

**Workshop on Bounded Rationality in Honor of Reinhard Selten,
Poznań University of Economics and Business, 5 October 2026**

and

Royal Holloway, University of London, 10 November 2026

Strategic Thinking: Theory and Evidence

**Vincent P. Crawford
University of Oxford, All Souls College,
and University of California San Diego**



UC San Diego



European Research Council
Established by the European Commission

10 September 2026

Prologue: Young Reinhard

1998 *EER*: “Basic concepts in game theory are often circular in the sense that they are based on definitions by implicit properties.... Boundedly rational strategic reasoning seems to avoid circular concepts. It directly results in a procedure by which a problem solution is found.”

1998 Conference on Mental Models in Games, Venice International University, Isola San Servolo, Venice (photos by Alessandro Rossi):



The talk is about modeling strategic thinking in games where equilibrium lacks a credible justification. It is based on the first part of my forthcoming Econometric Society monograph, *Strategic Thinking*, which builds on work with Bruno Broseta, Miguel Costa-Gomes, and Nagore Iriberri.

Bruno Broseta Dupre, now
Professor of Business, European
University of Valencia



Miguel Costa-Gomes, now Head
of the School of Economics and
Professor of Economics,
University of Edinburgh



Nagore Iriberri, now Ikerbasque
Research Professor, University of
the Basque Country



Introduction

Game theory already has a model of strategic thinking in Nash equilibrium.

Equilibrium builds in the assumption that each player's decisions are rational, which has considerable empirical support in nonstrategic applications.

But in most games, rationality yields weak restrictions on players' decisions.

This creates a demand for stronger behavioral assumptions, which is usually met by adding equilibrium's other assumption, that each player's decision is a best response to other players' decisions.

Defining equilibrium this way makes it possible to identify equilibrium decisions without regard to their rationales.

But only individuals think, so it is inappropriate to base an analysis of thinking on a basic assumption about the *relationship* among players' decisions.

For my purposes it is better to view equilibrium as “equilibrium in beliefs,” in which players independently form beliefs about each other's decisions and their best responses to their beliefs are self-confirming: rational expectations.

Together, equilibrium's rationality and rational-expectations assumptions yield far more precise predictions than common knowledge of rationality.

This precision makes it possible to address strategic questions with a theoretical clarity once attainable only for nonstrategic questions.

Equilibrium's clarity, generality, and tractability have made it the method of choice for theoretical strategic analyses in economics and other disciplines.

Empirically, however, equilibrium is only the start of the story.

Assuming that players are rational, equilibrium's rational-expectations assumption has two possible justifications: learning and thinking.

A learning justification requires that people have ample experience with closely analogous but strategically independent games.

The movie *Groundhog Day* provides an extreme but informative folk example:

- Its protagonist Phil Connors (played by Bill Murray) is stuck in a time loop, endlessly repeating the same day.
- Connors is free to make different choices in each new loop and can perfectly recall what happened in past loops.
- The other characters are also free to make different choices in each new loop; but unlike Connors, they *cannot* recall what happened in past loops.
- The endless loops give Connors unlimited opportunities to learn from experience with *closely analogous* but *strategically independent* games.
- He uses them to learn the best decisions in many different settings, even overcoming a very unpromising start to establish a romantic relationship.

Contrast *Groundhog Day* with a real example, the first FCC spectrum auctions:

- The auctions lacked close precedents, so bidders could not have learned their equilibrium bidding strategies from experience.
- The auctions had high potential revenues, which would be lost unless the design worked well the first time (and were lost with some later designs).
- The designers were well aware of auction theory's equilibrium-based Implications for revenue-maximizing sealed-bid or progressive designs (but without the spectrum's heterogeneity of goods and multi-good complementarities).
- Even so, instead of taking a design off the equilibrium-optimal shelf, the designers—informed by experiments—chose a novel progressive design, whose performance did not depend on bidders immediately coming close to equilibrium and therefore made far milder demands on their cognition.

In my view, many applications in which we reflexively assume equilibrium are more like the FCC spectrum auctions than *Groundhog Day*.

In applications where equilibrium lacks a credible learning justification, people can, in theory, still think their way to an equilibrium.

But a thinking justification for equilibrium requires far more than rationality:

The weakest sufficient conditions, for games in general, to ensure that players will focus their beliefs on an equilibrium are that their rationality and beliefs are common knowledge (Brandenburger 1992 *JEP*, Proposition 4).

Common knowledge of players' beliefs implies that they all have the same, “rational” expectations regarding each other's decisions.

Without learning, such rational expectations must come from common knowledge of the logic of some equilibrium.

