RES Easter School: Behavioural Economics
Brasenose College Oxford, 22-25 March 2015

Strategic Thinking I: Theory and Evidence

Vincent P. Crawford

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

(with thanks to Miguel Costa-Gomes, Nagore Iriberri, Colin Camerer, Teck-Hua Ho, and Juin-Kuan Chong)

UNIVERSITY OF OXFORD

UC San Diego

European Research Council
Established by the European Commission

Revised 20 March 2015

Strategic Thinking

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

(Nash equilibrium can be defined without reference to its interpretation, but it is best thought of as an "equilibrium in beliefs," in which players' strategies represent beliefs about others' strategies that are correct, given the rational strategy choices they imply.)

Strategic Thinking

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

(Nash equilibrium can be defined without reference to its interpretation, but it is best thought of as an "equilibrium in beliefs," in which players' strategies represent beliefs about others' strategies that are correct, given the rational strategy choices they imply.)

Equilibrium addresses the problem that in games, decision-theoretic rationality alone seldom restricts behavior enough to be useful: Even common knowledge of rationality implies only that players' strategies are rationalizable (Bernheim 1984 *Ecma* and Pearce 1984 *Ecma*), which often leaves behavior unrestricted.

Equilibrium makes more definite predictions by adding the "rational-expectations" assumption that players' beliefs are correct, and thus the same for all players.

Because many games have multiple equilibria, equilibrium refinements are often added, with the goal of making unique predictions based on theory alone.

Applications can safely assume equilibrium if a setting allows learning, if only long-run outcomes matter, and if equilibrium is unique.

Otherwise, outcomes are determined or influenced by strategic thinking.

Applications can safely assume equilibrium if a setting allows learning, if only long-run outcomes matter, and if equilibrium is unique.

Otherwise, outcomes are determined or influenced by strategic thinking.

Epistemic game theory provides conditions under which thinking based on iterated or common knowledge of rationality and beliefs will focus players' beliefs on an equilibrium, even in their initial responses to a game.

Applications can safely assume equilibrium if a setting allows learning, if only long-run outcomes matter, and if equilibrium is unique.

Otherwise, outcomes are determined or influenced by strategic thinking.

Epistemic game theory provides conditions under which thinking based on iterated or common knowledge of rationality and beliefs will focus players' beliefs on an equilibrium, even in their initial responses to a game.

But in many games such thinking is complex, and even people who are capable of it may doubt that others are capable, and so on, making a thinking justification for equilibrium behaviorally implausible.

People may then find simpler, nonequilibrium ways of thinking about the game.

Modeling strategic thinking more accurately can yield several benefits:

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

- It can challenge conclusions based on equilibrium or refinements in games where equilibrium is implausible without learning.

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

- It can elucidate learning from imperfect analogies, where assumptions about cognition determine which analogies players recognize.

Modeling strategic thinking more accurately can yield several benefits:

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

- It can challenge conclusions based on equilibrium or refinements in games where equilibrium is implausible without learning.

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

- It can elucidate learning from imperfect analogies, where assumptions about cognition determine which analogies players recognize.

However, even those who grant the desirability of modeling strategic thinking more accurately may doubt its feasibility:

- How can any model out-predict a rational-expectations notion?

- And how can one identify such a model among the huge set of possible non-equilibrium models?

But…

- There is now a large body of experimental research that studies strategic thinking by eliciting subjects' initial responses to games played as if in isolation (surveyed in Crawford, Costa-Gomes, and Iriberri 2013 *JEL*).

- The evidence suggests that people's thinking in novel or complex games does not follow the fixed-point or indefinitely iterated dominance reasoning that equilibrium often requires.

But…

- There is now a large body of experimental research that studies strategic thinking by eliciting subjects' initial responses to games played as if in isolation (surveyed in Crawford, Costa-Gomes, and Iriberri 2013 *JEL*).

- The evidence suggests that people's thinking in novel or complex games does not follow the fixed-point or indefinitely iterated dominance reasoning that equilibrium often requires.

  (Learning can still make people converge to something that *we* need fixed-point reasoning to characterize; the claim is that fixed-point reasoning doesn't directly describe people's thinking.)

But…

- There is now a large body of experimental research that studies strategic thinking by eliciting subjects' initial responses to games played as if in isolation (surveyed in Crawford, Costa-Gomes, and Iriberri 2013 *JEL*).

- The evidence suggests that people's thinking in novel or complex games does not follow the fixed-point or indefinitely iterated dominance reasoning that equilibrium often requires.

  (Learning can still make people converge to something that *we* need fixed-point reasoning to characterize; the claim is that fixed-point reasoning doesn't directly describe people's thinking.)

- The evidence also identifies systematic patterns that can be modeled.

In this lecture I will compare alternative models and then discuss the evidence.

I focus on normal-form games but discuss extensive-form games at the end.

Alternative normal-form models of strategic thinking

The leading alternatives are:

- Adding noise to equilibrium predictions ("equilibrium plus noise"), plus refinements such as risk- or payoff-dominance or "global games".

- Finitely iterated dominance and *k*-rationalizability (for strict dominance the two notions are equivalent in the two-person games I focus on here).

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

- "Level-*k*" models.

- Cognitive hierarchy ("CH") models.

Alternative normal-form models of strategic thinking

The leading alternatives are:

- Adding noise to equilibrium predictions ("equilibrium plus noise"), plus refinements such as risk- or payoff-dominance or "global games".

- Finitely iterated dominance and *k*-rationalizability (for strict dominance the two notions are equivalent in the two-person games I focus on here).

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

- "Level-*k*" models.

- Cognitive hierarchy ("CH") models.

All are general models of strategic behavior, applicable in any norma-form game.

All allow but do not assume equilibrium in all games.

The last three not only predict that deviations from equilibrium will sometimes occur, but also which games evoke them and the forms they are likely to take.

A. Equilibrium Plus Noise

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

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

Equilibrium plus noise often describes observed behavior well but sometimes misses systematic patterns in subjects' deviations from equilibrium.

B. Rationalizability and *k*-rationalizability

In two-person games (with differences in *n*-person games that are unimportant here), the implications of common or finitely iterated knowledge of players' rationality (without further restrictions on beliefs) are captured by indefinitely or finitely iterated strict dominance.

Their implications for players' strategy choices are captured by rationalizability or *k*-rationalizability, set-valued restrictions on individual player's choices (unlike equilibrium, which restricts the *relationship* between players' choices).

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

B. Rationalizability and *k*-rationalizability

In two-person games (with differences in *n*-person games that are unimportant here), the implications of common or finitely iterated knowledge of players' rationality (without further restrictions on beliefs) are captured by indefinitely or finitely iterated strict dominance.

Their implications for players' strategy choices are captured by rationalizability or *k*-rationalizability, set-valued restrictions on individual player's choices (unlike equilibrium, which restricts the *relationship* between players' choices).

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

Dominance-solvable game

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

Unique equilibrium without dominance

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

In games that are strictly dominance-solvable in k rounds, k-rationalizability implies that players have the same beliefs—with a qualification for mixed-strategy equilibria that is unimportant here—so any combination of k-rationalizable strategies is in equilibrium as in the first game above.

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

In games that are strictly dominance-solvable in *k* rounds, *k*-rationalizability implies that players have the same beliefs—with a qualification for mixed-strategy equilibria that is unimportant here—so any combination of *k*-rationalizable strategies is in equilibrium as in the first game above.

In other games, *k*-rationalizability and rationalizability allow deviations from equilibrium as in the second game above, where there is a “helix” of beliefs to make any strategy combination consistent with common knowledge of rationality.

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

Finitely iterated dominance and *k*-rationalizability are often consistent with the systematic patterns in subjects’ deviations from equilibrium.

C. Quantal Response Equilibrium ("QRE") and Logit QRE ("LQRE")

QRE seeks to capture the payoff-sensitivity of deviations from equilibrium that equilibrium plus noise sometimes misses.

In a QRE players' decisions are noisy, with the probability density of each decision increasing in its expected payoff.

But importantly, unlike in equilibrium plus noise (or in level-*k* and CH models), the payoffs are evaluated taking the noisiness of others' decisions into account.

C. Quantal Response Equilibrium ("QRE") and Logit QRE ("LQRE")

QRE seeks to capture the payoff-sensitivity of deviations from equilibrium that equilibrium plus noise sometimes misses.

In a QRE players' decisions are noisy, with the probability density of each decision increasing in its expected payoff.

But importantly, unlike in equilibrium plus noise (or in level-*k* and CH models), the payoffs are evaluated taking the noisiness of others' decisions into account.

A QRE model is closed by specifying a response distribution, which is logit in almost all applications. As its precision increases, QRE approaches equilibrium; and as its precision approaches 0, QRE approaches uniform randomization.

