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

Measuring cognition in economic decisions: How and why? II: Studying cognition via information search in game experiments

Vincent P. Crawford

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

(with large debts to Colin Camerer and Miguel Costa-Gomes)

Revised 7 March 2015

Introduction: Why use process data to study strategic thinking?

In principle, experimental design can separate the decisions implied by different models well enough to infer strategic thinking entirely from decisions.

But in economically interesting games, our ability to distinguish among models is often near the limit of experimental feasibility, and existing methods are fairly easily adapted to gather process data along with decision data.

Introduction: Why use process data to study strategic thinking?

In principle, experimental design can separate the decisions implied by different models well enough to infer strategic thinking entirely from decisions.

But in economically interesting games, our ability to distinguish among models is often near the limit of experimental feasibility, and existing methods are fairly easily adapted to gather process data along with decision data.

Monitoring subjects' searches for hidden but freely accessible payoff information provides an independent view of cognition and behavior, at low additional cost.

Independently separating decision rules' implications for search and decisions multiplies the power of the design to identify subjects' decision rules.

Introduction: Why use process data to study strategic thinking?

In principle, experimental design can separate the decisions implied by different models well enough to infer strategic thinking entirely from decisions.

But in economically interesting games, our ability to distinguish among models is often near the limit of experimental feasibility, and existing methods are fairly easily adapted to gather process data along with decision data.

Monitoring subjects' searches for hidden but freely accessible payoff information provides an independent view of cognition and behavior, at low additional cost.

Independently separating decision rules' implications for search and decisions multiplies the power of the design to identify subjects' decision rules.

Analyzing search data suggests an algorithmic view of how subjects process information into decisions, which yields more direct insight into cognition.

With careful design, search data can sometimes directly reveal the algorithms subjects use to make their decisions.

In these slides I describe three examples of this approach, highlighting modes of analysis and how search data changes our view of cognition.

Background reading

Johnson, Camerer, Sen, and Rymon, “Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining,” *Journal of Economic Theory* 2002 (“JC”; see also their cited 1993 paper)

Camerer and Johnson, “Thinking about Attention in Games: Backward and Forward Induction,” in Brocas and Carrillo (eds.), *The Psychology of Economic Decisions, Volume Two: Reasons and Choices*, Oxford University Press, 2004

Costa-Gomes, Crawford, and Broseta, “Cognition and Behavior in Normal-Form Games: An Experimental Study,” *Econometrica* 2001 (“CGCB”)

Costa-Gomes and Crawford, “Cognition and Behavior in Two-Person Guessing Games: An Experimental Study,” *American Economic Review* 2006 (“CGC”).

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

Wang, Spezio, and Camerer, “Pinocchio’s Pupil: Using Eyetracking and Pupil Dilation To Understand Truth-telling and Deception in Sender-Receiver Games” *American Economic Review* 2010 (“WSC”).

Brocas, Carrillo, Wang, and Camerer, “Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games,” *Review of Economic Studies* 2014 (“BCWC”).

JC’s Extensive-Form Alternating-Offers Bargaining Games

JC’s subjects played series of finite-horizon alternating-offers bargaining games, framed in extensive form.

When JC’s experiments were originally designed (early 1990s), it was known that subjects systematically deviated from the *purely pecuniarily self-interested* subgame-perfect equilibrium strategies in these games:

- Proposers often made offers more generous than such an equilibrium predicts.
- Responders often rejected small-to-moderate offers that would have given them positive pecuniary gains.

JC's Extensive-Form Alternating-Offers Bargaining Games

JC's subjects played series of finite-horizon alternating-offers bargaining games, framed in extensive form.

When JC's experiments were originally designed (early 1990s), it was known that subjects systematically deviated from the *purely pecuniarily self-interested* subgame-perfect equilibrium strategies in these games:

- Proposers often made offers more generous than such an equilibrium predicts.
- Responders often rejected small-to-moderate offers that would have given them positive pecuniary gains.

But there was controversy about whether those deviations were due to:

- Nonpecuniary motives, with responders' taking revenge for unfair offers by rejecting them (and proposers' beliefs that they would do so), or
- Cognitive limitations preventing subjects from identifying their subgame-perfect equilibrium offer or acceptance strategies.

JC addressed this issue by presenting a series of different alternating-offers bargaining games to subjects, within a publicly announced structure.

Each game was presented as a sequence of "pies" via MouseLab (DOS-based; web version at http://www.mouselabweb.org/), which normally concealed the pies but allowed subjects to look them up as often as desired, one at a time.

Subjects were not allowed to write down the pie sizes, variations across games made it impossible to remember them from previous plays, and their look-up frequencies made clear that they did not just scan and memorize them.

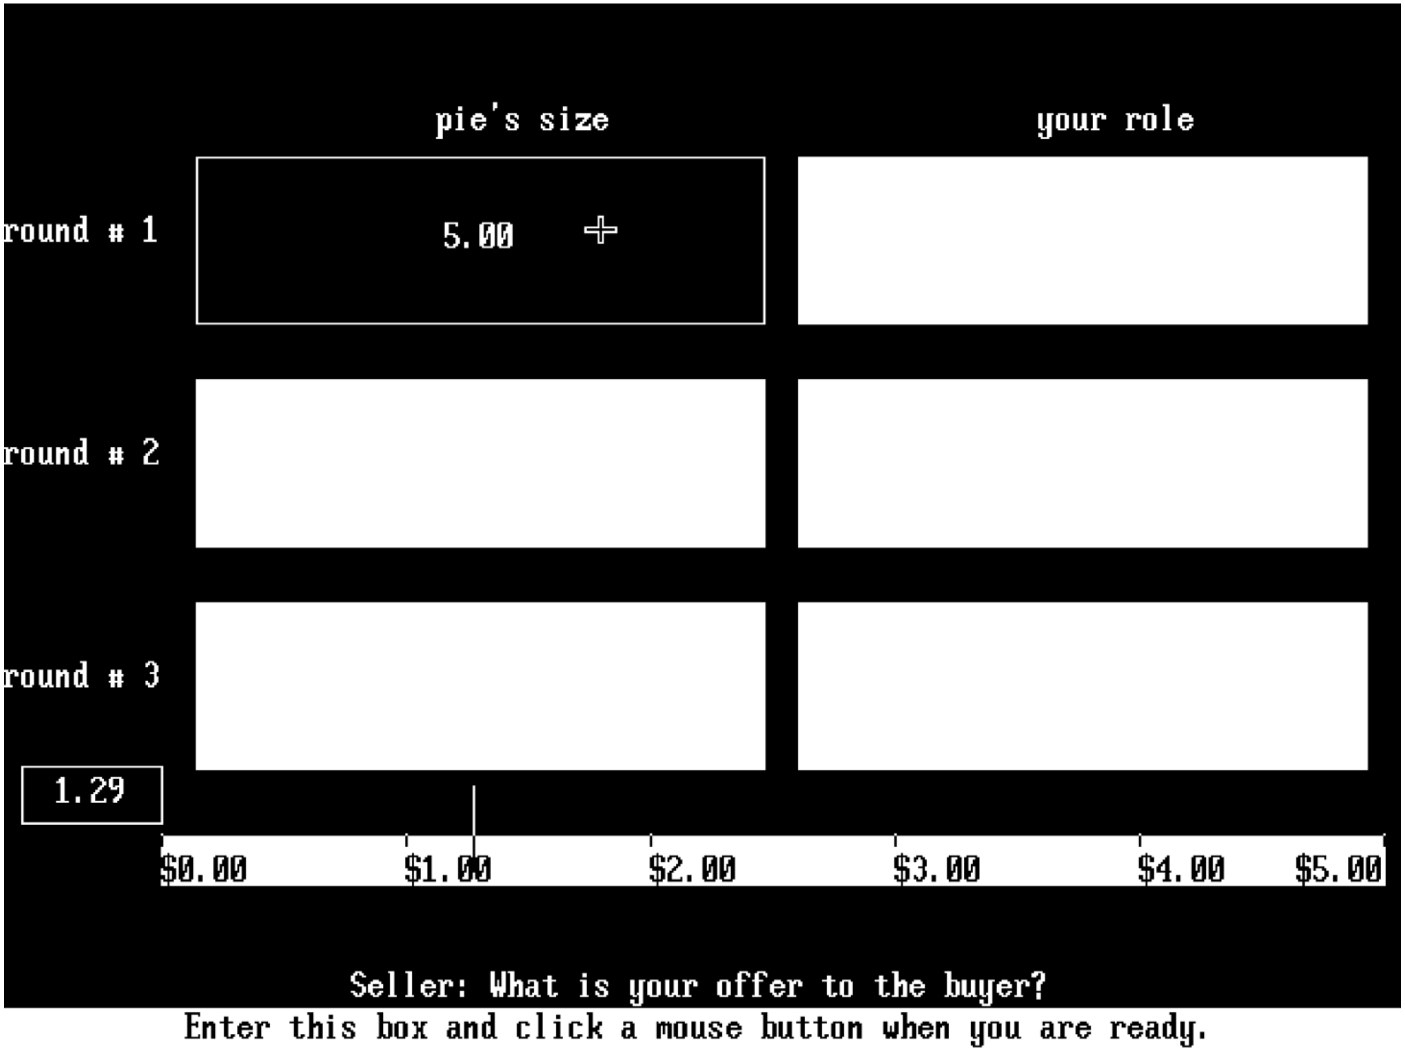

CJ's Figure 1. MouseLab Screen Display

JC argued that (even with privately observed revenge motives) backward induction is the easiest way to find a subgame-perfect (or sequential) equilibrium.