But in all but the simplest games, equilibrium logic requires fixed-point or indefinitely iterated-dominance reasoning, which are complex enough that even people who can follow them are unlikely to believe that others will follow them.

Empirically, in experiments that elicit subjects' initial responses to games, subjects often deviate systematically from equilibrium.

As Selten (1998) said, regarding such experiments:

“Boundedly rational strategic reasoning seems to avoid circular concepts. It directly results in a procedure by which a problem solution is found.”

A fortiori, the signs on some trucks that say, “If you can't see my mirrors, I can't see you,” suggest that some people must be reminded to consider others' cognition at all.

What should we do in applications where equilibrium's justification is weak?

Should we assume equilibrium because we hope its predictions will still be correct on average, or fear that without it there is no basis for analysis?

I will argue that when equilibrium lacks a credible learning or thinking justification, there is no reason to believe its predictions will be correct, even on average.

I will further argue that we can then use theory and experimental evidence to find *nonequilibrium* models that predict and correct equilibrium's biases, while preserving much of equilibrium's generality, precision, and tractability.

The goal is not to supplant equilibrium but to amend it in a useful way.

Better nonequilibrium models must answer (at least) two questions:

- How could *any* model out-predict a rational-expectations notion such as equilibrium?
- And how could we identify such a model within the enormous set of possible nonequilibrium models?

(By analogy with the first line of Tolstoi's *Anna Karenina*, "Happy families are all alike; every unhappy family is unhappy in its own way":

How could there be a *general* theory of unhappy families—or of nonequilibrium decisions?)

The rest of the talk proposes answers as follows:

- When a person plays a game in which s/he finds the logic of equilibrium inaccessible or unconvincing, s/he must still somehow decide what to do.
- If how s/he decides has empirical regularities in across games, they could identify a decision rule—a function or correspondence of the game that specifies a decision in each game—that explains her/his decisions.
- Although there are no guarantees in empirical work, experiments that study strategic thinking by eliciting subjects' initial responses to series of different games sometimes yield reliable estimates of individual subjects' decision rules.
- Subjects' estimated decision rules are heterogeneous, and most of them (noisily) follow one of a small number of discrete rules.

This suggests that the structure of strategic thinking is completely different than in some leading models, such as equilibrium or quantal response equilibrium with continuous errors.

- Importantly, the decision rules most people follow make their decisions conditionally predictable (with noise), game by game.
- Thus, estimating or calibrating a population's rule distribution yields a model that can predict the distribution of players' decisions, *effectively in advance*.

Such a model amends equilibrium analysis rather than supplanting it:

- In some games the rules most people follow mimic equilibrium decisions and the model establishes robustness of equilibrium-based predictions.
- In others people's rules deviate predictably from equilibrium and the model corrects equilibrium's bias, out-predicting its rational expectations.

Remarks

- My view that such models are worth further study is *not* based on a belief that they explain most of people's decisions; they stop well short of that.
- I do believe that we need theoretically and empirically grounded ways to think about applications in which equilibrium's justifications are weak.
- And it is surprising that *any* model can explain even a fraction of people's deviations from equilibrium: Such models are a glass half full.

- Although experiments play the leading empirical role in identifying nonequilibrium models, the models have implications for field data.
- The book discusses empirical analyses of entry in industrial organization, where firms are assumed to make equilibrium decisions with noise (e.g. Bresnahan and Reiss 1990 *REStud*, 1991 *JPE*).
- But firms don't get much practice entering markets, and if they follow nonequilibrium decision rules, equilibrium predictions may be biased.
- A nonequilibrium model can correct such bias (e.g. Goldfarb and Yang 2009 *J. Marketing Research*, Goldfarb and Xiao 2011 *AER*).

Outline

In the rest of the talk I develop these arguments for simultaneous-move, symmetric-information games.

(The book also considers some games with asymmetric information and others that are sequential in that they allow preplay communication.)

I analyze decisions at the individual level and assume people's decisions are rational (but possibly noisy) best responses to beliefs about others' decisions.

A model represents a player population as a frequency distribution over a set of decision rules. A player role is filled by an i.i.d. draw from the distribution.

The rest of the talk is organized as follows:

- Essential and negotiable desiderata for models of strategic thinking.
- Overview of leading models of strategic thinking.
- Assessing models in the light of desiderata and experimental evidence.
- Further remarks on QRE, LQRE, level- k , and CH models.

(The book goes on to discuss theoretical applications of level- k or CH models to a range of questions equilibrium analysis does not resolve, from outguessing to symmetry-breaking and assurance in coordination, salience, communication of intentions or private information, and mechanism design.)

