

LOOK-UPS AS THE WINDOWS OF THE STRATEGIC SOUL: STUDYING COGNITION VIA INFORMATION SEARCH IN GAME EXPERIMENTS¹

Vincent P. Crawford, University of California, San Diego

5 October 2006; revised 16 January 2007

1. Introduction

Because human decisions are the result of cognitive processes, theories of human behavior rest at least implicitly on assumptions about cognition. Neuroeconomics reflects the belief that using evidence on neural correlates of cognition will lead us to better theories of decisions.

Gul and Pesendorfer [2007], henceforth “GP”, argue that, on the contrary, because economic theory was intended to explain only decisions it should only be tested by observing decisions. They view neuroeconomics as a radical departure from economics in part because neural data concern involuntary, unconscious processes. Such processes are not decisions, so “our” theories cannot be about them.² Moreover, they argue, trying to extend our theories to explain neural data would require sacrificing important strengths of rational-choice analysis.

This paper attempts to narrow the gap between these views by discussing some recent experiments that elicit subjects’ initial responses to games with the goal of identifying the structure of their strategic thinking—subjects’ attempts to predict others’ decisions by taking their incentives into account.³ Strategic thinking can of course be studied in experiments that elicit decisions alone, via designs in which different models of cognition imply different decisions, as in Stahl and Wilson [1994, 1995] or “SW”; Nagel [1995]; and Ho, Camerer, and Weigelt [1998] or “HCW”. But the experiments I discuss study strategic thinking more directly, by monitoring and analyzing subjects’ searches for hidden but freely accessible payoff information, as in Camerer,

¹ This paper is based on joint work with Miguel Costa-Gomes, University of York, and Bruno Broseta, Red de Institutos Tecnológicos de la Comunidad Valenciana, particularly on Costa-Gomes and Crawford [2006, 2007]. I thank Miguel Costa-Gomes for our many discussions over the years, and Andrew Caplin for his very helpful comments on a previous draft. The experiments and analysis on which this paper is based were funded in part by the U.S. National Science Foundation under grant SES- 0100072 and the U.K. Economic & Social Research Council under grant R/000/22/3796.

² GP do allow an “inspirational” role for data other than decisions, but they exclude such data from theory-testing.

³ Why study strategic thinking when with enough experience in a stationary environment, even amoebas—or human reinforcement learners, who need not even know that they are playing a game—usually converge to equilibrium? Many applications of game theory involve situations with no clear precedents. (Should you sell U.S. airline stocks when the market re-opens after 9/11, or buy them on the anticipation that others will overreact?) Comparative statics and design questions inherently involve new games with new equilibria, which players cannot reach by copying behavior from analogous games. In such situations subjects’ initial responses are often plainly “strategic” but nonetheless deviate from equilibrium. Even in settings in which players can be expected to converge to equilibrium, the structure of strategic thinking can influence the rate of convergence and equilibrium selection.

Johnson, Rymon, and Sen [1993] and Johnson, Camerer, Sen, and Rymon [2002], henceforth collectively “CJ”; Costa-Gomes, Crawford, and Broseta [2001] or “CGCB”; and Costa-Gomes and Crawford [2006, 2007] or collectively “CGC”. My discussion draws extensively on CGC [2007], which reports and analyzes the information search data from CGC [2006].

CJ’s, CGCB’s, and CGC’s analyses of search rest on explicit models of cognition and therefore raise some of the same issues that GP raise about neuroeconomics. But unlike neural correlates of cognition, search is a voluntary, conscious process. Rational-choice analysis can therefore be used to describe it, eliminating one source of resistance to studying cognition. Further, the clarity of the insights into behavior these analyses yield is a “proof of concept” that shows how much can be gained by expanding the domain of analysis beyond decisions. Although the analysis of search data sidesteps some important issues raised by studying neural data, I hope that considering it will bring us closer to agreement on how, and whether, to do neuroeconomics.

