

*L2*s and *L3*s are less gullible than *L1*s, and the level-*k* model with estimated level frequencies fits better than equilibrium or any model of homogeneous thinking.

Level-*k* mechanism design

Crawford, Kugler, Neeman, and Pauzner (2009) discuss the design of revenue-maximizing independent-private-value auctions for level-*k* bidders, as in Myerson's (1981) equilibrium-based analysis.

Robust mechanism design is often thought of as getting the equilibrium-based optimum under weaker behavioral assumptions, but with a level-*k* or other model that predicts deviations it may be possible to do better.

For level-*k* bidders, revenue-equivalence fails.

All levels respect simple dominance, so with independent private values a second-price auction makes level-*k* bidders bid like equilibrium bidders.

In a first-price auction, *L1* bidders tend to overbid relative to equilibrium, *L2*s tend to underbid, etc. (Crawford and Iriberri 2007a).

Thus if *L1*s are known to predominate in the population, a first-price auction will yield more expected revenue than a second-price auction.

What famous model does this statue (in Singapore) represent?

Crawford (2021) revisits Myerson and Satterthwaite's (1983; "MS") classic equilibrium-based analysis of design of incentive-efficient mechanisms for bilateral trading with independent private values, replacing MS's equilibrium assumption with a level-*k* model.

MS's analysis seems at first to require the full power of their equilibrium assumption, which bundles four distinct behavioral assumptions: (i) decision-theoretic rationality, (ii) homogeneity of strategic thinking, and (iii) predictability and (iv) coordination/correctness of beliefs.

Crawford's level-*k* analysis unbundles those assumptions, retaining (i) decision-theoretic rationality and relaxing the behaviorally more questionable (ii) homogeneity of strategic thinking and, in a structured way, (iii)-(iv) predictability and coordination/correctness of beliefs.

(See also de Clippel, Saran, and Serrano's (2018) analysis of implementation with bounded depths of reasoning and Kneeland's (2022) analysis of level-*k* implementation, including bilateral trading.)

MS's analysis was inspired by Chatterjee and Samuelson's (1983; "CS") positive analysis of bilateral trading with independent private values via double auction.

CS and MS both study trading between a potential seller and buyer of an indivisible object, in exchange for money.

Seller's and buyer's von Neumann-Morgenstern utility functions are quasilinear in money: risk-neutral, with money values for the object.

Denote the buyer's value V and the seller's value C (for "cost").
V and C are independent, with positive densities f(V) and g(C) on their supports and distribution functions F(V) and G(C).

CS and MS allowed the densities to have any bounded overlapping supports, but without important loss of generality I take the supports to be identical and normalize them to [0, 1].

Equilibrium in the double auction

In the double auction:

- If the buyer's bid b ≥ the seller's ask a, the seller exchanges the object for a given weighted average of b and a

- CS allowed any weights between 0 and 1, but I take the weights to be equal, so the buyer acquires the object at price (a + b)/2, the seller's utility is (a + b)/2, and the buyer's is V - (a + b)/2

- If b < a, the seller retains the object, and no money changes hands

The double auction has many Bayesian equilibria. When f(V) and g(C) are uniformly distributed, CS identify a linear equilibrium, which also plays a central role in MS's analysis.

Denote the buyer's bidding strategy b(V) and the seller's asking strategy a(C), with * subscripts for the equilibrium strategies. In the linear equilibrium, with value densities supported on [0, 1],

$$b_*(V) = 2V/3 + 1/12$$

unless V < ¼, when $b_*(V)$ can be anything that precludes trade; and

$$a_*(C) = 2C/3 + ¼$$

unless C > ¾, when $a_*(C)$ can be anything that precludes trade.

Trade occurs if and only if 2V/3 + 1/12 ≥ 2C/3 + ¼, or V ≥ C + ¼; thus with positive probability the outcome is ex post inefficient (true in all equilibria).

The ex ante probability of trade is 9/32 ≈ 28% and the expected total surplus is 9/64 ≈ 0.14, less than the maximum individually rational probability of trade 50% and expected surplus 1/6 ≈ 0.17.

Ex ante-incentive-efficient mechanisms, assuming equilibrium

MS noted the inefficiency of double-auction equilibria and asked whether another choice of trading rules ("mechanism") could assure efficiency.

MS characterized ex-ante-*incentive*-efficient mechanisms in CS's trading environment, also requiring interim individual rationality.

Like CS, MS allowed general, independent private value distributions with strictly positive densities on ranges that overlap for the buyer and seller; but I will continue to take both value supports to be [0, 1].

MS assumed the designer can commit to any trading rules and traders will play any desired equilibrium in the game created by those rules.

The revelation principle shows that there is then no loss of generality in restricting attention to incentive-compatible direct mechanisms.

(A *direct* mechanism is one in which a player's decisions are conformable to value reports, with the outcome a function of the reports.

An *incentive-compatible* mechanism is one in which truthful reporting is an equilibrium.

An *(interim) individually rational* mechanism is one that yields buyer and seller expected utility ≥ 0 for every possible realization of their values.)