They also argued that backward induction is naturally associated with search patterns in which subjects first look at the last-period pie, then between the last and second-last pie, then between the second-last and first-period pie.

JC argued that (even with privately observed revenge motives) backward induction is the easiest way to find a subgame-perfect (or sequential) equilibrium.

They also argued that backward induction is naturally associated with search patterns in which subjects first look at the last-period pie, then between the last and second-last pie, then between the second-last and first-period pie.

Noting that, in theory, subjects could scan and memorize the pie sizes before making decisions, in which case search might be unrelated to cognition, JC ran a "robot/trained subjects" control with subjects playing against a computer, told it would play the subgame-perfect equilibrium, and trained to identify its strategies.

JC's robot/trained subjects made decisions close to the subgame-perfect equilibrium, while coming close to its characteristic search pattern.

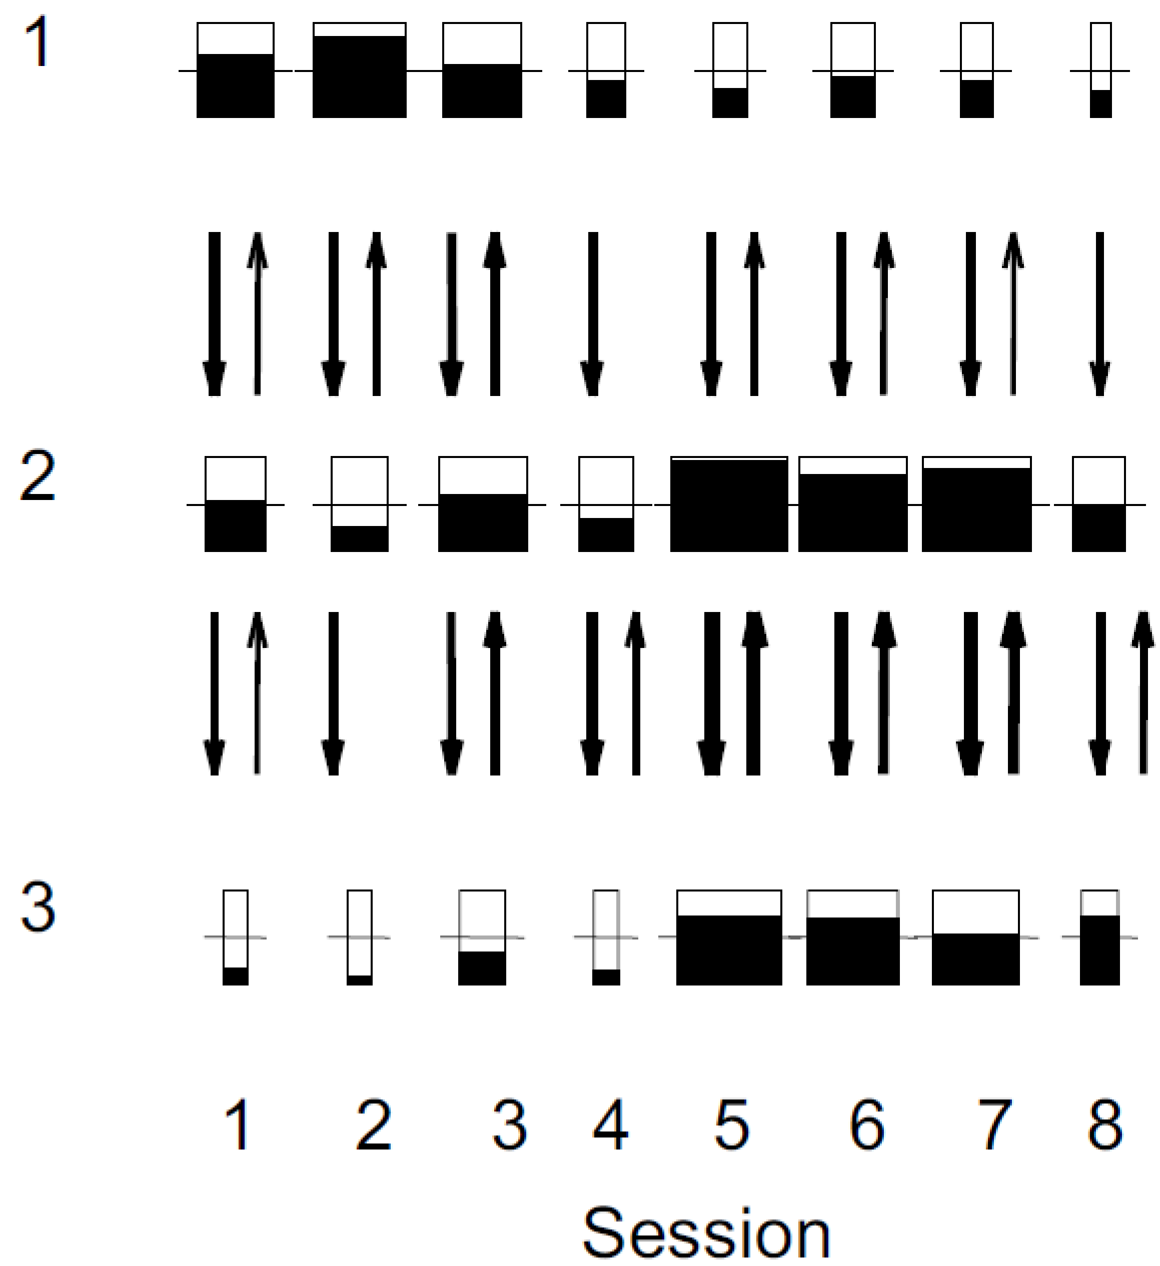

FIG. 6. Changes in information processing before training (Sessions 1–4) and after (Sessions 5–8).

JC argued that (even with privately observed revenge motives) backward induction is the easiest way to find a subgame-perfect (or sequential) equilibrium.

They also argued that backward induction is naturally associated with search patterns in which subjects first look at the last-period pie, then between the last and second-last pie, then between the second-last and first-period pie.

Noting that, in theory, subjects could scan and memorize the pie sizes before making decisions, in which case search might be unrelated to cognition, JC ran a "robot/trained subjects" control with subjects playing against a computer, told it would play the subgame-perfect equilibrium, and trained to identify its strategies.

JC's robot/trained subjects made decisions close to the subgame-perfect equilibrium, while coming close to its characteristic search pattern.

This and JC's other results suggest that there are strong regularities in search behavior, and that subjects' searches contain a lot of information about cognition.

By contrast, JC's Baseline subjects, playing against other baseline subjects, without training, deviated substantially from both subgame-perfect equilibrium decisions and backward-induction search, in positively correlated ways:

- About 10% never even looked at the last-period pie.

- Many deviated from the characteristic backward-induction search pattern.

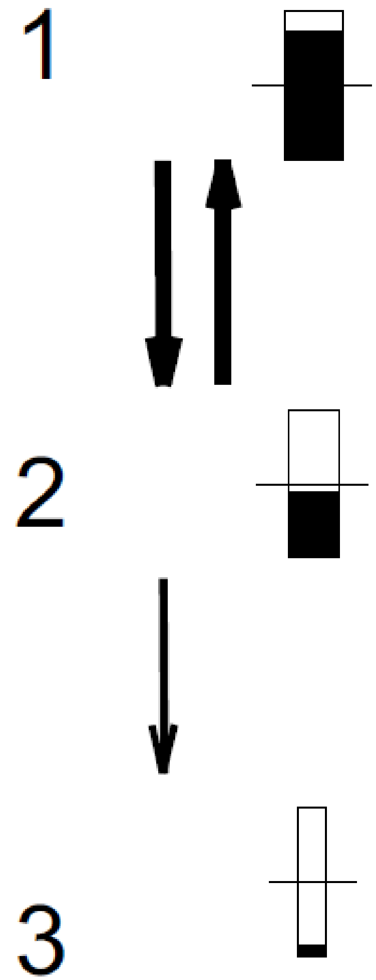

Excerpt from JC’s Figure 3: Baseine subjects.

By contrast, JC's Baseline subjects, playing against other baseline subjects, without training, deviated substantially from both subgame-perfect equilibrium decisions and backward-induction search, in positively correlated ways:

- About 10% never even looked at the last-period pie.

- Many deviated from the characteristic backward-induction search pattern.

- Subjects whose searches were further from backward induction made decisions further from their subgame-perfect equilibrium decisions.