Essential model desiderata

- A model should explain each player's initial responses across games via a single decision rule—a function (or correspondence) of the game that specifies a decision (or set of decisions) in each game.
- A player's decision rule should give a comprehensible, direct account of her/his decisions. (That is what makes the first desideratum restrictive.)

Selten: “Boundedly rational strategic reasoning...directly results in a procedure by which a problem solution is found.”

In modelling strategic thinking, there is no “as if” safety net:

The book argues that a decision rule that mimics a person's decisions in a set of games without explanation, or a “black-box” model of those decisions, even with cross-validation, is not reliable beyond them.

Negotiable desiderata

- A model's decision rules should avoid fixed-point reasoning.
- A model's decision rules should avoid contingent reasoning.
- A model's decision rules should make precise predictions, at least generically determining a unique distribution of decisions.
- (A new notion) A model's decision rules should be *self-informed*.
I call a rule self-informed if players can identify their rules' mean decisions using only information directly available to themselves.
(Requiring this only for *mean* decisions avoids proscribing rules that determine only decision distributions, like quantal response equilibrium.)
Directly available information includes theoretical predictions of others' decisions; it excludes empirical parameters pertaining to others.
Non-self-informed rules can be *logically* consistent, but they assume unlimited sophistication regarding others' decisions, while denying sophistication with regard to equilibrium logic: *behaviorally* inconsistent.

Leading models of strategic thinking

The leading models of strategic thinking all satisfy my essential desiderata but vary in compliance with my negotiable desiderata.

- Equilibrium plus noise, which adds possibly payoff-sensitive noise to equilibrium decisions, evaluating payoffs for exact equilibrium beliefs.

Experimental subjects' and other people's decisions are often highly sensitive to out-of-equilibrium payoffs.

Equilibrium plus noise assumes that players evaluate their decisions' payoffs assuming that others make noiseless equilibrium decisions.

This precludes sensitivity to out-of-equilibrium payoffs, even with payoff-sensitive noise.

The next models all seek to explain such sensitivity, in different ways.

- Rationalizability and k -rationalizability (Bernheim 1984 *Ecma* and Pearce 1984 *Ecma*), which relaxes equilibrium's rational expectations but respects common (or k levels of) knowledge of rationality.

In two-person games the rationalizable (or k -rationalizable) decisions are those that survive an indefinite number (or k) rounds of eliminating strictly dominated decisions. They are rarely unique. k can differ across players.

Rationalizability and k -rationalizability are *consistent with* sensitivity to out-of-equilibrium payoffs but do not truly seek to explain it.

- In a dominance- k model (Costa-Gomes et al. 2001 *Ecma*; “CGCB”; Costa-Gomes and Crawford 2006 *AER*; “CGC”), a D_k player does k rounds of eliminating dominated decisions (simultaneously for all players) and best responds to a uniform prior over others’ remaining decisions (possibly with noise). A player’s k remains the same across games of comparable complexity. k can differ across players.

In games with no dominated decisions D_k , $k = 1, 2, \dots$ make the same decisions as a level- k or CH $L1$, defined below.

D_{k-1} respects k -rationalizability and refines it to yield precise decisions.

- In a quantal response equilibrium (“QRE”; Rosenthal 1993 *JEBO*, McKelvey and Palfrey 1995 *GEB*), players’ decisions are noisy, with a decision’s density increasing in its expected payoff, evaluating payoffs taking other players’ noise into account (unlike in equilibrium plus noise). QRE is completed by specifying a response distribution with estimated precision parameter. (A logit distribution yields logit QRE (“LQRE”).)

QRE and LQRE generalize equilibrium’s rational expectations to allow players to noisily best respond to the noise in other players’ decisions.

- In a level- k model (Keynes 1936, Nagel 1995 *AER*; Stahl and Wilson 1994 *JEBO*, 1995 *GEB*; Ho et al. 1998 *AER*; “HCW”; CGCB; CGC) or a cognitive hierarchy model (“CH”; Camerer et al. 2004 *QJE*), an L_k player anchors her/his beliefs on a naïve model of others’ decisions called L_0 and iterates best responses k times. A player’s k remains the same across games of comparable complexity. k can differ across players.

A level- k L_1 best responds to L_0 , L_2 to L_1 , L_3 to L_2 , and so on.

A CH L_1 best responds to L_0 but higher CH levels respond to estimated mixtures (usually Poisson) of lower levels, possibly including L_0 . L_k Bayesian updates beliefs about others’ levels, assuming they are lower.

L_0 is usually taken to be uniform random over the player’s decisions.

(The book discusses independent-random or truthful L_0 specifications for games with asymmetric information or preplay communication.)

Higher level- k and CH levels’ decisions are usually allowed to be noisy.