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

C. Quantal Response Equilibrium ("QRE") and Logit QRE ("LQRE")

QRE seeks to capture the payoff-sensitivity of deviations from equilibrium that equilibrium plus noise sometimes misses.

In a QRE players' decisions are noisy, with the probability density of each decision increasing in its expected payoff.

But importantly, unlike in equilibrium plus noise (or in level-*k* and CH models), the payoffs are evaluated taking the noisiness of others' decisions into account.

A QRE model is closed by specifying a response distribution, which is logit in almost all applications. As its precision increases, QRE approaches equilibrium; and as its precision approaches 0, QRE approaches uniform randomization.

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

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

The resulting QRE, "logit QRE", or "LQRE" implies error distributions that respond to out-of-equilibrium payoffs in plausible ways; and often fits subjects' deviations from equilibrium better than an equilibrium-plus-noise model.

D. Level-*k* Models

In a level-*k* model people follow rules of thumb that:

- Anchor their beliefs in a naïve model of others' responses, called *L0*, usually uniform random over the feasible decisions (sometimes truthful, etc.); and

- Adjust their beliefs via a small number (*k*) of iterated best responses, so *L1* best responds to *L0*, *L2* to *L1*, and so on.

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

D. Level-*k* Models

In a level-*k* model people follow rules of thumb that:

- Anchor their beliefs in a naïve model of others' responses, called *L0*, usually uniform random over the feasible decisions (sometimes truthful, etc.); and

- Adjust their beliefs via a small number (*k*) of iterated best responses, so *L1* best responds to *L0*, *L2* to *L1*, and so on.

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

- *Lk* (for $k > 0$) is decision-theoretically rational, with an accurate model of the game; it departs from equilibrium only in deriving its beliefs from an oversimplified model of others' responses.

D. Level-*k* Models

In a level-*k* model people follow rules of thumb that:

- Anchor their beliefs in a naïve model of others' responses, called *L0*, usually uniform random over the feasible decisions (sometimes truthful, etc.); and

- Adjust their beliefs via a small number (*k*) of iterated best responses, so *L1* best responds to *L0*, *L2* to *L1*, and so on.

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

- *Lk* (for $k > 0$) is decision-theoretically rational, with an accurate model of the game; it departs from equilibrium only in deriving its beliefs from an oversimplified model of others' responses.

- *Lk* (for $k > 0$) respects *k*-rationalizability, hence in two-person games its decisions survive *k* rounds of iterated dominance.

- Thus *Lk* mimics equilibrium decisions in *k*-dominance-solvable games, but may deviate systematically in more complex games.

- A level-*k* model (with 0 weight on *L0*) can be viewed as a heterogeneity-tolerant refinement of *k*-rationalizability (but makes precise predictions).

The population level frequencies are treated as behavioral parameters and either estimated from the data or calibrated from previous estimates.

The estimated frequency of *L0* is normally small or 0 and the distribution of levels is concentrated on *L1*, *L2*, and *L3*.

The population level frequencies are treated as behavioral parameters and either estimated from the data or calibrated from previous estimates.

The estimated frequency of *L0* is normally small or 0 and the distribution of levels is concentrated on *L1*, *L2*, and *L3*.

It is usually assumed that *Lk* types make errors (rather than treating *L0* as an error structure), often taken to be logit with estimated precision, as in LQRE.

Thus the probability density of each type’s decision is increasing in its expected payoff, evaluated using the type’s model of others’ decisions.

The population level frequencies are treated as behavioral parameters and either estimated from the data or calibrated from previous estimates.

The estimated frequency of *L0* is normally small or 0 and the distribution of levels is concentrated on *L1*, *L2*, and *L3*.

It is usually assumed that *Lk* types make errors (rather than treating *L0* as an error structure), often taken to be logit with estimated precision, as in LQRE.

Thus the probability density of each type's decision is increasing in its expected payoff, evaluated using the type's model of others' decisions.

Unlike QRE, a level-*k* model treats deviations from equilibrium as an integral, deterministic part of the structure, rather than as responses to errors.

But it too responds to out-of-equilibrium payoffs in plausible ways, and often fits subjects' deviations from equilibrium better than an equilibrium-plus-noise model.

E. Cognitive Hierarchy ("CH") Models

In a CH model, a close relative of level-*k*, *Lk* best responds not to *Lk-1* alone but to a mixture of lower-level types; and the type frequencies are assumed to be determined by Poisson distribution with an estimated parameter (the average *k*).

(To an econometrician this specification may seem more natural than a level-*k* specification; but which better describes non-econometricans' thinking is an empirical question, on which the jury is still out.)

E. Cognitive Hierarchy ("CH") Models

In a CH model, a close relative of level-*k*, *Lk* best responds not to *Lk-1* alone but to a mixture of lower-level types; and the type frequencies are assumed to be determined by Poisson distribution with an estimated parameter (the average *k*).

(To an econometrician this specification may seem more natural than a level-*k* specification; but which better describes non-econometricans' thinking is an empirical question, on which the jury is still out.)

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

CH types make undominated decisions, but unlike level-*k* types, but a CH *L2* or higher might not respect *k*–rationalizability.

E. Cognitive Hierarchy ("CH") Models

In a CH model, a close relative of level-*k*, *Lk* best responds not to *Lk-1* alone but to a mixture of lower-level types; and the type frequencies are assumed to be determined by Poisson distribution with an estimated parameter (the average *k*).

(To an econometrician this specification may seem more natural than a level-*k* specification; but which better describes non-econometricans' thinking is an empirical question, on which the jury is still out.)

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

CH types make undominated decisions, but unlike level-*k* types, but a CH *L2* or higher might not respect *k*–rationalizability.

In a CH model *L1* and higher types are usually assumed not to make errors.

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

A CH model also responds to out-of-equilibrium payoffs in plausible ways and fits subjects' deviations from equilibrium better than an equilibrium-plus-noise model.

Experimental evidence from guessing and other normal-form games

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

—John Maynard Keynes, *The General Theory of Employment, Interest, and Money*

Experimental evidence from guessing and other normal-form games

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

—John Maynard Keynes, *The General Theory of Employment, Interest, and Money*

The Keynes quotation suggests thinking in which players anchor their beliefs in a model of others’ naive reactions and then iterate best responses a finite, heterogeneous number of times. (Keynes’ “fourth, fifth and higher degrees” are way too high to be realistic, but they may be only a coy reference to himself.)

I first discuss Nagel's (1995 *AER*); Ho, Camerer, and Weigelt's (1998 *AER*; "HCW"); and Bosch-Domènech et al.'s (2002 *AER*) analyses of *n*-person guessing games, which give a simple introduction to this literature.

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

I first discuss Nagel's (1995 *AER*); Ho, Camerer, and Weigelt's (1998 *AER*; "HCW"); and Bosch-Domènech et al.'s (2002 *AER*) analyses of *n*-person guessing games, which give a simple introduction to this literature.

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

Recall that in most experiments that study strategic thinking, game-theoretically naïve subjects play series of different but related games with randomly, anonymously paired partners and no feedback.

The goal is to suppress learning and repeated-game effects, to elicit subjects' initial responses to each game, "uncontaminated" by learning.

I first discuss Nagel's (1995 *AER*); Ho, Camerer, and Weigelt's (1998 *AER*; "HCW"); and Bosch-Domènech et al.'s (2002 *AER*) analyses of *n*-person guessing games, which give a simple introduction to this literature.

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

Recall that in most experiments that study strategic thinking, game-theoretically naïve subjects play series of different but related games with randomly, anonymously paired partners and no feedback.

The goal is to suppress learning and repeated-game effects, to elicit subjects' initial responses to each game, "uncontaminated" by learning.

CGC's design follows this practice.

Nagel's and HCW's subject groups repeatedly played *n*-person games, but in their larger groups it is plausible that subjects treated future influences as negligible, so that their first-period responses can be viewed as initial responses to games played as if in isolation, "uncontaminated" by strategic teaching.

In Nagel's and HCW's *n*-person guessing games, *n* subjects (*n* = 15-18 in Nagel, *n* = 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. (2002 *AER*), essentially the same games were played in the field, by more than 7500 volunteers recruited from subscribers of the newspapers *Financial Times*, *Spektrum der Wissenchaft*, or *Expansión*.

In Nagel's and HCW's *n*-person guessing games, *n* subjects (*n* = 15-18 in Nagel, *n* = 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. (2002 *AER*), essentially the same games were played in the field, by more than 7500 volunteers recruited from subscribers of the newspapers *Financial Times*, *Spektrum der Wissenchaft*, or *Expansión*.

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

In Nagel's and HCW's *n*-person guessing games, *n* subjects (*n* = 15-18 in Nagel, *n* = 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. (2002 *AER*), essentially the same games were played in the field, by more than 7500 volunteers recruited from subscribers of the newspapers *Financial Times*, *Spektrum der Wissenchaft*, or *Expansión*.