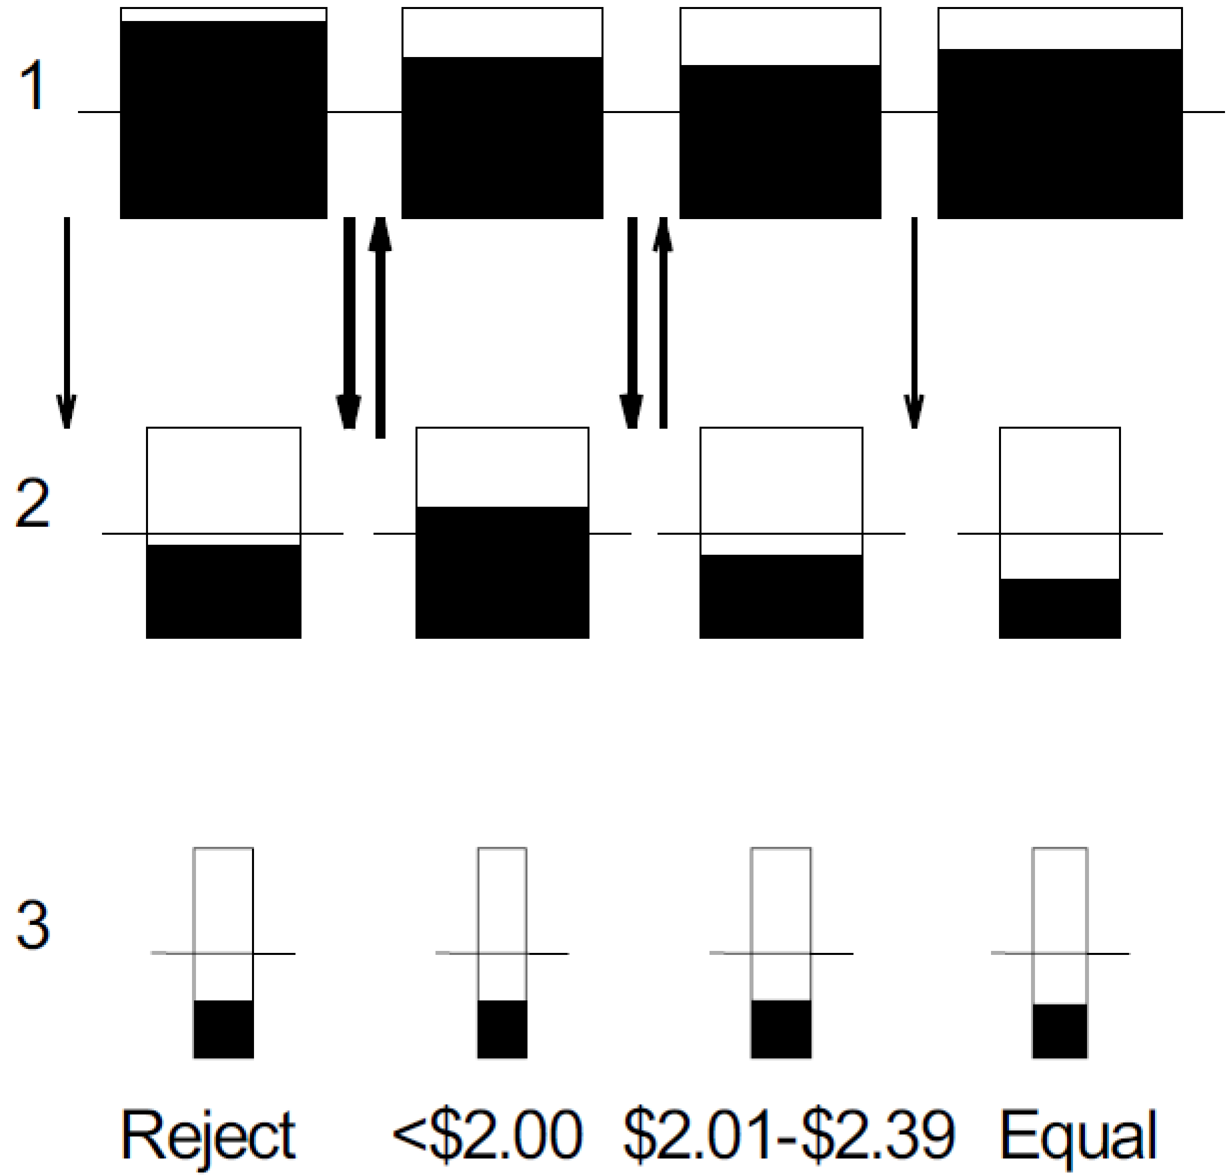

FIG. 5. Icon graphs for player 2 by rejection or size of offer accepted.

JC also found evidence of a mixture of “levels” in the subject population:

- *Level-0* subjects treat the first round as an ultimatum game
- *Level-1* subjects look one round ahead but truncate beyond that, and
- *Level-2* subjects look two rounds ahead as subgame-perfect equilibrium requires, hence are functionally equivalent here to *Equilibrium* subjects.

JC also found evidence of a mixture of “levels” in the subject population:

- *Level-0* subjects treat the first round as an ultimatum game
- *Level-1* subjects look one round ahead but truncate beyond that, and
- *Level-2* subjects look two rounds ahead as subgame-perfect equilibrium requires, hence are functionally equivalent here to *Equilibrium* subjects.

Level-2, *Level-1*, and *Level-0* subjects deviate progressively more and more from subgame-perfect equilibrium in search as well as decisions, so that a mixture of levels yields a positive correlation between search and decision deviations.

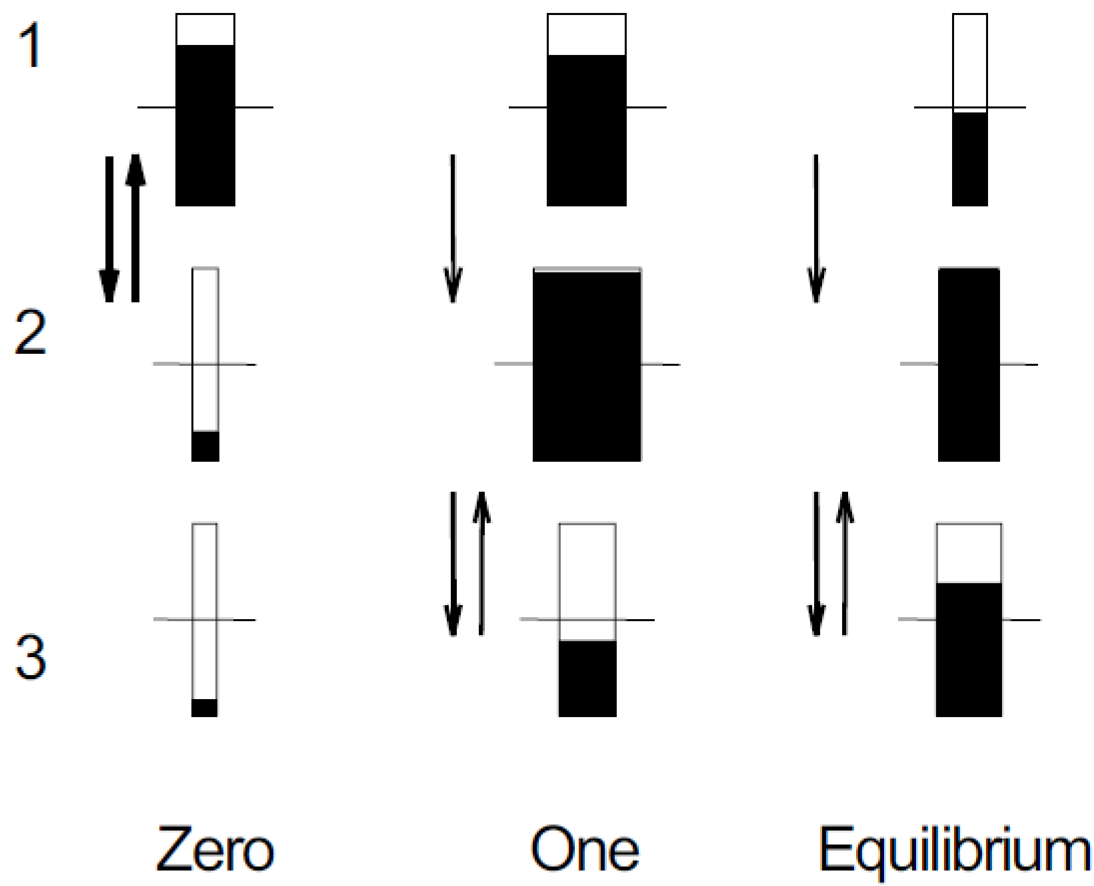

FIG. 7. Icon graphs of information processing (time, number of acquisitions and transitions) by type.

JC also found evidence of a mixture of “levels” in the subject population:

- *Level-0* subjects treat the first round as an ultimatum game
- *Level-1* subjects look one round ahead but truncate beyond that, and
- *Level-2* subjects look two rounds ahead as subgame-perfect equilibrium requires, hence are functionally equivalent here to *Equilibrium* subjects.

Level-2, *Level-1*, and *Level-0* subjects deviate progressively more and more from subgame-perfect equilibrium in search as well as decisions, so that a mixture of levels yields a positive correlation between search and decision deviations.

JC’s analysis suggests that the deviations from subgame-perfect equilibrium are due roughly half to revenge motives and half to cognitive limitations.

CGCB’s normal-form matrix games

CGCB’s subjects played a series of 18 matrix games, with various patterns of iterated dominance or unique pure-strategy equilibrium without dominance.

Within a publicly announced structure, each game was presented via MouseLab, as a matrix with payoffs separated horizontally and all subjects as Row players.

Subjects were not allowed to write down the payoffs, and their look-up frequencies made clear that they did not memorize them.