Level- k and CH L_1 s make the same decisions; higher levels may differ.

The book argues (on the grounds that $L1$ is also naïve and errors should be modeled flexibly) that $L0$ should be viewed only as the starting point of higher levels' models of others, not as a rule any actual player uses.

Unconstrained estimates of $L0$'s frequency are usually 0 or close to 0.

Level- k and CH rules above $L0$ are rational. They differ from equilibrium only in best responding to oversimplified models of others.

Level- k (but not CH) L_{k+1} respects k -rationalizability and refines it in a natural way (often different from Dk 's) that also yields precise decisions.

Level- k and CH rules both mimic equilibrium decisions in sufficiently simple games, while deviating systematically in more complex games.

- In a “representative-agent” variant of level- k models, which Nagel and HCW propose for their symmetric n -person guessing games, a player anchors her/his beliefs on a uniform random $L0$ that pertains to others' average guess (the payoff-relevant summary statistic in their games). One could also define a “representative-agent” variant of CH models.

Level- k models appear to have first been suggested by Keynes (1936, Chapter 12), in his famous analogy comparing speculative markets to a beauty contest:

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

The rules Keynes mentions are heterogeneous, discrete, and based on finitely iterated best responses to lower levels—just like level- k (but not quite CH) rules.

Keynes might have read von Neumann (1928, in German), but he was surely less tempted by the fixed-point reasoning of rational expectations than we are.

My negotiable desiderata discriminate sharply among models. Simplifying from the book, only level- k is self-informed, precise, and avoids fixed-point and contingent reasoning. CH is all of these but self-informed.

	Precise Predictions	Set-Valued Predictions
Self-Informed	Equilibrium Level-k <i>Dk</i>	<i>Rationalizability</i> <i>k-Rationalizability</i>
Not Self-Informed	QRE LQRE CH	

Leading decision rules' desiderata
(Rules that don't require contingent reasoning are in **bold**; rules that don't require fixed-point reasoning are in *italic*.)

Incidentally, adaptive learning models such as reinforcement, beliefs-based, and EWA learning can all be plausibly implemented in self-informed ways. But sophisticated learning models such as Selten's (1991) anticipatory learning can be implemented only in non-self-informed ways.

Further remarks on QRE, LQRE, level- k , and CH models

QRE and LQRE replace equilibrium's fixed-point reasoning in the space of beliefs about decisions (possibly mixed but with limited support) with fixed-point reasoning in the far larger space of full-support decision *distributions*. If equilibrium logic is cognitively taxing, QRE's logic is doubly taxing.

QRE and LQRE are not self-informed because their mean decisions are sensitive to the precision of the error distribution (even if its form is known)—far more sensitive in some games than for noisy individual decisions:

QRE and LQRE's potential to explain sensitivity to out-of-equilibrium payoffs unavoidably relies on empirical parameters of others' decision distributions.

QRE and LQRE implicitly assume that a player who, by hypothesis, cannot precisely identify her/his own equilibrium decisions, is sophisticated enough to identify the *dispersion* of other players' deviations from equilibrium.

QRE's and LQRE's lack of self-informedness turns what is usually thought of as their main strength into a weakness:

Haile et al. 2008 *AER* show that if there is only one observation per game-player pair and error distributions are unrestricted, by varying the distributions QRE can be made to fit any dataset—however irregular—perfectly.

With more observations, varying the error distributions stops short of allowing a perfect fit, but QRE is still unusually sensitive to assumptions on the form of the error distribution. We have no theory of its form, and it has been specified entirely by convention (logit or probit), without testing.

(The book gives a more extensive critique, including of recent attempts to remedy QRE's and LQRE's lack of self-informedness.)

Level- k and CH decision rules avoid fixed-point reasoning by allowing players' models of others to be simpler than their models of themselves.

Level- k and CH rules' iterated best responses also avoid contingent reasoning (including iterated elimination of dominated decisions).

Level- k models achieve self-informedness by modeling people's mean responses to out-of-equilibrium payoffs via the deterministic structures of their decision rules, with theoretically specified models of others' decisions.

CH models do not quite achieve self-informedness because the estimated parameter of the level distribution unavoidably depends on others' decisions. The violation may be mild because people may often have more accurate beliefs about others' levels than about the noisiness of others' decisions.

Level- k models are fully consistent with both of Selten's (1998) dicta:

“Boundedly rational strategic reasoning seems to avoid circular concepts” in that they avoid fixed-point reasoning; and

“It directly results in a procedure by which a problem solution is found” in that each Lk rule is an algorithm for identifying its decision in any game.

CH models are just as consistent, except for their lack of self-informedness.