In each case the subject who guessed closest to a target (*p* = 1/2, 2/3, or 4/3 in Nagel; *p* = 0.7, 0.9, 1.1, or 1.3 in 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 game theory.

For definiteness, consider Nagel's leading treatment:

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

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

- The structure was publicly announced.

For definiteness, consider Nagel's leading treatment:

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

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

- The structure was publicly announced.

These games have a unique equilibrium, in which all players guess 0.

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

For definiteness, consider Nagel's leading treatment:

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

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

- The structure was publicly announced.

These games have a unique equilibrium, in which all players guess 0.

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

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

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

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

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

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

The rationality-based argument for this “all–0” equilibrium is stronger than many equilibrium arguments: it depends “only” on iterated knowledge of rationality (though not finitely), 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?

The rationality-based argument for this “all–0” equilibrium is stronger than many equilibrium arguments: it depends “only” on iterated knowledge of rationality (though not finitely), 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 so, and Bosch-Domènech et al.’s (who had much more time to reflect, and who could consult with others) fairly rarely did so.

Most subjects’ guesses respected from 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) and Bosch-Domènech et al.’s Figure 1 illustrate these points:

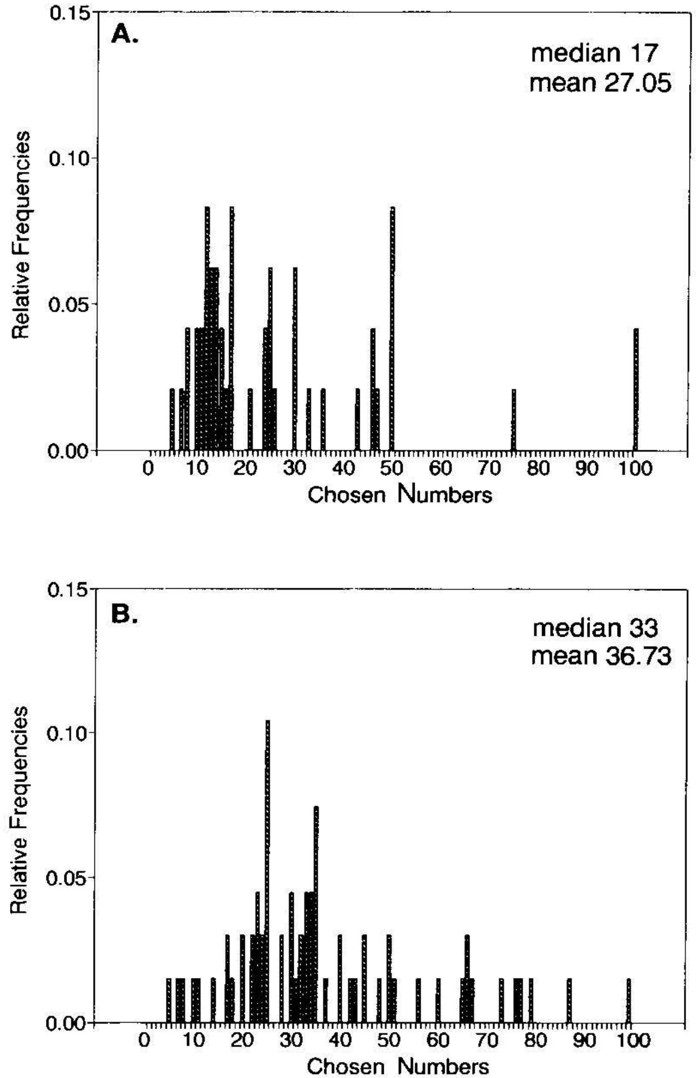

Nagel's Figure 1 (top *p* = 1/2, bottom *p* = 2/3)

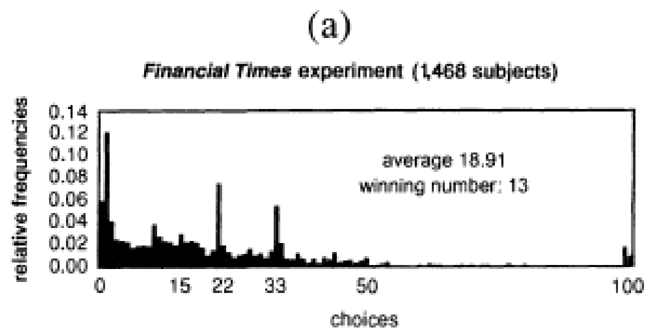

FIGURE 1. RELATIVE FREQUENCIES OF CHOICES IN THREE NEWSPAPER EXPERIMENTS

Bosch-Domènech et al.'s Figure 1 ($p = 2/3$)

These data resemble neither equilibrium plus noise nor QRE.

(For logit or any reasonable error distribution, though by Haile et al’s 2008 *AER* result we could make the data an exact QRE for an unreasonable distribution.)

These data resemble neither equilibrium plus noise nor QRE.

(For logit or any reasonable error distribution, though by Haile et al's 2008 *AER* result we could make the data an exact QRE for an unreasonable distribution.)

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

The distributions of subjects' guesses have spikes that track $50p^k$ for k = 1, 2, 3 across treatments with various *p*s, respecting 0-3 rounds of iterated dominance, so that guesses respect *k*-rationalizability for at most small values of *k*.

These data resemble neither equilibrium plus noise nor QRE.

(For logit or any reasonable error distribution, though by Haile et al's 2008 *AER* result we could make the data an exact QRE for an unreasonable distribution.)

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

The distributions of subjects' guesses have spikes that track $50p^k$ for $k = 1, 2, 3$ across treatments with various *p*s, respecting 0-3 rounds of iterated dominance, so that guesses respect *k*-rationalizability for at most small values of *k*.

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

But separation of decision rules is comparatively weak in these games, despite their large strategy spaces, because each subject played only one game (although there was useful between-subjects variation across treatments).

But separation of decision rules is comparatively weak in these games, despite their large strategy spaces, because each subject played only one game (although there was useful between-subjects variation across treatments).

For example, many theorists infer from the spikes in the data that subjects explicitly performed finitely iterated dominance, as we teach students to do.

In this interpretation, called *Dk*, a player does *k* rounds of iterated dominance and then best responds to a uniform prior over other players' remaining strategies.

But separation of decision rules is comparatively weak in these games, despite their large strategy spaces, because each subject played only one game (although there was useful between-subjects variation across treatments).

For example, many theorists infer from the spikes in the data that subjects explicitly performed finitely iterated dominance, as we teach students to do.

In this interpretation, called *Dk*, a player does *k* rounds of iterated dominance and then best responds to a uniform prior over other players' remaining strategies.

Dk respects *k*-rationalizability by construction, adding a selection via the prior.

In Nagel's and HCW's games *Dk* guesses $([0+100p^k]/2)p \equiv 50p^k$.

(In games without dominance *Dk* for $k > 0$ makes the same decisions as *L1*.)

Inferring finitely iterated dominance is premature, because in these games *Lk+1* makes exactly the same guesses: $[(0+100)/2]p^{k+1} \equiv ([0+100p^k]/2)p \equiv 50p^k$.

(Note that it is *Lk*+1 that is *Dk*'s cousin, not *Lk*: an unimportant quirk of notation.)

Lk+1 respects *k*-rationalizability, often with a different selection, though not here.

Inferring finitely iterated dominance is premature, because in these games *Lk+1* makes exactly the same guesses: $[(0+100)/2]p^{k+1} \equiv ([0+100p^k]/2)p \equiv 50p^k$.

(Note that it is *Lk*+1 that is *Dk*'s cousin, not *Lk*: an unimportant quirk of notation.)

Lk+1 respects *k*-rationalizability, often with a different selection, though not here.

In other experiments, including some of HCW's, *Dk*'s and *Lk*+1's guesses are weakly separated, and the results are inconclusive on which rule subjects follow.

In CGC's experiment discussed next, *Dk*'s and *Lk*+1's guesses are strongly separated, and the results very strongly favor *Lk+1* rules:

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

Aside on specifying *L0* in *n*-person games

These lectures focus mainly on two-person games, but in *n*-person games it can matter whether *L0* is independent across players or correlated.

The limited available evidence (e.g. HCW and Costa-Gomes, Crawford, and Iriberri 2009 *JEEA*) suggests that most subjects have highly correlated, "representative agent"-like, models of others.

(This may be related to the phenomenon of informational naiveté discussed elsewhere in these lectures.)

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

End of aside

CGC's (2006 *AER*) two-person guessing game experiments

In CGC's design (as in Stahl and Wilson's 1994 *JEBO*, 1995 *GEB* and Costa-Gomes, Crawford, and Broseta's 2001 *Ecma* designs), each subject played a series of different games, with learning and repeated-game effects suppressed.