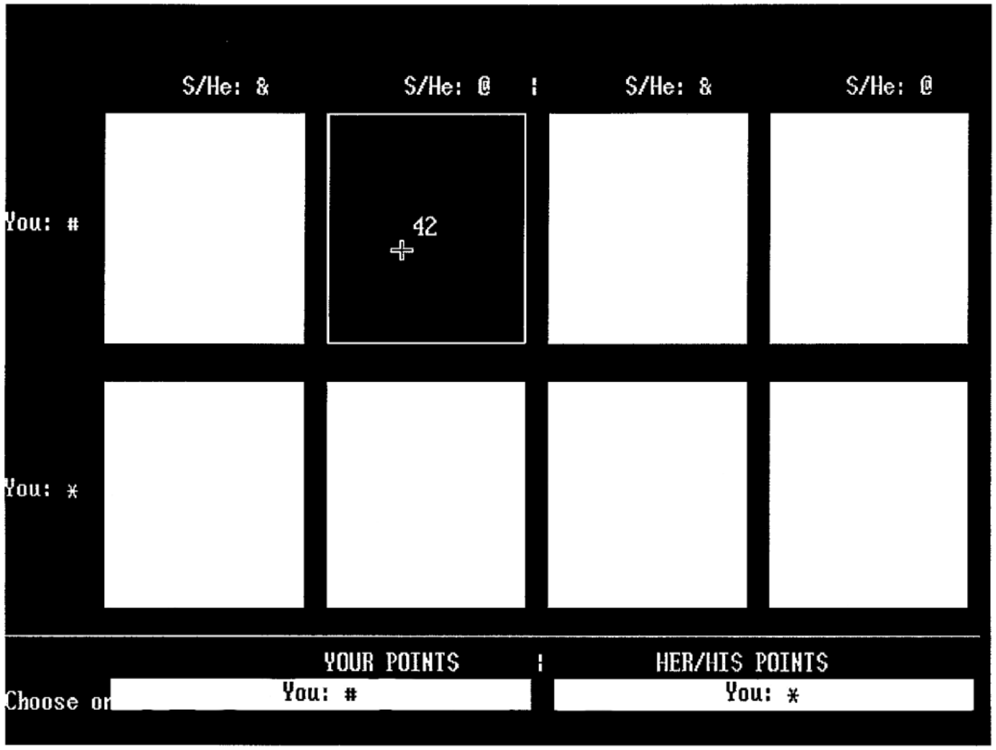

CGCB's Figure 1. MouseLab Screen Display (for a 2×2 game)

CGC's normal-form two-person guessing games

CGC's subjects played a series of 16 two-person guessing games.

Each player in a given game had his own limits and target, and the targets and limits varied independently across players and games.

Players were not required to guess between their limits:

Guesses outside their limits were automatically adjusted up to the lower or down to the upper limit as needed (enhances separation of rules' search implications).

A player's payoff increased with the closeness of his adjusted guess to his target times the other's adjusted guess.

Within a publicly announced structure, each game was presented via MouseLab.

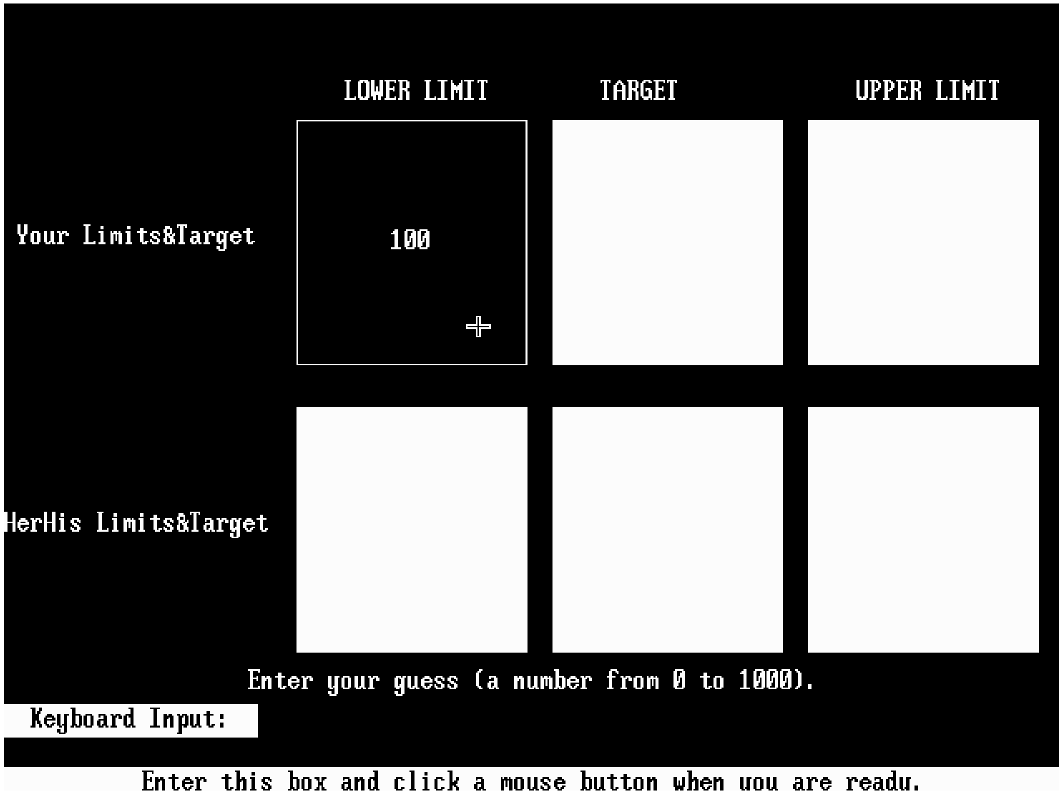

CGC's Figure 6. Screen Shot of the MouseLab Display

Within a publicly announced structure, each game was presented via MouseLab.

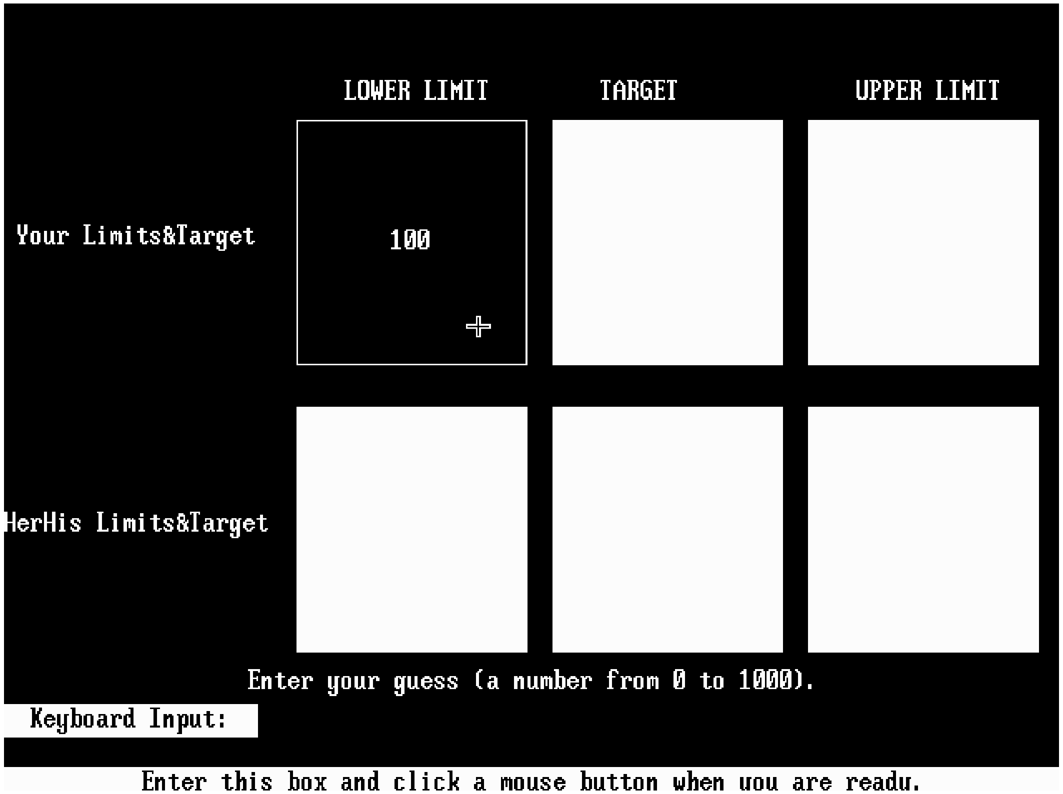

CGC's Figure 6. Screen Shot of the MouseLab Display

Again subjects were not allowed to write down the payoffs, and the frequencies of repeated look-ups made clear that they did not memorize them.

Like JC's and CGCB's designs, CGC's maintains control over subjects' motives for search by making information from previous plays irrelevant to current plays.

CGC’s design combines the strengths of JC’s and CGCB’s designs for studying cognition via search.

In each case low search costs and free access to the payoff parameters made the games’ structures effectively public knowledge (except for responders’ possible revenge motives in JC’s design), so the results can be used to test theories of behavior in complete-information versions of the games.

In each case the design independently separates rules’ implications for search and decisions.

CGC’s design maintains the simplicity of JC’s and CGCB’s, and its simple parametric structure makes rules’ search implications independent of the game.

By contrast with JC's one-dimensional search, CGCB's and CGC's designs make search multidimensional, which makes it potentially more informative.

The analysis of search, however, involves some choices.

CGCB's and CGC's search analyses were organized around theories of cognition that more readily suggest roles for which look-ups subjects make, in which orders, than for numbers of look-ups, transition frequencies, or durations.

(No claim that durations are irrelevant was intended, just that they don't deserve the priority they have been given. CGCB (Table IV) do present some results on durations, under the heading of "gaze times.")

JC also studied look-up orders (transitions between pies) but also gave weight to look-up durations and the numbers of look-ups of each pie ("acquisitions").

Others, like Rubinstein (2007 *Economic Journal*), considered only durations.

Gabaix, Laibson, Moloche, and Weinstein (2006 *American Economic Review*) focused on numbers of look-ups and some measures of look-up order.

By contrast with JC's one-dimensional search, CGCB's and CGC's designs make search multidimensional, which makes it potentially more informative.