(The book also discusses level- k and CH models that distinguish people's abilities from their optimal levels given beliefs about others' abilities.)

Folk game theory: level- k versus CH and equilibrium in outguessing

The book links its applications to common intuitions about strategic questions by illustrating each with one or more quotations from “folk game theory”.

My term is meant to suggest an analogy with folk physics—untrained people’s intuitive beliefs about the laws of physics.

Folk physics imperfectly reflects real physics but yields insights into cognition.

Folk game theory imperfectly reflects traditional game theory but yields insights into the cognition behind its empirical counterpart, behavioral game theory.

If a model of strategic thinking should give a comprehensible, direct account of a player’s decisions, then folk game theory is data!

I now give a folk illustration of level- k thinking in an outguessing game with payoff asymmetry, contrasting it with CH and then equilibrium thinking.

Early in M.M. Kaye's novel *The Far Pavilions*, the main male character, Ash/Ashok, escapes and tries to evade pursuers along a north-south road.

Both choose, strategically simultaneously, between north and south.

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

South is warm, while north are the Himalayas with winter coming; so both Ash and the pursuers gain one extra unit of payoff for choosing south.

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

Ash's mentor tells him "...ride hard for the north, since they will be sure you will go southward where the climate is kinder...."

Ash follows the advice, the pursuers go south, and Ash escapes.

The mentor's advice is to choose a level- k (not CH) $L3$ response to a uniform random $L0$. His wording supports all three of this statement's claims:

- 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 between north and south is decisive only if they do not differ in the—more important to both—probability that Ash is caught.

- Given that pursuers model Ash as an $L1$, only level- k (not CH) $L2$ pursuers are "sure [Ash] will go southward".

CH $L2$ pursuers are not sure because they believe Ash might be an $L0$.

- Thus, Ash's mentor is modeling pursuers as level- k (not CH) $L2$, advising Ash to choose his level- k (not CH) $L3$ response to pursuers' $L2$ decision.

A CH $L3$'s decision differs from a level- k $L3$'s for some level frequencies.

Real applications seldom have a mentor telling us how players think.

But we can still specify a level- k model and use data to estimate its level frequencies econometrically or experience to calibrate their likely values.

Estimates of the population decision-rule distribution reflect strategic thinking's heterogeneity and are usually concentrated on $L1$, $L2$, and $L3$.

Suppose, for instance, that each player role is filled from a 50-30-20 mixture of $L1$ s, $L2$ s, and $L3$ s; and there are no decision errors.

Ash then goes north with probability $\frac{1}{2}$ and pursuers go south with probability $\frac{4}{5}$.

With independence, the likelihood of the novel's "observed" outcome {Ash north, pursuers south} is $\frac{2}{5}$, better than the $\frac{1}{4}$ that random decisions would imply.

By contrast, the game 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 $\Pr\{\text{south}\}$, $p^* = 1/4$, and the pursuers' $\Pr\{\text{south}\}$, $q^* = 3/4$. The observed outcome has likelihood $9/16 >$ the level- k likelihood of $2/5$.

Although the mixed-strategy equilibrium assigns a higher likelihood to the observed outcome than a level- k model, most people find equilibrium's logic counterintuitive:

- A player's equilibrium mixing probability is determined not by her/his own payoffs, but by equilibrium's need to make the other player's mixing consistent with rationality.
- Equilibrium mixing probabilities respond to the game's payoff asymmetry in an intuitive way for pursuers ($q^* = 3/4 > 1/2$, the equilibrium probability with no asymmetry); but in a counterintuitive way for Ash, who goes south with lower probability when its potential payoffs are higher ($p^* = 1/4 < 1/2$). The implication for Ash's mixing follows from equilibrium's fixed-point logic, but people's responses to such games tend to follow intuition in both roles.

The level- k model's heterogeneous strategic thinking roughly “quasi-purifies” the mixed-strategy equilibrium, with uncertainty about players' thinking rather than randomization driving strategic uncertainty; and it thereby avoids equilibrium's counterintuitive comparative statics.

Far Pavilions parallels the famous outguessing example in Conan Doyle's (1894) story "The Final Problem", which von Neumann and Morgenstern (1953, pp. 176-178) chose to illustrate their "minimax" mixed-strategy equilibrium for zero-sum two-person games: "Sherlock Holmes is as good as 48% dead when his train pulls out from Victoria Station."

However, although Holmes notes that Moriarty is his intellectual equal and the question is difficult, Holmes does not connect the difficulty to the lack of a pure-strategy equilibrium or the game's payoff asymmetry.

Is it unfair to ask for that as early as 1894? Search for Waldegrave (in 1713!) at https://en.wikipedia.org/wiki/Game_theory.