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

CGC's design combines the variation of games with Nagel's and HCW's large strategy spaces to increase the power of subjects' decisions to reveal thinking.

Their two-person games more fully engage strategic thinking, because subjects know their partners do not view themselves as negligible parts of the interaction.

Two-person games also avoid the ambiguity of interpretation noted above regarding whether *L0* is independent across players or correlated.

Subjects were randomly and anonymously paired to play a series of 16 different two-person guessing games, without feedback.

Each player had his own lower and upper limit, both strictly positive.

(Positive limits make the games finitely dominance-solvable, in 2 to 53 rounds.)

Each player also had 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 varied independently across players and games, with targets both less than one, both greater than one, or "mixed".

(In Nagel's and HCW's designs the targets and limits were always the same for both players, and varied at most between subjects across treatments.)

These games have essentially unique equilibria, determined by players' lower (upper) limits when the product of targets is less (greater) than one.

These games have essentially unique equilibria, determined by players' lower (upper) limits when the product of targets is less (greater) than one.

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

The product of targets is 1.05 > 1, and the equilibrium is therefore determined by players' upper limits: The first player guesses his upper limit 500, but the second player guesses 750 (= 500 × his target 1.5), below his upper limit 900.

No guess is dominated for the first player, but any guess outside [450, 750] is dominated for the second player. Given that, any guess outside [315, 500] is iteratively dominated for the first player, and so on until the equilibrium at (500, 750) is reached after 22 iterations.

These games have essentially unique equilibria, determined by players' lower (upper) limits when the product of targets is less (greater) than one.

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

The product of targets is 1.05 > 1, and the equilibrium is therefore determined by players' upper limits: The first player guesses his upper limit 500, but the second player guesses 750 (= 500 × his target 1.5), below his upper limit 900.

No guess is dominated for the first player, but any guess outside [450, 750] is dominated for the second player. Given that, any guess outside [315, 500] is iteratively dominated for the first player, and so on until the equilibrium at (500, 750) is reached after 22 iterations.

The discontinuity of the equilibrium correspondence stress-tests equilibrium and enhances its separation from other types: Equilibrium responds much more strongly to the difference between games where the product of targets is slightly greater or less than one than behaviorally plausible other decision rules do.

CGC'S data analysis
As suggested by previous work, CGC assumed that each subject's guesses were determined, up to errors, by a single type in all 16 games.

CGC's types all build in risk-neutrality and rule out social preferences:

- *L1*, *L2*, and *L3*, as defined above

  (In CGC's design, *L2* and *L3* are not separated from their CH counterparts; see CGC, footnote 36.)

- *D1* and *D2*, as defined above

- *Equilibrium*, which makes its (essentially unique) equilibrium decision

- *Sophisticated*, which best responds to the probability distributions of others' decisions, proxied by the observed population frequencies

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

● CGC's design yields very strong separation of types' guesses.

(a_i,b_i, and p_i are a player's lower limit, upper limit, and target, and a_j,b_j, and p_j are his partner's.)

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

Game	a_i	b_i	p_i	a_j	b_j	p_j	*L1*	*L2*	*L3*	*D1*	*D2*	*Eq*	*So*
1	100	900	1.5	300	500	0.7	600	525	630	600	611.25	750	630
2	300	900	1.3	300	500	1.5	520	650	650	617.5	650	650	650
3	300	900	1.3	300	900	1.3	780	900	900	838.5	900	900	900
4	300	900	0.7	100	900	1.3	350	546	318.5	451.5	423.15	300	420
5	100	500	1.5	100	500	0.7	450	315	472.5	337.5	341.25	500	375
6	100	500	0.7	100	900	0.5	350	105	122.5	122.5	122.5	100	122
7	100	500	0.7	100	500	1.5	210	315	220.5	227.5	227.5	350	262
8	300	500	0.7	100	900	1.5	350	420	367.5	420	420	500	420
9	300	500	1.5	300	900	1.3	500	500	500	500	500	500	500
10	300	500	0.7	100	900	0.5	350	300	300	300	300	300	300
11	100	500	1.5	100	900	0.5	500	225	375	262.5	262.5	150	300
12	300	900	1.3	300	900	1.3	780	900	900	838.5	900	900	900
13	100	900	1.3	300	900	0.7	780	455	709.8	604.5	604.5	390	695
14	100	900	0.5	300	500	0.7	200	175	150	200	150	150	162
15	100	900	0.5	100	500	0.7	150	175	100	150	100	100	132
16	100	900	0.5	100	500	1.5	150	250	112.5	162.5	131.25	100	187

- Of CGC's 88 subjects, 43 subject's guesses complied *exactly* with one type's guesses in from 7 to 16 games: 20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*.

- Of CGC’s 88 subjects, 43 subject’s guesses complied *exactly* with one type’s guesses in from 7 to 16 games: 20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*.

- The size of CGC’s strategy spaces, with 200 to 800 possible exact (within 0.5) guesses across 16 different games, makes exact compliance very powerful evidence for a subject’s apparent type, so that his type can confidently be identified from his guessing “fingerprint”, without econometrics (with a qualification for *Equilibrium*, discussed below).

  (By contrast, there are usually many possible reasons for choosing one of a few strategies in a small matrix game; and even in Nagel’s games, rules as cognitively disparate as *Dk* and *Lk*+1 make identical guesses.)

- Of CGC's 88 subjects, 43 subject's guesses complied *exactly* with one type's guesses in from 7 to 16 games: 20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*.

- The size of CGC's strategy spaces, with 200 to 800 possible exact (within 0.5) guesses across 16 different games, makes exact compliance very powerful evidence for a subject's apparent type, so that his type can confidently be identified from his guessing "fingerprint", without econometrics (with a qualification for *Equilibrium*, discussed below).

  (By contrast, there are usually many possible reasons for choosing one of a few strategies in a small matrix game; and even in Nagel's games, rules as cognitively disparate as *Dk* and *Lk*+1 make identical guesses.)

- Further, because CGC's types build in risk-neutral, self-interested rationality, the deviations from equilibrium of subjects with high exact compliance with level-*k* types must be caused not by irrationality, risk aversion, altruism, spite, or confusion, but by their simplified models of others.

- Of CGC's 88 subjects, 43 subject's guesses complied *exactly* with one type's guesses in from 7 to 16 games: 20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*.

- The size of CGC's strategy spaces, with 200 to 800 possible exact (within 0.5) guesses across 16 different games, makes exact compliance very powerful evidence for a subject's apparent type, so that his type can confidently be identified from his guessing "fingerprint", without econometrics (with a qualification for *Equilibrium*, discussed below).

  (By contrast, there are usually many possible reasons for choosing one of a few strategies in a small matrix game; and even in Nagel's games, rules as cognitively disparate as *Dk* and *Lk*+1 make identical guesses.)

- Further, because CGC's types build in risk-neutral, self-interested rationality, the deviations from equilibrium of subjects with high exact compliance with level-*k* types must be caused not by irrationality, risk aversion, altruism, spite, or confusion, but by their simplified models of others.

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

- Of CGC's 88 subjects, 43 subject's guesses complied *exactly* with one type's guesses in from 7 to 16 games: 20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*.

- The size of CGC's strategy spaces, with 200 to 800 possible exact (within 0.5) guesses across 16 different games, makes exact compliance very powerful evidence for a subject's apparent type, so that his type can confidently be identified from his guessing "fingerprint", without econometrics (with a qualification for *Equilibrium*, discussed below).

  (By contrast, there are usually many possible reasons for choosing one of a few strategies in a small matrix game; and even in Nagel's games, rules as cognitively disparate as *Dk* and *Lk*+1 make identical guesses.)

- Further, because CGC's types build in risk-neutral, self-interested rationality, the deviations from equilibrium of subjects with high exact compliance with level-*k* types must be caused not by irrationality, risk aversion, altruism, spite, or confusion, but by their simplified models of others.

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

- But all but 14 violated simple dominance in less than 20% of the games (versus 38% for random guesses); and econometric estimates of their types are concentrated on *L1*, *L2*, *L3*, and *Equilibrium* in similar proportions.

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

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

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

Aside on apparent *Equilibrium* subjects

Doubts remain about the subjects with high exact compliance with *Equilibrium*; consider the 8 subjects with near-*Equilibrium* fingerprints.

- 85% of those 8 subjects' deviations from equilibrium were in games with mixed targets (on the right half of Figure 4 below).

  (By contrast, *L1*s' exact compliance with *L1* guesses is almost the same with and without mixed targets; though *L2*s' and *L3*s' is lower with mixed targets.)

- Those subjects, whose exact *Equilibrium* compliance was off the chart, are following a rule that only mimics *Equilibrium*, and only without mixed targets.