The analysis of search, however, involves some choices.

CGCB's and CGC's search analyses were organized around theories of cognition that more readily suggest roles for which look-ups subjects make, in which orders, than for numbers of look-ups, transition frequencies, or durations.

(No claim that durations are irrelevant was intended, just that they don't deserve the priority they have been given. CGCB (Table IV) do present some results on durations, under the heading of "gaze times.")

JC also studied look-up orders (transitions between pies) but also gave weight to look-up durations and the numbers of look-ups of each pie ("acquisitions").

Others, like Rubinstein (2007 *Economic Journal*), considered only durations.

Gabaix, Laibson, Moloche, and Weinstein (2006 *American Economic Review*) focused on numbers of look-ups and some measures of look-up order.

CGCB and CGC also argued that cognition is sufficiently idiosyncratic and searches are sufficiently noisy that they are best studied at the individual level.

Rubinstein, Gabaix et al. and WSC studied search at high levels of aggregation. BCWC took an intermediate view, looking for clusters of subjects.

JC's, CGCB's, and CGC's analyses take a procedural view of decision-making, in which a subject follows one of a set of decision rules in all games.

His rule determines his search, and his rule and search determine his decision.

(Because a rule's search implications depend not only on what decisions it specifies, but why, something like a rule-based model seems necessary here.)

JC's, CGCB's, and CGC's analyses take a procedural view of decision-making, in which a subject follows one of a set of decision rules in all games.

His rule determines his search, and his rule and search determine his decision.

(Because a rule's search implications depend not only on what decisions it specifies, but why, something like a rule-based model seems necessary here.)

The possible rules and their search implications provide bases for the enormous spaces of possible decision and search sequences.

This structure makes it possible to identify links between a subject's cognition, search, and decisions; and makes it meaningful to ask whether a subject's searches deviated from equilibrium in the "same direction" as his decisions.

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

- *L1* best responds to a uniform random *L0*, *L2* best responds to *L1*, and so on.

- *D1* (*D2*) does one round (two rounds) of deletion of dominated decisions and then best responds to a uniform prior over the other's remaining decisions.

- *Equilibrium* makes its equilibrium decision.

- *Sophisticated* best responds to the probabilities of other's decisions (proxied in the data analysis by subjects' observed frequencies).

CGC’s rules all build in risk-neutrality and rule out social preferences:

● *L1* best responds to a uniform random *L0*, *L2* best responds to *L1*, and so on.

● *D1* (*D2*) does one round (two rounds) of deletion of dominated decisions and then best responds to a uniform prior over the other's remaining decisions.

● *Equilibrium* makes its equilibrium decision.

● *Sophisticated* best responds to the probabilities of other’s decisions (proxied in the data analysis by subjects’ observed frequencies).

Each of these decision rules is naturally associated with algorithms that process payoff information into decisions.

The analyses use those algorithms as models of cognition, deriving a rule’s search implications under simple assumptions about how it determines search.

In theory a subject can search in any order, memorize the information, and then make his decisions—in which case search will reveal nothing about cognition.

But there are strong empirical regularities in search behavior.

The goal is to stylize these regularities via enough assumptions to extract the signal from the noise in searches; but not so many that they distort its meaning.

In theory a subject can search in any order, memorize the information, and then make his decisions—in which case search will reveal nothing about cognition.

But there are strong empirical regularities in search behavior.

The goal is to stylize these regularities via enough assumptions to extract the signal from the noise in searches; but not so many that they distort its meaning.

JC (implicitly) and CGCB impose two such assumptions:

● Occurrence: If your rule's decision depends on a particular piece of hidden information, then you must have looked at it at least once; and

● Adjacency: The two pieces of hidden information associated with the most basic operations your rule's decision depends on must be adjacent in your look-up sequence.

In theory a subject can search in any order, memorize the information, and then make his decisions—in which case search will reveal nothing about cognition.

But there are strong empirical regularities in search behavior.

The goal is to stylize these regularities via enough assumptions to extract the signal from the noise in searches; but not so many that they distort its meaning.

JC (implicitly) and CGCB impose two such assumptions:

● Occurrence: If your rule's decision depends on a particular piece of hidden information, then you must have looked at it at least once; and

● Adjacency: The two pieces of hidden information associated with the most basic operations your rule's decision depends on must be adjacent in your look-up sequence.

CGC derive each rule's characteristic look-up sequence in a way that subsumes Occurrence and Adjacency, and use its density in a subject's actual sequence.

For CGCB's subjects (framed as Rows), assuming Occurrence and Adjacency:

- Up-down transitions in own payoffs are associated with decision-theoretic rationality
- Left-right transitions in other's payoffs are associated with thinking about the other subject's incentives
- Transitions from own to other's payoffs and back for the same decision combination are associated with interpersonal (fairness) comparisons

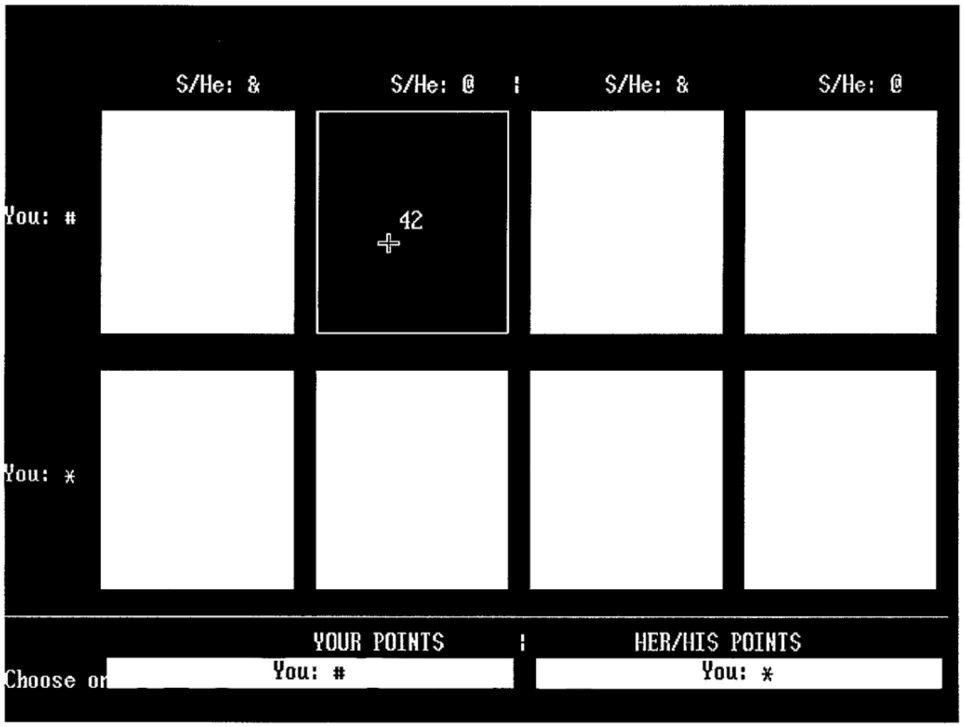

CGCB's Figure 1. MouseLab Screen Display (for a 2×2 game)

In CGCB's data, the most frequent rules estimated from decisions alone are *L1* (CGCB's *Naïve*, which is not separated from *Optimistic* (maximax) by decisions in their design) and *L2*, each nearly half of the population.

Incorporating search compliance into the econometric estimates (using an error-rate model not explained here) shifts the estimated rule distribution toward *L1*, at the expense of *Optimistic* (maximax) and *D1*.

In CGCB’s data, the most frequent rules estimated from decisions alone are *L1* (CGCB’s *Naïve*, which is not separated from *Optimistic* (maximax) by decisions in their design) and *L2*, each nearly half of the population.

Incorporating search compliance into the econometric estimates (using an error-rate model not explained here) shifts the estimated rule distribution toward *L1*, at the expense of *Optimistic* (maximax) and *D1*.

The shift occurs because *L1*’s search implications explain more of the variation in subjects’ searches and decisions than *Optimistic*'s, which are too unrestrictive to be useful in the sample; and because *L1*’s search implications explain more of the variation in subjects’ searches and decisions than *D1*’s, which are more restrictive than *Optimistic*'s, but only weakly correlated with subjects’ decisions.

D1 also loses some frequency to *L2*, even though their decisions are weakly separated in CGCB’s design, because *L2*'s search implications explain much more of the variation in subjects' searches and decisions.

In CGCB’s data, the most frequent rules estimated from decisions alone are *L1* (CGCB’s *Naïve*, which is not separated from *Optimistic* (maximax) by decisions in their design) and *L2*, each nearly half of the population.

Incorporating search compliance into the econometric estimates (using an error-rate model not explained here) shifts the estimated rule distribution toward *L1*, at the expense of *Optimistic* (maximax) and *D1*.

The shift occurs because *L1*’s search implications explain more of the variation in subjects’ searches and decisions than *Optimistic*'s, which are too unrestrictive to be useful in the sample; and because *L1*’s search implications explain more of the variation in subjects’ searches and decisions than *D1*’s, which are more restrictive than *Optimistic*'s, but only weakly correlated with subjects’ decisions.

D1 also loses some frequency to *L2*, even though their decisions are weakly separated in CGCB’s design, because *L2*'s search implications explain much more of the variation in subjects' searches and decisions.

Overall, CGCB’s analysis of decisions and search yields a significantly different interpretation of behavior than their analysis of decisions alone.

Including search suggests a very simple view of behavior, with *L1* and *L2* making up 65-90% of the population, and *D1* 0% (or 20% if one doubts CGCB's model of search). (CGC’s subsequent work suggests that 0% is closer to correct.)