And Conan Doyle does not try to identify Holmes's or Moriarty's thinking. (In games, "intellectual equal" requires fixed-point reasoning; and the first existence theorems came only ten or fifteen years later.)

Experiments that elicit subjects' initial responses to games

Nagel, HCW, and Bosch-Domènech et al. 2002 *AER*

In Nagel's, HCW's, and Bosch-Domènech et al.'s n -person guessing games, subjects made guesses between lower and upper limits. The subject who guessed closest to a target p times the group average guess won a prize.

The games' structures were publicly announced and the games are dominance-solvable (except with target $4/3$), with a unique equilibrium "all-0".

The targets and limits were always the same for both players; they varied at most between subjects, across different treatments.

Nagel's and HCW's subjects played the same game repeatedly in the same group. Their *first-round* guesses can be viewed as initial responses if they treated their own influences on the future as negligible, which is plausible.

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




FIGURE 1. CHOICES IN THE FIRST PERIOD: A) SESSIONS 1–3 ($p = 1/2$); B) SESSIONS 4–7 ($p = 2/3$); C) SESSIONS 8–10 ($p = 4/3$)

Nagel's Figure 1. Distributions of subjects' first-round guesses (top $p = 1/2$, middle $p = 2/3$, bottom $p = 4/3$)




FIGURE 1. RELATIVE FREQUENCIES OF CHOICES
IN THREE NEWSPAPER EXPERIMENTS

**Bosch-Domènech et al.'s Figure 1. Distributions of subjects' guesses
($p = 2/3$)**

The leading games, with lower and upper limits 0 and 100 and targets < 1 , are dominance-solvable in an infinite number of rounds, with 0 the unique equilibrium or rationalizable decision. Yet initially, in each case:

- Subjects rarely made their equilibrium or rationalizable guesses.
- Few subjects respected k -rationalizability for $k > 3$.
- Subjects' guesses were highly heterogeneous.

Importantly:

- Subjects' aggregate guess distributions had spikes—visible for the first time in the design's large decision spaces—that track $50p^k$ for $k = 1, 2, 3$, even across treatments with different subjects and different target p s.

The spikes were between subjects: not subject idiosyncrasies but a feature of thinking that generalizes across subjects and games.

This discrete, partly deterministic structure is very different than in equilibrium or quantal response equilibrium for any plausible error distribution, and any continuous error distribution.

It is also inconsistent with subjects following heterogeneous rules along a continuum, such as equilibrium with a continuous payoff parameter.

Nagel, HCW, and Bosch-Domènech et al. attributed the spikes to level- k decision rules based on a representative-agent, uniform random $L0$.

Many theorists, however, view them as evidence that subjects consciously performed finitely iterated elimination of dominated decisions.

Making finitely iterated elimination of dominated decisions precise, as in CGCB's and CGC's dominance- k decision rules, Nagel's, HCW's, and Bosch-Domènech et al.'s main treatments perfectly confound L_{k+1} and D_k (the difference in indices is an unimportant quirk of my notation):

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

Because L_{k+1} and D_k make different decisions in other games, it's important to know which of them, or which *other* decision rule, subjects were following.

Stahl and Wilson 1994 *JEBO*, 1995 *GEB* and Costa-Gomes et al. 2001 *Ecma* (“CGCB”)

Despite Nagel's, HCW's, and Bosch-Domènech et al.'s large decision spaces, their one initial response per subject yields only weak separation of decision rules (not only in confounding $Lk+1$ and Dk).

Their dominance-solvable guessing games also confound rationalizability and equilibrium's fixed-point reasoning.

Stahl and Wilson and CGCB took first steps toward identifying subjects' decision rules more precisely, boosting power by having subjects play series of different two-person matrix games.

Subjects were randomly and anonymously paired, with no feedback until the end, to suppress repeated-game effects and rote learning.

(“Eureka” learning was possible, but rare.)

Stahl and Wilson's and CGCB's two-person games separated rationalizability and k -rationalizability from equilibrium.

They also make subjects' thinking fully two-sided, as in most applications.

Nagel's and HCW's games are like me thinking about the stock market: I know that "it" isn't thinking about me and that greatly simplifies my thinking, while also making it less informative about my thinking in other applications.

Most applications are like Warren Buffet thinking about the stock market and experiments should reflect the doubly recursive nature of strategic thinking.

This is important even if—especially if—we end up rejecting decision rules that reflect equilibrium's fixed-point reasoning.

(CGCB also used a procedural model of decision making to analyze their subjects' searches for hidden payoff information, as discussed in the book.)

In Stahl and Wilson's and CGCB's designs, a subject's profile of responses to the series of games discriminates among decision rules.