When traders have quasilinear utility functions, denoting their value reports v and c (distinct from their true values V and C), the payoff-relevant aspects of an outcome are determined by two functions:

- p(v, c), the probability that the object is transferred, and

- x(v, c), the expected monetary payment from buyer to seller

Although these outcome functions depend only on reported values, traders' utilities are determined by their true values.

MS's Theorem 1 uses the conditions for incentive-compatibility and individual rationality to derive an "incentive budget constraint" (my term, not theirs), which summarizes a given mechanism's expected surplus cost of creating incentives for traders to reveal their true values.

Corollary 1 uses that constraint to show that no incentive-compatible, individually rational mechanism can assure ex post Pareto-efficiency.

MS's Theorem 2 studies the problem of choosing the outcome functions $p(\cdot, \cdot)$ and $x(\cdot, \cdot)$ to maximize the sum of traders' ex ante expected utilities subject to the incentive budget constraint, whose conditions characterize the outcome functions of ex-ante-incentive-efficient mechanisms.

In CS's example with uniform value densities, MS's Theorem 2 yields a closed-form solution for an incentive-compatible, ex-ante-incentive-efficient mechanism, which transfers the object if and only if the reported values satisfy $v \geq c + 1/4$, at price $(v + c + 1/2)/3$—just as happens in the linear equilibrium of the double auction with uniform densities.

Thus, even though the double auction is not incentive-compatible, traders' linear equilibrium bidding strategies in it shade to mimic the outcome of truthful reporting in MS's ex-ante-incentive-efficient mechanism.

Put another way, with uniform value densities the double auction is outcome-equivalent to an ex-ante-incentive-efficient mechanism. (This does not generalize beyond uniform value densities.)

Level-*k* analysis of the double auction

Following Crawford and Iriberri (2007a), Crawford (2021) applied the level-*k* model to CS's double auctions with uniform value densities, denoting the buyer's bidding strategy $b_i(V)$ and the seller's asking strategy $a_i(C)$, where subscripts denote levels i = 1,2.

L1 and *L2* buyers and sellers have the same value-shading slope, 2/3, as equilibrium bidders, but *L1*s (*L2*s) are more (less) aggressive.

For *L1*s double auction outcomes are less efficient than for equilibrium bidders, while for *L2*s double auction outcomes are more efficient.

- Is there a mechanism that could counteract *L1*s' aggression to reduce the inefficiency of the double auction for *L1*s?

- And is there a mechanism that could take advantage of *L2*s' lack of aggression to preserve or improve upon the efficiency of the double auction for *L2*s?

Mechanism design for level-*k* traders

Throughout the analysis of level-*k* design, I restrict attention to direct mechanisms, I define incentive-efficiency notions for a designer's correct beliefs, but I derive incentive constraints from traders' level-*k* beliefs.

To begin to answer the above questions, consider MS's ex-ante-incentive-efficient mechanism with uniform value densities.

For it, *L1* buyers' and sellers' incentive constraints happen to coincide with equilibrium buyers' and sellers' incentive constraints, because *L1*s anticipate uniform random responses and the value densities are uniform. By induction, this coincidence also holds for *Lk* incentive constraints.

Thus MS's closed-form solution for an incentive-compatible, ex-ante-incentive-efficient mechanism with uniform value densities remains valid for any populations of level-*k* buyers and sellers, whose levels are known or not: MS's result for this case is completely robust to level-*k* thinking.

But for level-*k* traders the mechanism must be implemented in incentive-compatible form: The “raw” double auction, though outcome-equivalent to MS’s mechanism for equilibrium traders, is not equivalent for *L1* traders.

The revelation principle fails for level-*k* traders because the choice of mechanism influences the correctness of level-*k* beliefs.

This influence allows MS’s mechanism to improve outcomes for *L1*s via what I call tacit exploitation of predictably incorrect beliefs:

“Tacit” because the designer need not lie, only to present the mechanism.

“Exploitation” in a benign sense of extracting more surplus and using it to benefit traders by increasing the efficiency of trade (unlike in auctions).

“Predictably incorrect” to the extent that the level-*k* model is correct and the designer knows traders’ levels.

More generally, the failure of the revelation principle implies a substantive choice regarding whether or not to require level-*k*-incentive compatibility.

Some analysts have argued that incentive compatibility is essential for applications, e.g. for school choice or combinatorial auctions.

Others are willing in some contexts to consider non-incentive-compatible mechanisms like the Boston Mechanism or first-price auctions.

I consider both cases, but I focus on when incentive compatibility is required, distinguishing whether traders' levels are observable or not.

(The paper also contains some discussion of design relaxing level-*k*-incentive-compatibility, still assuming that traders best respond.)

Returning to the second question about auctions above, if incentive compatibility is required it's not possible to exploit *L2*s' lack of aggression in the double auction: For, the double auction is not incentive compatible, and we have seen that with uniform value densities the best level-*k*-incentive-compatible mechanism is MS's equilibrium-incentive-compatible, ex-ante-incentive-efficient mechanism.