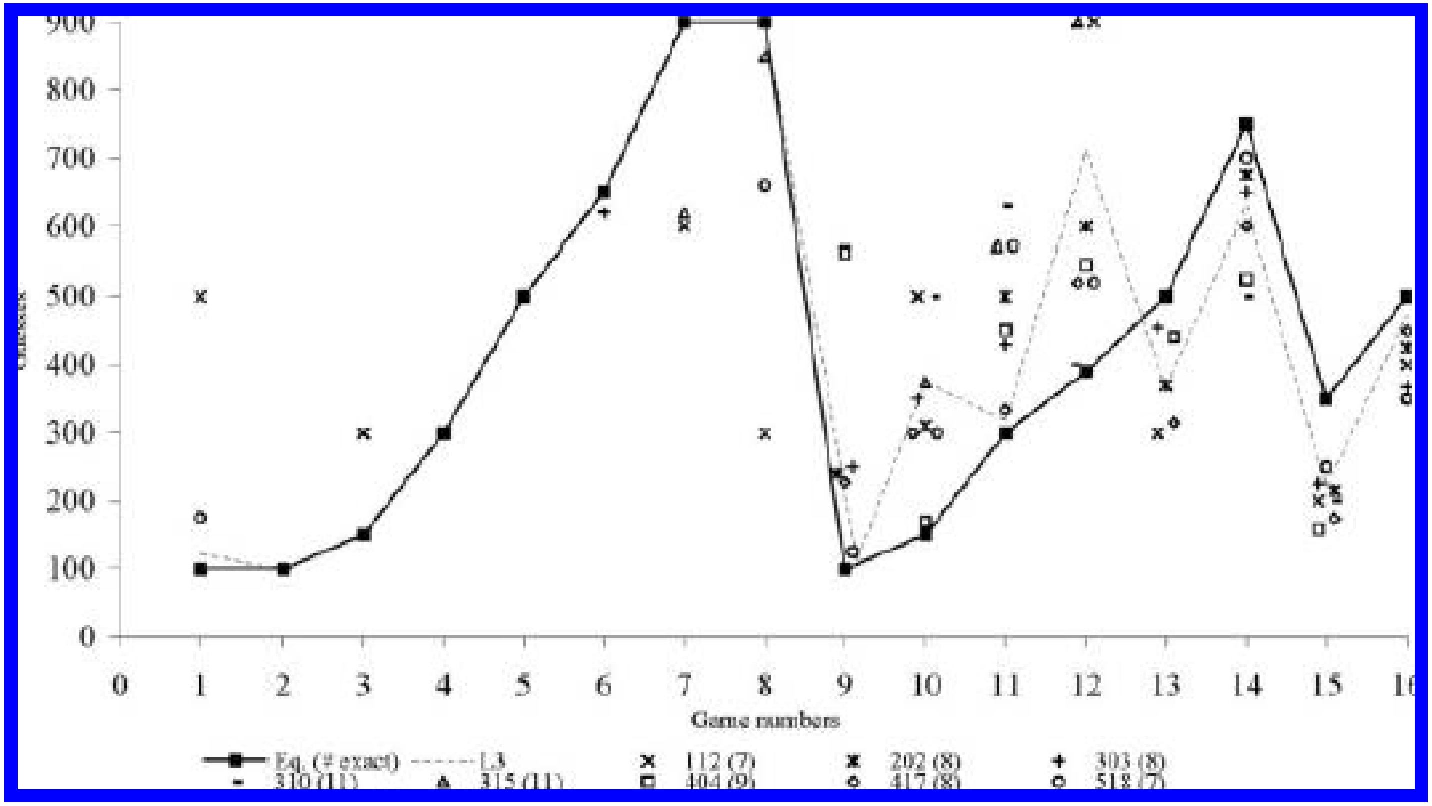

FIGURE 4. "FINGERPRINTS" OF EIGHT APPARENT *EQUILIBRIUM* SUBJECTS

Notes: Only deviations from *Equilibrium*'s guesses are shown. Of these subjects' 128 guesses, 69 (54 percent) were exact *Equilibrium* guesses.

Aside on apparent *Equilibrium* subjects

Doubts remain about the subjects with high exact compliance with *Equilibrium*; consider the 8 subjects with near-*Equilibrium* fingerprints.

- 85% of those 8 subjects' deviations from equilibrium were in games with mixed targets (on the right half of Figure 4 below).

  (By contrast, *L1*s' exact compliance with *L1* guesses is almost the same with and without mixed targets; though *L2*s' and *L3*s' is lower with mixed targets.)

- Those subjects, whose exact *Equilibrium* compliance was off the chart, are following a rule that only mimics *Equilibrium*, and only without mixed targets.

- Yet the ways we teach people to identify equilibria (best-response dynamics, equilibrium checking, iterated dominance) work equally well with and without mixed targets: Whatever those subjects were doing, it's something we haven't thought of (the questionnaires don't help).

- *Equilibrium* Robot/Trained Subjects' compliance was as high with as without mixed targets, so training eliminates whatever those subjects were doing.

End of aside

To sum up, CGC's results show that a level-*k* or CH model with a uniform random *L0* and *L1*, *L2*, *L3*, and possibly *Equilibrium* subjects explains a large fraction of subjects' deviations from equilibrium.

Although about half of CGC's subjects' deviations from equilibrium remain unexplained by such a model, CGC's specification test (discussed in other slides) suggests that there are not significant numbers of subjects following other rules.

CGC's conclusions are consistent with the results of most previous studies of initial responses to normal-form games with neutral framing, just more precise.

(Some later studies have questioned their portability, but the jury is still out.)

Thus it may be reasonable to take the part of the structure that can be identified as a basis for modeling, and treat the remaining deviations as errors.

Extensive-form games

Beard and Beil (1994 *MS*) start with Rosenthal's (1981 *JET*) game, framed in extensive form (displayed here in normal form, with A the row player).

	l	**r**
L	**y** **x**	**y** **x**
R	**v** **0**	**w** **z**

- “Real-time play” (as opposed to “strategy method”).

- Player A can opt out (L) with payoffs x for A and y for B; or give player B the move (R). B then has two choices: l with payoffs 0 for A, v for B; or r with payoffs z for A and w for B; where z > x and w > v (y > or < w).

- The unique subgame-perfect equilibrium is (R, r) (which uniquely survives iterated weak dominance), but A players who think B is not certain to play r are tempted by L; thus the game tests for reliance on other’s dominance.

Assuming subjects maximize expected pecuniary payoffs (no risk aversion, no social preferences), intuition suggests:

(H1) A players should be more willing to play R when x is lower (R is less risky).

(H2) A players should be more willing to play R when w – v is higher (B has more incentive to choose r).

(H3) A players should be more willing to play R when y is lower (B is less likely to resent A's choice of R and choose l), or w and v are higher (B is more likely to reciprocate A's choice of R by choosing r).

Subgame-perfect equilibrium doesn't say any of (H1)-(H3), but extensive- or normal-form QRE might (McKelvey and Palfrey 1995 *GEB,* 1998 *EE*), maybe with social preferences that depend on others' choices (Rabin 1993 AER).

Beard and Beil's treatments test these intuitions, mostly holding the critical probability that B chooses l that makes A indifferent between L and R near 0:

- Can test (H1) by comparing Treatments 1, 2, and 3.
- Can test (H2) by comparing Treatments 1 and 4.
- Can test (H3) by comparing Treatments 1 and 5.

Beard and Beil's Treatments

(A,B) $payoffs		Player *A* plays *R*		
Treatment	Player *A* plays *L*	Player *B* plays *l*	Player *B* plays *r*	Critical probability+
1	(9.75, 3.00)	(3.00, 4.75)	(10.00, 5.00)	3.57%
2	(9.00, 3.00)	(3.00, 4.75)	(10.00, 5.00)	14.29%
3	(7.00, 3.00)	(3.00, 4.75)	(10.00, 5.00)	42.86%
4	(9.75, 3.00)	(3.00, 3.00)	(10.00, 5.00)	3.57%
5	(9.75, 6.00)	(3.00, 4.75)	(10.00, 5.00)	3.57%
6	(9.75, 5.00)	(5.00, 9.75)	(10.00, 10.00)	5.00%
7*	(58.50, 18.00)	(18.00, 28.50)	(60.00, 30.00)	3.57%

*Subjects in treatment 7 received the stated payoffs only with probability 1/6.

+Probability of nonmaximizing play of l by B that a risk-neutral payoff-maximizing A subject needs for L to be optimal.

Results

- 97.8% of B subjects made choices that maximized own pecuniary rewards.
- Despite predictability of B subjects' decisions, most A subjects opted out.
- (H1)-(H3) were all correct: The rate of opting out varied across treatments in a coherent manner, suggesting that payoffs had a significant, intuitive effect on subjects' willingnesses to rely on the self-interested behavior of others.
- Experience as a B player was associated with significantly greater willingness to rely on the other's maximization in the role of an A player.