They analyzed the data within subjects, assuming that a single decision rule governed a subject's decisions across the games s/he played, which is more informative than fitting the aggregate decision distributions.

(The book discusses the econometric details, following CGCB.)

- Few subjects systematically violated rationality.
- But few respected equilibrium's fixed-point reasoning; and few were sophisticated in the sense of noisily best responding to others' decisions.
- Compliance with equilibrium and respect for k -rationalizability declined steadily with k and the complexity of the game (e.g. CGCB's Table II).

Costa-Gomes and Crawford 2006 *AER* (“CGC”)

Despite the greater discriminatory power of Stahl and Wilson's and CGCB's series of different games, their coarse decision spaces still yielded only weak separation of decision rules (including but not limited to L_{k+1} and D_k).

This forced them to rely more heavily than necessary on a priori specification of subjects' possible decision rules and on structural econometric estimation.

CGC further boosted power by combining Stahl and Wilson's and CGCB's series of different games with Nagel's and HCW's larger decision spaces.

As explained in the book and Crawford (2008), CGCB and CGC also monitored subjects' searches for hidden payoff information along with their decisions. They used a procedural, algorithmic model of decision-making to jointly analyze subjects' searches and decisions, further multiplying the power of their designs to reveal subjects' decision rules.

CGC's subjects played a series of 16 different two-person guessing games, again with random and anonymous pairing and no feedback to suppress repeated-game effects and rote learning. (Eureka learning was again rare.)

Here I discuss only the Baseline treatment, in which subjects played the games against other subjects. (The book discusses other treatments.)

In the games, each player has her/his own lower and upper limit; and her/his own target. A player's payoff increases with the closeness of her/his guess to her/his target times the other's guess. (Guesses were automatically adjusted to stay with limits, to enhance separation of information search implications.)

Unlike in Nagel's and HCW's n -person guessing games, subjects' targets and limits varied independently across players and games, with targets either both less than one, both greater than one, or mixed.

CGC's games are finitely dominance-solvable, with essentially unique equilibria determined by players' lower (upper) limits when the product of their targets $< (>) 1$.

Dominance-solvability usually takes many rounds, so subjects' thinking was unlikely to be shaped by it; their thinking was effectively doubly recursive.

Equilibrium's location is determined by the product of players' targets: more subtle than in Nagel's and HCW's games and can go against intuition.

(Even so, most subjects in a Robot/Trained Subjects treatment who were trained and rewarded for equilibrium decisions could identify them.)

The discontinuity of the equilibrium correspondence when the product of targets equals one stress-tests equilibrium and more strongly separates it from behaviorally plausible nonequilibrium decision rules.

A subject's guesses in the series of 16 large strategy spaces form a strategic "fingerprint" that often reveals her/his decision rule without econometrics.

(In the table, the rule *Sophisticated* best responds to the observed frequency distributions of other subjects' observed decisions—a non-vicious circle.)

Game Player's limits v opponent's limits, player's target v opponent's target	Player's guess for decision rule						
	<i>L1</i>	<i>L2</i>	<i>L3</i>	<i>D1</i>	<i>D2</i>	<i>Eq.</i>	<i>Soph.</i>
1. Limits [100, 500] v [100, 900], targets 0.7 v 0.5	350	105	122.5	122.5	122.5	100	122
2. Limits [100, 900] v [100, 500], targets 0.5 v 0.7	150	175	100	150	100	100	132
3. Limits [100, 900] v [300, 500], targets 0.5 v 0.7	200	175	150	200	150	150	162
4. Limits [300, 500] v [100, 900], targets 0.7 v 0.5	350	300	300	300	300	300	300
5. Limits [300, 500] v [300, 900], targets 1.5 v 1.3	500	500	500	500	500	500	500
6. Limits [300, 900] v [300, 500], targets 1.3 v 1.5	520	650	650	617.5	650	650	650
7. Limits [300, 900] v [300, 900], targets 1.3 v 1.3	780	900	900	838.5	900	900	900
8. Limits [300, 900] v [300, 900], targets 1.3 v 1.3	780	900	900	838.5	900	900	900
9. Limits [100, 900] v [100, 500], targets 0.5 v 1.5	150	250	112.5	162.5	131.25	100	187
10. Limits [100, 500] v [100, 900], targets 1.5 v 0.5	500	225	375	262.5	262.5	150	300
11. Limits [300, 900] v [100, 900], targets 0.7 v 1.3	350	546	318.5	451.5	423.15	300	420
12. Limits [300, 500] v [300, 900], targets 1.3 v 0.7	780	455	709.8	604.5	604.5	390	695
13. Limits [300, 500] v [100, 900], targets 0.7 v 1.5	350	420	367.5	420	420	500	420
14. Limits [100, 900] v [300, 500], targets 1.5 v 0.7	600	525	630	600	611.25	750	630
15. Limits [100, 500] v [100, 500], targets 0.7 v 1.5	210	315	220.5	227.5	227.5	350	262
16. Limits [100, 500] v [100, 500], targets 1.5 v 0.7	450	315	472.5	337.5	341.25	500	375