CGC measure a subjects' search compliance with a decision rule as the density of the rule's characteristic look-up sequence in the subject's observed sequence.

CGC then incorporate search compliance into the econometric estimates using an error-rate model similar to CGCB's, not explained here.

A rule's characteristic look-up sequence in one of CGC's games is based on its minimal search implications, as derived from the rule's *ideal guesses*, those the rule would imply if the game did not limit the player's guesses.

(Recall that guesses outside limits were automatically adjusted up to the lower or down to the upper limit as needed, to enhance separation of rules' search implications. With automatic adjustment and CGC's quasiconcave payoffs, a rule's ideal guesses are all a subject needs to know to implement the rule.)

Evaluating a formula for a rule's ideal guess requires a series of operations, some of which are *basic* in that they logically precede any other operation.

Like JC and CGCB, CGC derived rules' search implications assuming that subjects perform basic operations one at a time via adjacent look-ups, remember their results, and otherwise rely on repeated look-ups rather than memory.

Basic operations are then represented in a look-up sequence by adjacent pairs that can appear in either order, but cannot be separated by other look-ups.

Evaluating a formula for a rule's ideal guess requires a series of operations, some of which are *basic* in that they logically precede any other operation.

Like JC and CGCB, CGC derived rules' search implications assuming that subjects perform basic operations one at a time via adjacent look-ups, remember their results, and otherwise rely on repeated look-ups rather than memory.

Basic operations are then represented in a look-up sequence by adjacent pairs that can appear in either order, but cannot be separated by other look-ups.

E.g. *L1*'s ideal guess is $p^i[a^j+b^j]/2$, where p^i is its own target and a^j and b^j are other's lower and upper limits. *L1*'s characteristic look-up sequence is $\{[a^j, b^j], p^i\}$.

In this formula, *L1*'s only basic operation is $[a^j+b^j]$, part of averaging other's limits, is grouped within square brackets to show that a^j and b^j cannot be separated.

Other operations, whose look-ups grouped within curly brackets or parentheses, can appear in any order and their look-ups can be separated.

(*L1* does not need to look up its own limits because it can enter its ideal guess and rely on automatic adjustment to ensure that its adjusted guess is optimal.)

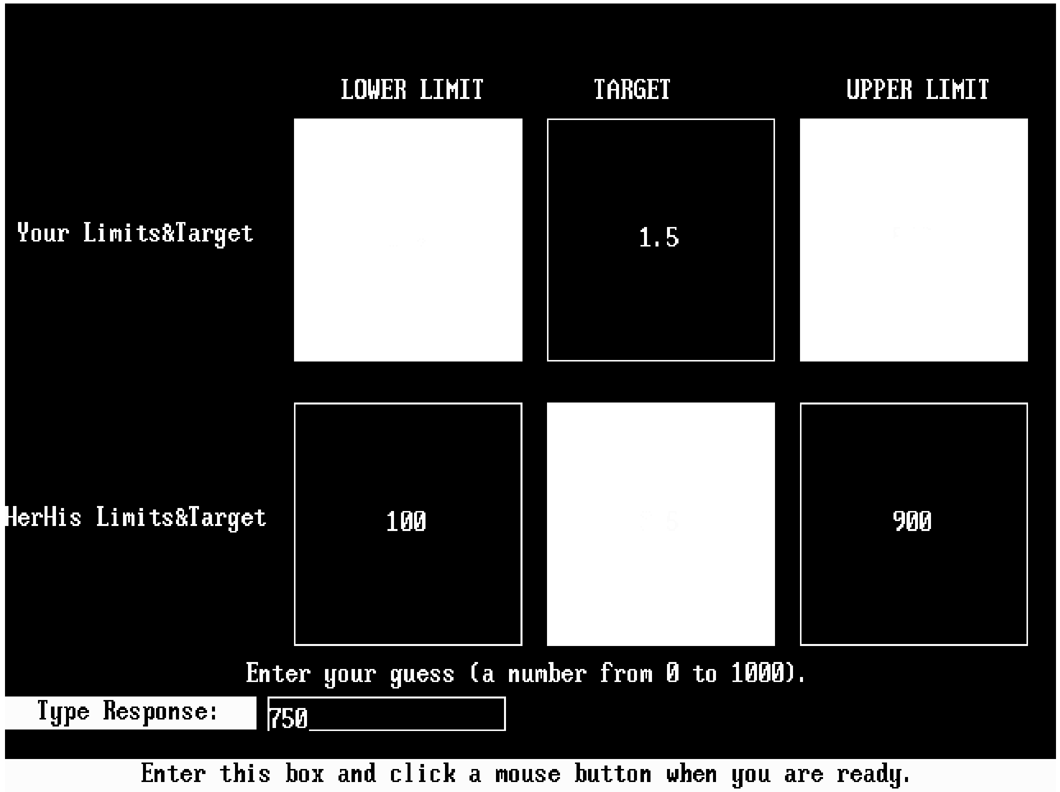

***L1*'s search implications** (subjects couldn't open more than one box at a time)

L1's ideal guess: $p^i[a^j+b^j]/2 = 750$.

L1's search $\{[a^j, b^j], p^i\} \equiv \{[4, 6], 2\}$ in the box numbers MouseLab records.

	a	***p***	***b***
You (*i*)	1	2	3
S/he (*j*)	4	5	6

MouseLab Box Numbers

L2's ideal guess is $p^i R(a^j,b^j;\ p^j[a^i+b^i]/2)$, where p^i is its own target, a^i and b^i are its own lower and upper limits, a^j and b^j are other's lower and upper limits, and $R(\cdot;\ \cdot)$ is the automatic adjustment function.

L2's model of other's *L1* guess is $p^j[a^i+b^i]/2$.

L2's characteristic look-up sequence is $\{([a^i,b^i],p^j),a^j,b^j,p^i\}$.

(*L2* needs to look up its own limits only to predict other's *L1* guess; like *L1* it can enter its own ideal guess and rely on automatic adjustment to its optimal guess.)

(*L2* needs to look up other's limits a^j and b^j to predict other's *L1 adjusted* guess.)

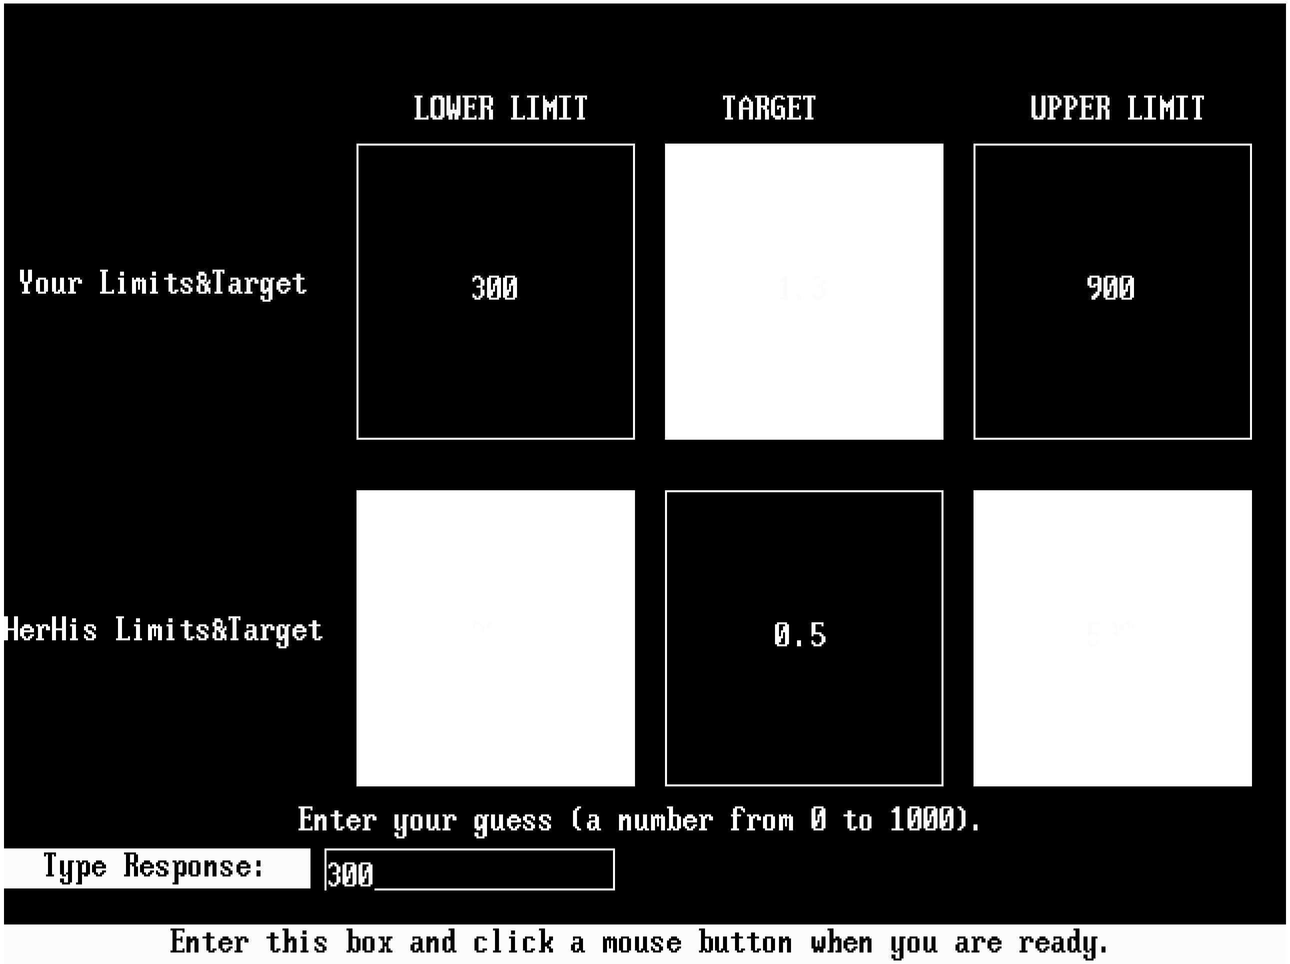

_L2_’s search implications: first step

L2’s model of its partner’s *L1* guess: $p^j[a^i+b^i]/2 = 300$.