			A chose *R*		
Treatment	# of pairs	*A* chose *L*	*B* chose *l*	*B* chose *r*	% secure by A
1	35	23	2	10	65.7%
2	31	20	0	11	64.5%
3	25	5	0	20	20.0%
4	32	15	0	17	46.9%
5	21	18	0	3	85.7%
6	26	8	0	18	30.7%
7*	30	20	0	10	66.7%

Other work tests key theoretical extensive-form notions:

- *Subgame-perfectness/backward induction* (a form of iterated weak dominance in the normal form; generalizations perfect Bayesian or sequential equilibrium).
- *Forward induction* (another form of iterated weak dominance in the normal form).

Experiments by Cooper, DeJong, Forsythe, and Ross (1990, 1994) and others show that forward induction falls short of describing most subjects' thinking.

Camerer, Johnson, Rymon, and Sen 1993 and Johnson, Camerer, Sen, and Rymon 2002 *JET* (discussed in other slides) studied backward induction in three-period alternating-offers bargaining games, finding widespread deviations in subjects' decisions and information searches from subgame-perfect equilibrium.

Binmore, McCarthy, Ponti, L. Samuelson, and Shaked 2002 *JET* dissect backward induction into (i) rationality, (ii) subgame-consistency (play in a subgame is independent of its position in a larger game), and (iii) truncation consistency (replacing a subgame with its equilibrium payoffs does not affect play elsewhere in the game). Via comparisons across alternating-offers bargaining games with carefully related structure, controlling for social preferences, they find widespread violations of subgame and truncation consistency, which are roughly consistent with heterogeneous, finite look-ahead.

Extensive-form games with communication

J. Wang, Spezio, and Camerer (2010 *AER*; “WSC”; see also CCGI, 2013 *JEL*, Section 9.3.2) used eyetracking to study the use of cheap talk to signal private information in discrete Crawford and Sobel (1982 *Ecma*) sender-receiver games.

Here the main puzzle from previous work is that:

- Senders and receivers deviate systematically from equilibrium in the direction of “overcommunication”; i.e. senders are more truthful and receivers more credulous than in equilibrium with no cost of lying; and

- Despite those deviations, Crawford and Sobel’s equilibrium-based comparative statics result, that more communication is possible, the closer are the sender’s and receiver’s preferences, is strongly confirmed.

Extensive-form games with communication

J. Wang, Spezio, and Camerer (2010 *AER*; “WSC”; see also CCGI, 2013 *JEL*, Section 9.3.2) used eyetracking to study the use of cheap talk to signal private information in discrete Crawford and Sobel (1982 *Ecma*) sender-receiver games.

Here the main puzzle from previous work is that:

- Senders and receivers deviate systematically from equilibrium in the direction of “overcommunication”; i.e. senders are more truthful and receivers more credulous than in equilibrium with no cost of lying; and

- Despite those deviations, Crawford and Sobel’s equilibrium-based comparative statics result, that more communication is possible, the closer are the sender’s and receiver’s preferences, is strongly confirmed.

WSC resolve the puzzle by finding that their subjects’ information searches and decisions are close to the predictions of a level-*k* model with *L0* anchored not in uniform randomness but in truthfulness, in the style of Crawford’s (2003 *AER*) level-k model of signaling of intended decisions (CGCI, Section 9.1).

WSC's design implements discretized sender-receiver games that vary the closeness of sender's and receiver's preferences, while eyetracking senders.

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

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

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

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

- The Sender's ideal outcome is A = S + b (or 5, if S + B > 5).

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

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

Crawford and Sobel characterized the possible equilibrium relationships between Sender’s observed S and Receiver’s choice of A.

For models with continuous state and action spaces that generalize Wang et al.’s examples (except for discreteness), they showed that all equilibria are “partition equilibria”, in which the Sender partitions the set of states into contiguous groups and tells the Receiver, in effect, only which group his observation lies in.

Crawford and Sobel characterized the possible equilibrium relationships between Sender's observed S and Receiver's choice of A.

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

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

Crawford and Sobel characterized the possible equilibrium relationships between Sender's observed S and Receiver's choice of A.

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

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

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

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

In Crawford and Sobel's analysis, because talk is cheap, there is nothing to tie down the meanings of messages in equilibrium.

What is tied down are the possible equilibrium relationships between Sender's observed S and Receiver's choice of A.

The meanings of messages are implicit in the beliefs that support the equilibrium.

In Crawford and Sobel's analysis, because talk is cheap, there is nothing to tie down the meanings of messages in equilibrium.

What is tied down are the possible equilibrium relationships between Sender's observed S and Receiver's choice of A.

The meanings of messages are implicit in the beliefs that support the equilibrium.

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

In equilibrium there can be no lying or deception as often occurs in real communication, only intentional vagueness (which also occurs).

In Crawford and Sobel's analysis, because talk is cheap, there is nothing to tie down the meanings of messages in equilibrium.

What is tied down are the possible equilibrium relationships between Sender's observed S and Receiver's choice of A.

The meanings of messages are implicit in the beliefs that support the equilibrium.

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

In equilibrium there can be no lying or deception as often occurs in real communication, only intentional vagueness (which also occurs).

But behaviorally, with a clear correspondence between state and message as I WSC's design, messages are understood literally, even if not always believed.

This motivates WSC's and Crawford's 2003 level-*k* analyses with *L0* anchored in truthfulness rather than uniform randomness.

When $b = 0$ or 1 in WSC's design (sender's and receiver's preferences are close enough) there are multiple equilibria; WSC focus on the most informative one.

When $b = 0$, the most informative equilibrium has $M = S$ and $A = S$: perfect truth-telling, credulity, and communication, as is plausible with identical preferences.

When b = 0 or 1 in WSC's design (sender's and receiver's preferences are close enough) there are multiple equilibria; WSC focus on the most informative one.

When b = 0, the most informative equilibrium has M = S and A = S: perfect truth-telling, credulity, and communication, as is plausible with identical preferences.

When b = 1, the most informative equilibrium has Senders sending M = 1 when S = 1 but M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and Receivers choosing A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}: The difference in preferences then causes imperfect communication even in the most informative equilibrium.

When b = 0 or 1 in WSC's design (sender's and receiver's preferences are close enough) there are multiple equilibria; WSC focus on the most informative one.

When b = 0, the most informative equilibrium has M = S and A = S: perfect truth-telling, credulity, and communication, as is plausible with identical preferences.

When b = 1, the most informative equilibrium has Senders sending M = 1 when S = 1 but M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and Receivers choosing A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}: The difference in preferences then causes imperfect communication even in the most informative equilibrium.

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

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

When b = 0 or 1 in WSC's design (sender's and receiver's preferences are close enough) there are multiple equilibria; WSC focus on the most informative one.

When b = 0, the most informative equilibrium has M = S and A = S: perfect truth-telling, credulity, and communication, as is plausible with identical preferences.

When b = 1, the most informative equilibrium has Senders sending M = 1 when S = 1 but M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and Receivers choosing A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}: The difference in preferences then causes imperfect communication even in the most informative equilibrium.

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

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

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

When b = 0 Wang et al.'s Senders almost always set M = S and their Receivers almost always set A = M: near the perfect communication predicted by the most informative equilibrium with identical preferences.

Figure 1 (next slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 0.

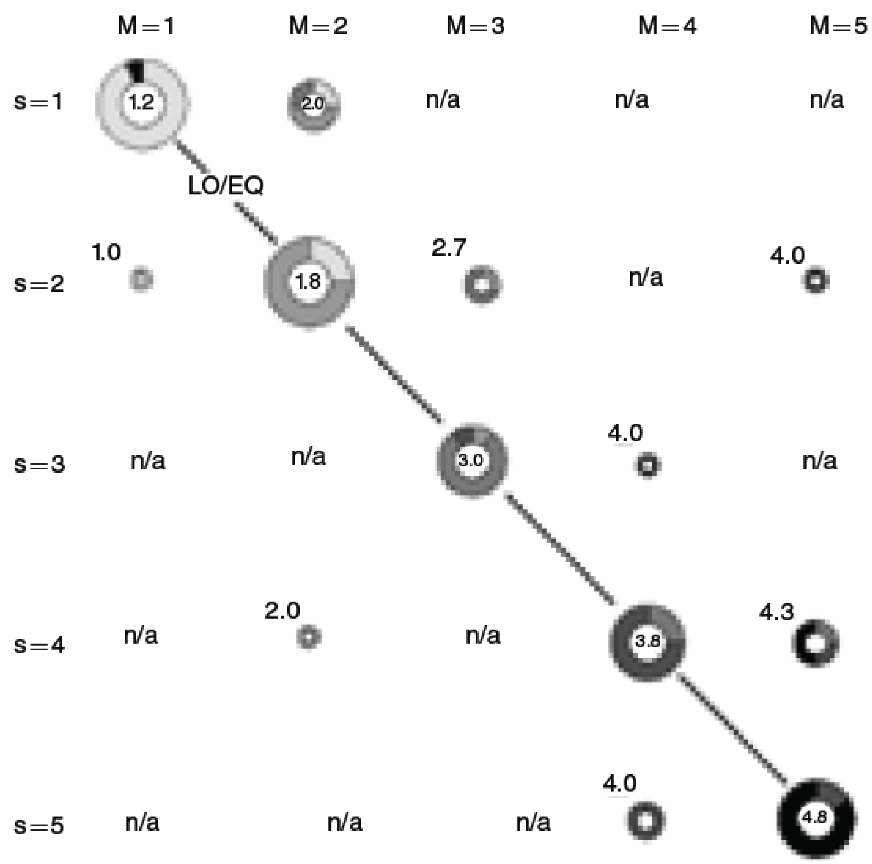

FIGURE 1. RAW DATA PIE CHARTS ($b = 0$)
(HIDDEN BIAS–STRANGER)