Design requiring level-*k*-incentive-compatibility when traders' levels are observable

With non-uniform value densities, when level-*k*-incentive-compatibility is required and traders' levels are observable, I assume that the designer can custom-tailor the mechanism to traders' levels.

An analysis like MS's then goes through, with qualitatively similar results but differences in detail (except for uniform value densities).

But MS's result that no incentive-compatible mechanism can ensure ex post efficient trade with probability one does not quite go through: For some value densities it may be possible to ensure ex post efficient trade.

Design requiring level-*k*-incentive-compatibility when traders' levels are unobservable

When level-*k*-incentive-compatibility is required and traders' levels are unobservable, a direct mechanism is not rich enough to screen both traders' levels and their private values.

Then, posted-price mechanisms eliminate conflict between different levels' incentive constraints, and only they are incentive-compatible.

It is easy to characterize the optimal posted price, which can be implemented is a distribution-free way, satisfying Wilson's (1987) critique.

Posted-price mechanisms rule out the sensitivity to reported values of mechanisms that are efficient in the set of level-*k*-incentive-compatible mechanisms when traders' levels are observable.

But at least with near-uniform value densities the cost in expected total surplus cost is modest, a fall from 0.14 to 0.125.

The analyses with and without observable levels requiring level-*k*-incentive-compatibility trace the need for robust implementation to the unpredictability of strategic thinking, a plausible rationale for robustness.

The analysis also brings expected surplus-maximizing mechanisms closer to mechanisms used in applications, where the unpredictability of people's thinking often seems to exert a major influence on design.

Some references

Adam Brandenburger, "Knowledge and Equilibrium in Games," *Journal of Economic Perspectives* 6 (1992), 83-101

Isabelle Brocas, Juan Carrillo, Stephanie Wang, and Colin Camerer, "Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games," *Review of Economic Studies* 81 (2014), 944–970

Alexander Brown, Colin Camerer, and Dan Lovallo, "To Review or Not To Review? Limited Strategic Thinking at the Movie Box Office," American Economic Journal: Microeconomics 4 (2012), 1-28

Colin Camerer, Teck-Hua Ho, and Juin-Kuan Chong, "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics* 119 (2004), 861-898

Miguel Costa-Gomes and Vincent Crawford "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study", *American Economic Review* 96 (2006), 1737-1768

Miguel Costa-Gomes, Vincent Crawford, and Bruno Broseta, "Cognition and Behavior in Normal-Form Games: An Experimental Study", *Econometrica* 69 (2001), 1193-1235

Vincent Crawford, "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions," *American Economic Review* 93 (2003), 133-149

Vincent Crawford, "Look-ups as the Windows of the Strategic Soul: Studying Cognition via Information Search in Game Experiments," Chapter 10 in Andrew Caplin and Andrew Schotter, editors, *Perspectives on the Future of Economics: Positive and Normative Foundations*, Oxford University Press, 2008; http://econweb.ucsd.edu/%7Evcrawfor/5Oct06NYUCognitionSearchMain.pdf

Vincent Crawford, "New Directions for Modelling Strategic Behavior: Game-theoretic Models of Communication, Coordination, and Cooperation in Economic Relationships," *Journal of Economic Perspectives* 30 (2016), 131-150

Vincent Crawford, "Let's Talk It Over: Coordination via Preplay Communication with Level-*k* Thinking," *Research in Economics* 70 (2017), 20-31

Vincent Crawford, "Efficient Mechanisms for Level-*k* Bilateral Trading," *Games and Economic Behavior* 127 (2021), 80-101

Vincent Crawford, Miguel Costa-Gomes, and Nagore Iriberri, "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature* 51 (2013), 5-62

Vincent Crawford and Nagore Iriberri, "Level-*k* Auctions: Can Boundedly Rational Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica* 75 (2007), 1721–1770

Vincent Crawford and Joel Sobel, "Strategic Information Transmission", *Econometrica* 50 (1982), 1431-1451

Tore Ellingsen and Robert Östling, "When Does Communication Improve Coordination?", *American Economic Review* 100 (2010), 1695-1724

Erik Eyster and Matthew Rabin, "Cursed Equilibrium," *Econometrica* 73 (2005), 1623-1672

Paul Milgrom and Nancy Stokey, "Information, Trade and Common Knowledge," *Journal of Economic Theory* 26 (1982), 17-27

Paul Milgrom and Robert Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica* 50 (1982), 1089-1122

Rosemarie Nagel, "Unraveling in Guessing Games: An Experimental Study," *American Economic Review* 85 (1995), 1313-1326

Reinhard Selten, "Features of Experimentally Observed Bounded Rationality," *European Economic Review* 42 (1998), 413-436

Dale Stahl and Paul Wilson, "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior and Organization* 25 (1994), 309-327

Joseph Tao-yi Wang, Michael Spezio, and Colin Camerer, "Pinocchio's Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games," *American Economic Review* 100 (2010), 984-1007