L2’s ideal guess: $p^iR(a^j,b^j; p^j[a^i+b^i]/2) = 450$.

L2’s search $\{([a^i,b^i],p^j),a^j,b^j,p^i\} \equiv \{([1, 3], 5), 4, 6, 2\}$ in MouseLab box numbers.

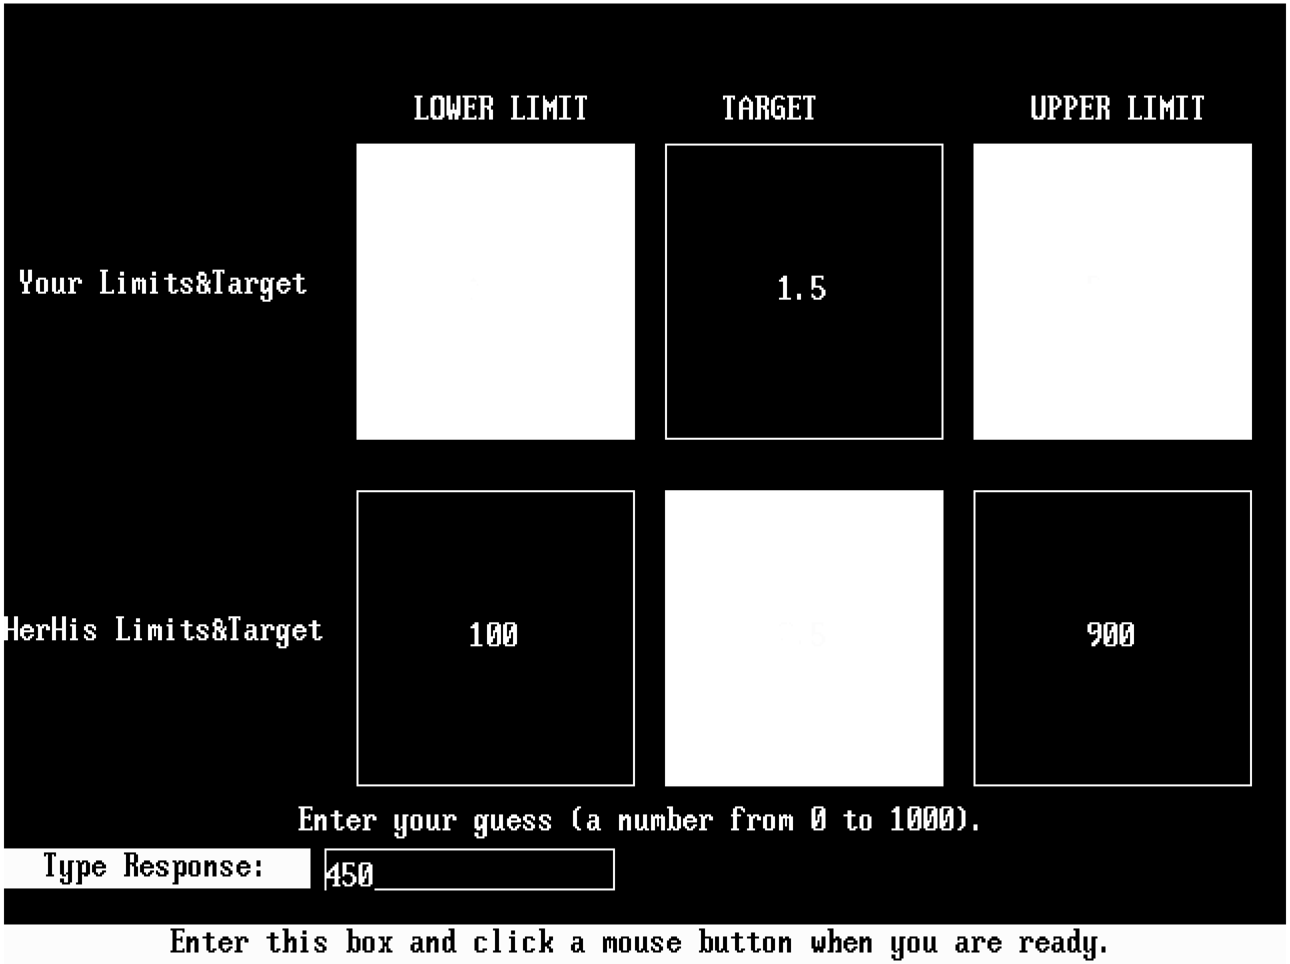

_L2_'s search implications: second step

L2's ideal guess: $p^i R(a^j,b^j;\, p^j[a^i+b^i]/2) = 450$.

L2's search $\{([a^i,b^i],p^j),a^j,b^j,p^i\} \equiv \{([1, 3], 5), 4, 6, 2\}$ in MouseLab box numbers.

Equilibrium can use any workable method to find its ideal guess; CGC allowed any method, and sought the one with minimal search requirements.

Equilibrium-checking (conjecturing guesses and checking them for consistency with equilibrium) is less demanding than other methods, but requires more luck than all but a few subjects appeared to have.

Equilibrium can use any workable method to find its ideal guess; CGC allowed any method, and sought the one with minimal search requirements.

Equilibrium-checking (conjecturing guesses and checking them for consistency with equilibrium) is less demanding than other methods, but requires more luck than all but a few subjects appeared to have.

CGC therefore allowed an *Equilibrium* player to use both targets to determine whether equilibrium is determined by upper or lower limits, and then to enter its own target times other's lower (upper) limit when the product of targets is < (>) 1, which CGC showed ensures that the player's adjusted guess is in equilibrium.

This has the same search requirements as equilibrium-checking except that it requires the targets to be adjacent; and thereby avoids the need for luck.

Equilibrium can use any workable method to find its ideal guess; CGC allowed any method, and sought the one with minimal search requirements.

Equilibrium-checking (conjecturing guesses and checking them for consistency with equilibrium) is less demanding than other methods, but requires more luck than all but a few subjects appeared to have.

CGC therefore allowed an *Equilibrium* player to use both targets to determine whether equilibrium is determined by upper or lower limits, and then to enter its own target times other's lower (upper) limit when the product of targets is < (>) 1, which CGC showed ensures that the player's adjusted guess is in equilibrium.

This has the same search requirements as equilibrium-checking except that it requires the targets to be adjacent; and thereby avoids the need for luck.

Equilibrium's ideal guess is then $p^i a^j$ if $p^i p^j < 1$ or $p^i b^j$ if $p^i p^j > 1$, and its search implications are $\{[p^i,p^j],a^j\} \equiv \{[2, 5], 4\}$ if $p^i p^j < 1$ or $\{[p^i,p^j],b^j\} \equiv \{[2, 5], 6\}$ if $p^i p^j > 1$.

Unlike in CGCB's and CJ's designs, in CGC's design *Equilibrium*'s search implications are just as simple as *L1*'s, and simpler than other rules'.

TABLE 4—TYPES' IDEAL GUESSES AND RELEVANT LOOK-UPS

Type	Ideal guess	Relevant look-ups
L1	$p^i[a^j + b^j]/2$	$\{[a^j, b^j], p^i\} \equiv \{[4, 6], 2\}$
L2	$p^iR(a^j, b^j; p^j[a^i + b^i]/2)$	$\{([a^i, b^i], p^j), a^j, b^j, p^i\} \equiv \{([1, 3], 5), 4, 6, 2\}$
L3	$p^iR(a^j, b^j; p^jR(a^i, b^i; p^i[a^j + b^j]/2))$	$\{([a^j, b^j], p^i), a^i, b^i, p^j\} \equiv \{([4, 6], 2), 1, 3, 5\}$
D1	$p^i(\max\{a^j, p^ja^i\} + \min\{p^jb^i, b^j\})/2$	$\{(a^j, [p^j, a^i]), (b^j, [p^j, b^i]), p^i\} \equiv \{(4, [5, 1]), (6, [5, 3]), 2\}$
D2	$p^i[\max\{\max\{a^j, p^ja^i\}, p^j\max\{a^i, p^ia^j\}\}$ $+ \min\{p^j\min\{p^ib^j, b^i\}, \min\{p^jb^i, b^j\}\}]/2$	$\{(a^i, [p^i, a^j]), (b^i, [p^i, b^j]), (a^j, [p^j, a^i]), (b^j, [p^j, b^i]), p^j, p^i\}$ $\equiv \{(1, [2, 4]), (3, [2, 6]), (4, [5, 1]), (6, [5, 3]), 5, 2\}$
Eq.	$\{a^i$ if $p^ia^j \leq a^i$ or $\min\{p^ia^j, b^i\}$ if $p^ia^j > a^i\}$ if $p^ip^j < 1$ or $\{b^i$ if $p^ib^j \geq b^i$ or $\max\{a^i, p^ib^j\}$ if $p^ib^j < b^i\}$ if $p^ip^j > 1$	$\{[p^i, p^j], a^j\} \equiv \{[2, 5], 4\}$ if $p^ip^j < 1$ or $\{[p^i, p^j], b^j\}$ $\equiv \{[2, 5], 6\}$ if $p^ip^j > 1$
Soph.	[no closed-form expression; *Sophisticated*'s search implications are the same as *D2*'s]	$\{(a^i, [p^i, a^j]), (b^i, [p^i, b^j]), (a^j, [p^j, a^i]), (b^j, [p^j, b^i]), p^j, p^i\}$ $\equiv \{(1, [2, 4]), (3, [2, 6]), (4, [5, 1]), (6, [5, 3]), 5, 2\}$

Notes: The most basic operations are represented by the innermost look-ups, grouped within square brackets; these can appear in any order, but *may not* be separated by other look-ups. Other operations are represented by look-ups grouped within parentheses or curly brackets; these can appear in any order, and *may* be separated by other look-ups. *Equilibrium*'s minimal search implications are derived not directly from *Equilibrium*'s ideal guesses, but from p^ia^j when $p^ip^j < 1$ and p^ib^j when $p^ip^j > 1$ via Observation 1 (see on-line Appendix H).