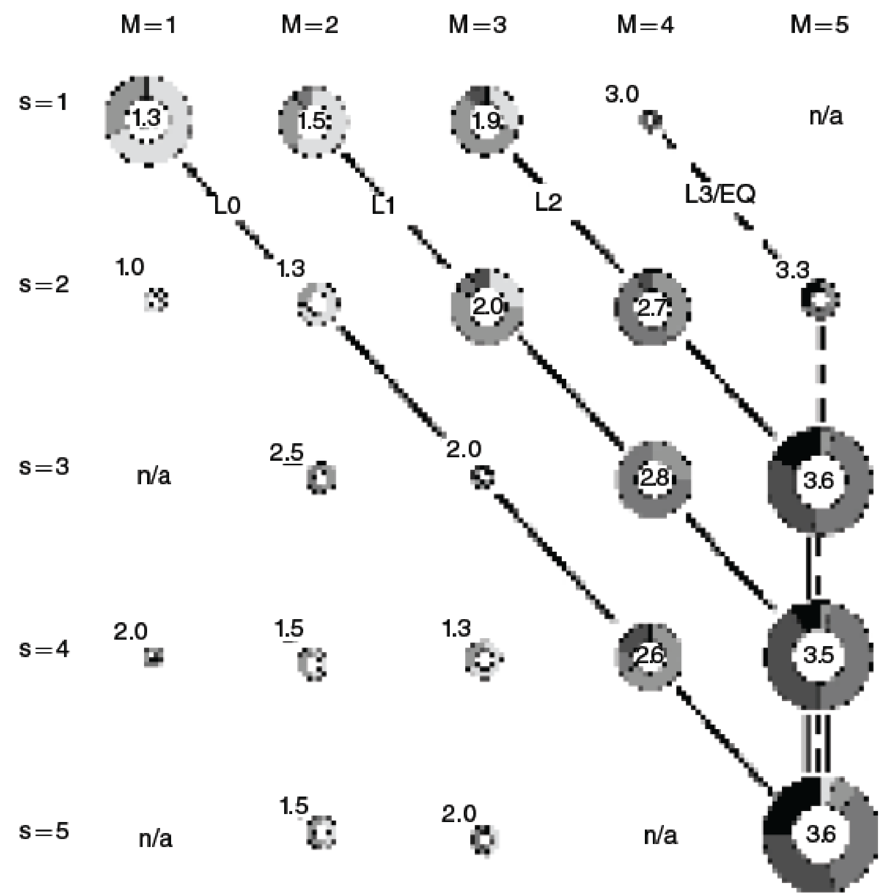

FIGURE 2. RAW DATA PIE CHART ($b = 1$)
(HIDDEN BIAS–STRANGER)

(A circle's size shows the Sender's message frequencies. A circle's darkness and the numbers inside show the Receiver's action frequencies.)

As b increases from 0 to 1 or 2, the amount of information transmitted decreases, as predicted by Crawford and Sobel's equilibrium comparative statics.

Figure 2 (previous slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 1.

As b increases from 0 to 1 or 2, the amount of information transmitted decreases, as predicted by Crawford and Sobel's equilibrium comparative statics.

Figure 2 (previous slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 1.

When b = 1, in the most informative robust equilibrium, Sender's message is M = 1 when S = 1 and M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and Receiver chooses A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}: Thus, in equilibrium the distributions of messages and actions are the same for S = 2, 3, 4, or 5.

As b increases from 0 to 1 or 2, the amount of information transmitted decreases, as predicted by Crawford and Sobel's equilibrium comparative statics.

Figure 2 (previous slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 1.

When b = 1, in the most informative robust equilibrium, Sender's message is M = 1 when S = 1 and M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and Receiver chooses A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}: Thus, in equilibrium the distributions of messages and actions are the same for S = 2, 3, 4, or 5.

However, in Figure 2:

- Senders almost always exaggerate the truth (messages above diagonal), apparently trying to move Receivers from Receivers' ideal action A = S toward Senders' ideal action A = S + 1 (or 5, if S + 1 > 5).

As b increases from 0 to 1 or 2, the amount of information transmitted decreases, as predicted by Crawford and Sobel's equilibrium comparative statics.

Figure 2 (previous slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 1.

When b = 1, in the most informative robust equilibrium, Sender's message is M = 1 when S = 1 and M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and Receiver chooses A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}: Thus, in equilibrium the distributions of messages and actions are the same for S = 2, 3, 4, or 5.

However, in Figure 2:

- Senders almost always exaggerate the truth (messages above diagonal), apparently trying to move Receivers from Receivers' ideal action A = S toward Senders' ideal action A = S + 1 (or 5, if S + 1 > 5).

- Even so, there is some information in Senders' messages (message distributions shift rightward going down, so positively correlated with the state).

As b increases from 0 to 1 or 2, the amount of information transmitted decreases, as predicted by Crawford and Sobel's equilibrium comparative statics.

Figure 2 (previous slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 1.

When b = 1, in the most informative robust equilibrium, Sender's message is M = 1 when S = 1 and M = {2, 3, 4, 5} when S = 2, 3, 4, or 5; and Receiver chooses A = 1 when M = 1 and A = 3 or 4 when M = {2, 3, 4, 5}: Thus, in equilibrium the distributions of messages and actions are the same for S = 2, 3, 4, or 5.

However, in Figure 2:

- Senders almost always exaggerate the truth (messages above diagonal), apparently trying to move Receivers from Receivers' ideal action A = S toward Senders' ideal action A = S + 1 (or 5, if S + 1 > 5).

- Even so, there is some information in Senders' messages (message distributions shift rightward going down, so positively correlated with the state).

- Receivers are deceived to some extent (average A usually > S).

Figure 3 (next slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 2.

In the essentially unique, most informative equilibrium when b = 2, M = {1, 2, 3, 4, 5}, so equilibrium message distributions would look the same for all five rows; and equilibrium actions would be concentrated on A = 3.