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

With 200 to 800 possible exact guesses in 16 games, a subject's exact compliance at these levels is very powerful evidence for a decision rule:

Uniformly random decisions track a given decision rule's decisions exactly (rounding to the nearest integer) in all 16 games once in 8.7×10^{42} tries. (I have not done the calculation for "7 or more" games but the odds are long.)

Because CGC's rules build in risk-neutral, selfish rationality and a uniform random *L0*, the *L1*, *L2*, and *L3* subjects' deviations from equilibrium must be due to their simplified models of others, not irrationality, risk aversion, altruism, spite, confusion, or exotic *L0*s.

CGC's other 45 subjects made guesses that were coherent, with violations of simple dominance less than 20% (versus 38% for random guesses).

But those subjects' guesses conformed less closely to a known decision rule.

CGC estimated subjects' decision rules econometrically subject by subject, assuming each subject's guesses are determined by a single decision rule in all 16 games. (Testing suggests that this assumption is rarely violated.)

To allow for the high frequencies of a rule's exact guesses, CGC allowed L_k for $k \geq 1$ to make mean-0, "spike-logit" decision errors, but otherwise with payoff-sensitive logit errors with estimated precisions, as in LQRE.

CGC's specification of possible decision rules balances the risk of omitting empirically relevant rules against that of overfitting via accidental correlations with included but empirically irrelevant rules.

CGC (Section II.D) operationalize this trade-off via a specification test that can detect empirically important omitted rules and accidental correlations. There are very few of those.

For the 45 subjects with less exact compliance with a decision rule, econometric estimates of their rules are concentrated on $L1$, $L2$, $L3$, and *Equilibrium* in roughly the same proportions as for the 43 with high exact compliance.

Unconstrained estimates of $L0$'s frequency are zero or small.

The estimates strongly favor L_{k+1} over D_k decision rules: Nagel's and HCW's spikes are level- k spikes and subjects respected k -rationalizability for small k as an unconscious by-product of following low-level level- k rules.

Decisions that don't follow $L1$, $L2$, $L3$, or *Equilibrium* rules lack discernible structure, so that it may be reasonable to treat them as errors.

CGC's results have been replicated by Fragiadakis, Knoepfle, and Niederle (2016) and Chegere, Fafchamps, and Falco (2024).

See however Georganas, Healy, and Weber's 2015 *JET* and Cooper et al.'s 2024 *AEJ Micro* results partly contradicting CGC's (discussed in the book).

Could a model such as QRE, CH, or level- k with a different $L0$ explain CGC's results better than CGC's model?

With regard to QRE, CGC's econometric model nests LQRE:

CGC's piecewise linear payoff function (pp. 1762-3) makes it optimal for a player to respond as if her/his beliefs about others' decisions were concentrated on the median of their distribution. Thus QRE (with *any* error distribution) reduces to *Equilibrium* (without noise).

In CGC's estimates *Equilibrium* subjects, even with noise, are rare (11/88).

Moreover, there is clear evidence that the subjects whose decisions were closest to equilibrium were following rules that mimic equilibrium in the games with unmixed targets but deviate with mixed targets (Crawford 2008).

Imposing homogeneity, a level- k or CH model with spike-logit errors fits CGC's subjects' decisions much better than LQRE or *Equilibrium*.

With regard to CH, level- k and CH $L1$ rules make the same decisions; and CGC's design doesn't strongly separate level- k and CH $L2$ or $L3$ decisions.

However, CH Lk rules for $k \geq 2$ are *not* self-informed: Their mean decisions unavoidably depend on the distribution of other players' levels, which, even if Poisson, has an empirical parameter that must be estimated or calibrated.

Thus, CH Lk rules for $k \geq 2$ implicitly assume that subjects know the distribution of others' levels, which they lack data to estimate.

That said, the practical import of CH's lack of self-informedness depends on how uncertain people are about the level distribution and how much that affects the models' predictions: In some games, not much; in others, a lot.

And CH rules for $k \geq 2$ are information omnivores, while level- k rules for $k \geq 2$ are on a level- k diet, which CGC's subjects followed pretty strictly. As a result CGC's analysis of subjects' searches strongly favors level- k over CH.