Note that although most theorists instinctively identify *Lk* with *Dk-1*, which both respect *k* rounds of iterated dominance, they are cognitively very different:

- *Lk* starts with a naïve prior over the other's decisions and iterates the best-response mapping.

- *Dk-1* starts with reasoning based on iterated knowledge of rationality and closes the process with a naïve prior.

This difference is obscured in a design as simple as Nagel's, and does not show up clearly from decisions even in more powerful designs.

Note that although most theorists instinctively identify *Lk* with *Dk-1*, which both respect *k* rounds of iterated dominance, they are cognitively very different:

- *Lk* starts with a naïve prior over the other's decisions and iterates the best-response mapping.

- *Dk-1* starts with reasoning based on iterated knowledge of rationality and closes the process with a naïve prior.

This difference is obscured in a design as simple as Nagel's, and does not show up clearly from decisions even in more powerful designs.

But in CGC's design these rules are separated clearly via search implications:

In Table 4:

- *L2*'s characteristic sequence is $\{([a^i,b^i],p^j),a^j,b^j,p^i\} \equiv \{([1, 3], 5), 4, 6, 2\}$.

- *D1*'s characteristic sequence is $\{(a^j,[p^j,a^i]),(b^j,[p^j,b^i]),p^i\} \equiv \{(4,[5,1]),(6,[5,3]),2\}$.

Search data

CGC's Baseline subjects played the game against other subjects.

CGC's Robot/Trained Subjects played the same games, but with each subject trained in and rewarded for following a rule: *L1*, *L2*, *L3*, *D1*, *D2*, or *Equilibrium*.

Search data

CGC's Baseline subjects played the game against other subjects.

CGC's Robot/Trained Subjects played the same games, but with each subject trained in and rewarded for following a rule: *L1*, *L2*, *L3*, *D1*, *D2*, or *Equilibrium*.

Search data for R/TS and Baseline subjects, chosen for high compliance with their rule's guesses (*not* compliance with any theory of search) suggest that:

- There is little difference between the look-up sequences of R/TS and Baseline subjects of a given rule (assigned rule for R/TS, apparent rule for Baseline).

- Table 4's relevant look-ups for a rule are dense in the search sequences for subjects with that rule (apparent or assigned), and the algorithms many are can be read directly from their searches, at least for the simpler rules.

- *Equilibrium* and *D2* subjects are stressed out but usually get decisions right.

Table 10.2. Selected Robot/Trained Subjects' Information Searches.

Subject	Type/Alt[a]	Game 1[b]	Game 2[b]
904	L1 (16)	1234564623	1234564321
1716	L1 (16)	14646213464623	46246213
1807	L1 (16)	462513	46213225
1607	L2 (16)	1354621313	1354613546213
1811	L2 (16)	1344465213*46	13465312564231356252
2008	L2 (16)	1113131313135423	13131356622333
1001	L3 (16)	46213521364*24623152	4621356425622231462562*62
1412	L3 (16)	1462315646231	462462546231546231
805	D1 (16)	1543564232132642	51453561536423
1601	D1 (16)	25451436231	5146536213
804	D1 (3)/L2 (16)	1543465213	5151353654623
1110	D2 (14)	1354642646*313	135134642163451463211136 414262135362*146546
1202	D2 (15)	246466135464641321342462 4226461246255*1224654646	123645132462426262241356 462*135242424661356462
704	DEq (16)	123456363256565365626365 6526514522626526	123456525123652625635256 262365456
1205	Eq (16)	12345642465256252563 52*465	123456244565565263212554 14666265425144526*31
1408	Eq (15)	12312345644563213211	1234564561236435241
2002	Eq (16)	142536125365253616361454 61345121345263	1436253614251425236256563

L1	{[4,6],2}
L2	{([1,3],5),4,6,2}
L3	{([4,6],2),1,3,5}
D1	{(4,[5,1], (6,[5,3]),2}
D2	{(1,[2,4]),(3,[2,6]),(4,[5,1],(6,[5,3]),5,2}
Eq	{[2,5],4} if pr. tar.<1,{[2,5],6} if > 1

Table 10.3. Selected Baseline Subjects' Information Searches.

Subject	Type/Alt[a]	Game 1[b]	Game 2[b]	Game 3[b]
101	L1 (15)	146246213	46213	462*46
118	L1 (15)	24613462624132*135	2462622131	24624246413*426
413	L1 (14)	123456545612346З*	12356462213*	264231
108	L2 (13)	135642	1356423	1356453
206	L2 (15)	533146213	53146231	5351642231
309	L2 (16)	1352	1352631526*2*3	135263
405	L2 (16)	144652313312546232 12512	1324562531564565 4546312315656262	3124565231*123654 55233**513
210	L3 (9) Eq (9) D2(8)	123456123456213456 254213654	1234564655622316 54456*2	1234556456123
302	L3 (7) Eq (7)	221135465645213213 45456*541	2135465662135454 6321*26654123	265413232145563214 563214523*654123
318	L1 (7) D1 (5)	13245646525213242* 1462	132465132*462	1346521323*4
417	Eq (8) L3 (7) L2 (5)	252531464656446531 6412524621213	25523662*3652435 63	5213636415265263* 652
404	Eq (9) L2 (6)	462135464655645515 21354*135462426256 356234131354645	46246135252426131 5463562	462135215634*52
202	Eq (8) D2 (7) L3 (7)	123456254613621342 *525	1234564456132554 6251356523	1234561235623
310	Eq (11)	123126544121565421 254362*21545 4*	1235462163262314 56*62	123655463213
315	Eq (11)	213465624163564121 325466	1346521246536561 213	132465544163*3625

L1	{[4,6],2}
L2	{([1,3],5),4,6,2}
L3	{([4,6],2),1,3,5}
D1	{(4,[5,1], (6,[5,3]),2}
D2	{(1,[2,4]),(3,[2,6]),(4,[5,1],(6,[5,3]),5,2}
Eq	{[2,5],4} if pr. tar.<1,{[2,5],6} if > 1

Recall that the large strategy spaces and varying targets and limits across CGC's games yield very strong separation via decisions: strategic "fingerprinting":

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

- CGC's other 45 main subjects' rules are less apparent from guesses; but *L1*, *L2*, *L3*, and *Equilibrium* are still the only ones in econometric estimates.

Recall that the large strategy spaces and varying targets and limits across CGC's games yield very strong separation via decisions: strategic "fingerprinting":

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

- CGC's other 45 main subjects' rules are less apparent from guesses; but *L1*, *L2*, *L3*, and *Equilibrium* are still the only ones in econometric estimates.

Most subjects' rules can be econometrically better identified by decisions and search than by decisions alone, and many can be identified from search alone (CGC, Tables 7A-B).

Adding search changes only a few subjects' estimated rules, with the guesses-and-search estimate resolving a tension in favor of the search-only estimate.

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.

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.

Even so, for some subjects, search is an important check on decisions:

- Baseline subject 309, with 16 exact *L2* guesses, misses some of *L2*'s relevant look-ups, avoiding deviations from *L2* only by luck (s/he later has a Eureka! moment between games 5 and 6, and from then on complies perfectly).

- Baseline subject 415 (not shown, CGC fn. 43) is an *L1* who fails Adjacency because s/he can remember three numbers at once, and in CGC's search analysis is therefore misclassified as *D1* (only clear failure in 71 subjects).

Other noteworthy analyses of search data

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

WSC find that both search and decision data are close to the predictions of a level-*k* model with *L0* anchored in truthfulness, in the style of Crawford's 2003 *AER* level-*k* analysis of signaling of intended decisions (CGCI, Section 9.1).

Such a level-*k* model explains two puzzling results from previous experiments:

- 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 costs of lying; and

- Despite the 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.

Other noteworthy analyses of search data

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

WSC find that both search and decision data are close to the predictions of a level-*k* model with *L0* anchored in truthfulness, in the style of Crawford's 2003 *AER* level-*k* analysis of signaling of intended decisions (CGCI, Section 9.1).

Such a level-*k* model explains two puzzling results from previous experiments:

- 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 costs of lying; and

- Despite the 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.

BCWC (Brocas, Carrillo, Wang, and Camerer, 2014 *REStud*) clustering analysis of search in zero-sum betting experiments, which confirms the level-*k* interpretation of most subjects' betting suggested by their decision data.