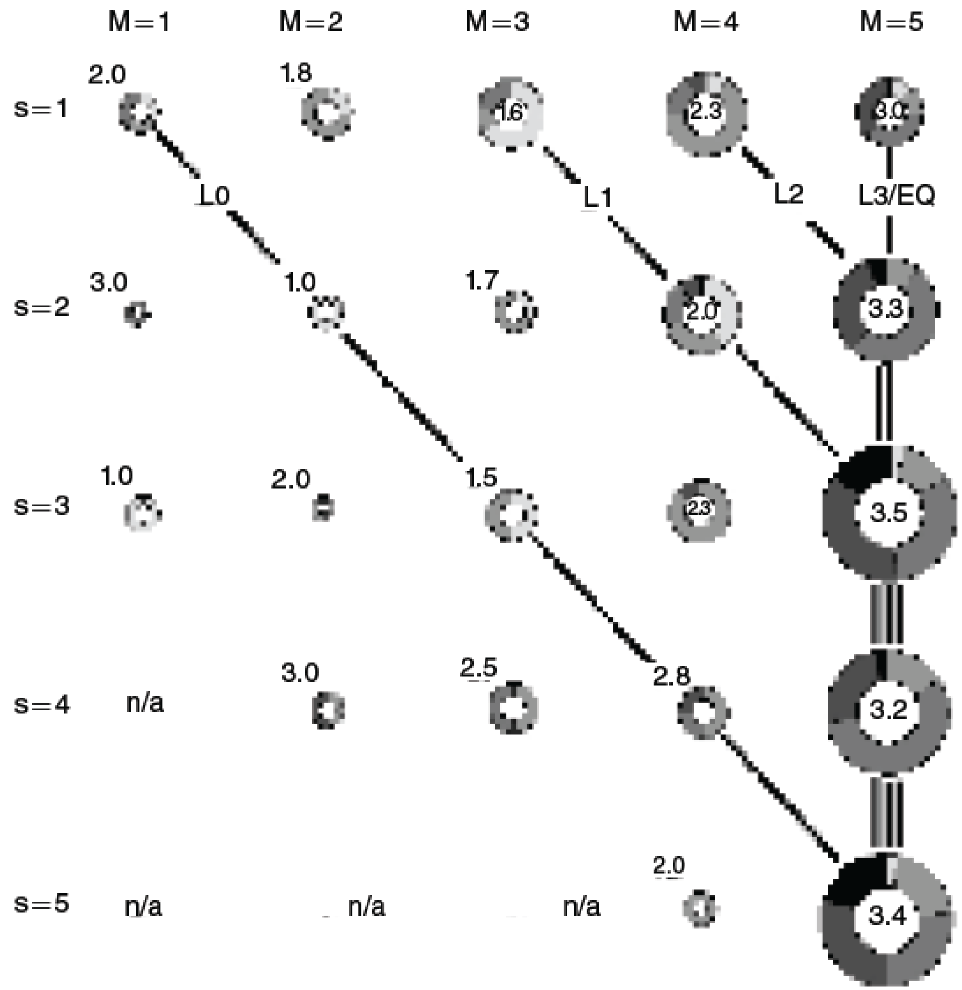

FIGURE 3. RAW DATA PIE CHART ($b = 2$)
(HIDDEN BIAS–STRANGER)

Figure 3 (previous slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 2.

In the essentially unique, most informative equilibrium when b = 2, M = {1, 2, 3, 4, 5}, so equilibrium message distributions would look the same for all five rows; and equilibrium actions would be concentrated on A = 3.

However, in Figure 3:

● Most Senders again exaggerate the truth (messages above the diagonal), apparently trying to move Receivers from Receivers' ideal action A = S toward Senders' ideal action A = S + 2 (or 5, if S + 2 > 5).

Figure 3 (previous slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 2.

In the essentially unique, most informative equilibrium when b = 2, M = {1, 2, 3, 4, 5}, so equilibrium message distributions would look the same for all five rows; and equilibrium actions would be concentrated on A = 3.

However, in Figure 3:

- Most Senders again exaggerate the truth (messages above the diagonal), apparently trying to move Receivers from Receivers' ideal action A = S toward Senders' ideal action A = S + 2 (or 5, if S + 2 > 5).

- Even so, there is some information in Senders' messages (message distributions shift rightward going down, so positively correlated with the state).

Figure 3 (previous slide) shows the Sender's message frequencies and the Receiver's action frequencies as functions of the observed state S when b = 2.

In the essentially unique, most informative equilibrium when b = 2, M = {1, 2, 3, 4, 5}, so equilibrium message distributions would look the same for all five rows; and equilibrium actions would be concentrated on A = 3.

However, in Figure 3:

- Most Senders again exaggerate the truth (messages above the diagonal), apparently trying to move Receivers from Receivers' ideal action A = S toward Senders' ideal action A = S + 2 (or 5, if S + 2 > 5).

- Even so, there is some information in Senders' messages (message distributions shift rightward going down, so positively correlated with the state).

- Receivers are usually deceived to some extent (average A usually > S).

What kind of model can explain results like this?

WSC, following Cai and Wang (2006 *GEB*), propose a level-*k* explanation in the style of Crawford's (2003 *AER*) analysis of preplay communication of intentions:

Anchor beliefs in a truthful Sender *L0*, which sets M = S; and a credulous Receiver *L0* (which also best responds to an *L0* Sender), setting A = M.

L1 Senders best respond to *L0* Receivers, inflating messages by b: M = S + b (up to M = 5), so *L0* Receivers choose S + b, yielding Sender's ideal action given S.

L1 Receivers (in WSC's numbering convention) best respond to *L1* Senders, discounting messages by b, normally setting A = M – b, yielding Receivers' ideal action given M = S + b of S.

("Normally" reflects WSC's assumption that *L1* Receivers take into account that when b = 2, *L1* senders with S = 3, 4, or 5 all send M = 5, so *L1* Receivers, knowing S is equally likely to be 3, 4, or 5, choose A = 4 not A = M – 2b = 3.)

TABLE 1—BEHAVIORAL PREDICTIONS OF THE LEVEL-k MODEL

Sender message (condition on state)						Receiver action (condition on message)					
State	1	2	3	4	5	Message	1	2	3	4	5
$b = 0$											
$L0/EQ$ sender	1	2	3	4	5	$L0/EQ$ receiver	1	2	3	4	5
$b = 1$											
$L0$ sender	1	2	3	4	5	$L0$ receiver	1	2	3	4	5
$L1$ sender	2	3	4	5	5	$L1$ receiver	1	1	2	3	4
$L2$ sender	3	4	5	5	5	$L2$ receiver	1	1	1	2	4
EQ sender	4	5	5	5	5	EQ receiver	1	1	1	1	4
$SOPH$ sender	3	4	5	5	5	$SOPH$ receiver	1	2	2	2	4
$b = 2$											
$L0$ sender	1	2	3	4	5	$L0$ receiver	1	2	3	4	5
$L1$ sender	3	4	5	5	5	$L1$ receiver	1	1	1	2	4
$L2$ sender	4	5	5	5	5	$L2$ receiver	1	1	1	1	4
EQ sender	5	5	5	5	5	EQ receiver	1	1	1	1	3
$SOPH$ sender	5	5	5	5	5	$SOPH$ receiver	2	2	2	2	3

Notes: $L0$ senders are truthful and $L0$ receivers best respond to $L0$ senders by following the message. $L1$ senders best respond to $L0$ receivers, while $L1$ receivers best respond to $L1$ senders, and so on. Note that when $b = 2$, due to discreteness both $L2$ and $EQ\,(= L3)$ senders best respond to $L1$ receivers.

L2 Senders best respond to *L1* Receivers, inflating messages by 2b: M = S + 2b (up to M = 5), so *L1* Receivers set A = M – b = S + b, yielding Senders' ideal action given S.

L2 Receivers best respond to *L2* Senders, discounting messages by 2b, normally setting A = M – 2b, yielding Receivers' ideal action given M = S + 2b of S.

("Normally" reflects WSC's assumption that *L2* Receivers take into account that when b = 1, *L2* senders with S = 3, 4, or 5 all send M = 5, so *L2* Receivers, knowing S is equally likely to be 3, 4, or 5, choose A = 4 not A = M – 2b = 3.

L2 Receivers also take into account that when b = 2, *L2* senders with S = 2, 3, 4, or 5 send M = 5, so *L2* Receivers, knowing S is equally likely to be 2, 3, 4, or 5, choose A = 4 not A = M – 2b = 3.)

Econometric estimation classifies 18% of 16 Sender subjects as *L0*, 25% *L1*, 25% *L2,* 14% *Sophisticated*, and 18% *Equilibrium* (note different type definitions).

Figures 2 and 3 show why. (When b = 1, *L1*, *L2*, and *Eq* all predict M = 5 when S = 4 or 5; and when b = 2, *L1*, *L2*, and *Eq* all predict M = 5 when S = 3, 4, or 5.)

Notes: The true states are in rows, and senders' messages are in columns. Each cell contains the average action taken by the receivers and a pie chart breakdown of the actions. Actions are presented in a gray scale, ranging from white (action 1) to black (action 5). The size of the pie chart is proportional to the number of occurrences for the corresponding state and message.

TABLE 4— LEVEL-*k* CLASSIFICATION RESULTS

Session	ID	log L	*k*	Exact	lambda	Treatment
1	1	−46.23	*SOPH*	0.64	0.06	eyetracked subject #1
1	2	−25.99	*L1*	0.87	0.00	eyetracked subject #2
1	3	−15.98	*L2*	0.91	0.44	open box
2	1	−37.32	*L1*	0.60	0.52	eyetracked subject #3
2	2	−37.34	*EQ*	0.73	0.52	open box (eyetracked to round 20)
2	3	−25.70	*SOPH*	0.83	0.07	open box
3	1	−68.84	n/a	0.13	0.01	eyetracked subject #4
3	2	−17.71	*SOPH*	0.89	0.12	eyetracked subject #5
3	3	−54.73	*EQ*	0.60	0.03	open box
4	1	−50.86	*L1*	0.51	0.04	eyetracked subject #6
4	3	−25.22	*EQ*	0.82	0.48	open box
5	1	−22.26	*L1*	0.89	0.02	eyetracked subject #7
5	2	−35.77	*L2*	0.78	0.03	eyetracked subject #8
5	3	−25.17	*EQ*	0.87	0.04	open box
6	1	−16.27	*L2*	0.91	0.43	eyetracked subject #9
6	2	−42.02	*SOPH*	0.62	0.13	eyetracked subject #10
6	3	−52.17	*L0*	0.62	0.01	open box