CJ’s, CGCB’s, and CGC’s analyses suggest a concrete answer to GP’s challenge, Why study cognition if our goal is only to understand and predict decisions? In CGC’s decision data, for instance, most subjects deviate systematically from equilibrium, and their deviations are not well explained by noisy generalizations of equilibrium such as McKelvey and Palfrey’s [1995] quantal response equilibrium (“QRE”). Following SW and CGCB, CGC described their behavior via a structural non-equilibrium model in which each subject’s decisions are determined, in all games, by one of a small set of decision rules or *types* (as they are called in this literature). The possible types were restricted to general principles of decision-making, so that the theory’s predictions would be potentially as portable to new games as equilibrium predictions. CGC showed that to the extent that subjects’ deviations from equilibrium decisions can be distinguished from randomness, which is considerable, they are best explained by types that are rational, self-interested, and understand the game, but which base their beliefs on simplified models of others’ decisions. In other words, subjects’ deviations from equilibrium had mainly to do with how they model others’ decisions, not with nonstandard preferences or irrationality.⁴ Further, CGC’s analysis of subjects’ searches for hidden payoff information shows that extending the domain of the theory to include cognition and information search as well as conventional decisions allows more powerful tests and

⁴ This conclusion is consistent with SW’s, Nagel’s, HCW’s, and CGCB’s results, but their evidence is less clear.

precise identification of subjects' types, sometimes directly revealing the algorithms subjects use to process payoff information into decisions and/or distinguishing intended decisions from errors.⁵

One could still choose to use subjects' decisions alone to model their behavior via revealed preference, as in GP's proposal. This requires a generalization of their proposal, because as stated it refers only to individual decision problems. In games, given the lack of a rational-choice model of non-equilibrium beliefs, revealed preference would need to be augmented by assumptions about players' beliefs, presumably by assuming that beliefs are in equilibrium.⁶ However, as CGC's and previous experiments make clear, such an approach is unlikely to predict reliably beyond sample: In initial responses to new games, subjects' simplified models of others' decisions would yield different patterns of deviation from equilibrium. Only by coincidence would those patterns be well described by equilibrium (or QRE) with subjects' preferences as inferred from previous games. No empirically serious model of initial responses to games can ignore cognition, and relegating data other than conventional decisions to an inspirational role arbitrarily limits the use of a powerful tool for understanding and predicting decisions and may yield a biased view of behavior.

This critique highlights an important issue in judging GP's proposal. GP take the main problem of economics to be uncovering preferences, with rationality in the decision-theoretic sense (and in games, equilibrium) assumed. By contrast, CGC's and previous experiments on initial responses to games test (and mostly affirm) rationality; induce preferences (and leave little room for doubts about risk or social preferences); and focus instead on discovering the general principles that govern strategic thinking. Although many applications involve games whose players have enough clear precedents to justify equilibrium via learning, applications involving novel strategic situations, and where strategic thinking is the main source of uncertainty, are not uncommon. The costs and benefits of GP's proposal should be evaluated in the latter applications as well as in the cases involving individual decisions they use to develop and illustrate their proposal.

The rest of the paper is organized as follows. Section 2 begins by reviewing CJ's and CGCB's designs and results, with the goal of introducing the key design and modeling issues in studying cognition via information search in the contexts in which they first emerged. Section 3

⁵ CGCB's and CJ's analyses make this point in different ways. Camerer [2007], Caplin [2007], and Schotter [2007] argue cogently, in complementary ways, for the use of non-choice data, and outline frameworks to guide their use in analyses. Koszegi and Rabin [2007] discuss using decision data to distinguish intended decisions from errors.

⁶ Or, if the rules and possible preferences allow multiple equilibria, the equilibrium identified by some agreed-upon selection principle. Although GP's paper focuses entirely on individual decisions, private communications suggest that they accept the need for extending their proposal to games by assuming equilibrium.

reviews CGC's [2006] use of decision data to identify subjects' types and the evidence that the main source of their deviations from equilibrium is cognitive, not preference-based. Section 4 introduces CGCB's and CGC's [2006, 2007] model of cognition and search and uses CGC's search data to illustrate its use in interpreting search data. Section 5 highlights questions raised by CGC's (2006) analysis of decisions that seem likely to continue to resist analysis via decisions alone, but which search analysis might answer. Section 6 outlines a deeper explanation of the assumptions that underlie CGCB's and CGC's model of cognition and search, which views search strategies as rational decisions under plausible assumptions about the benefits and costs of search and constraints on working memory.⁷ Section 7 is the conclusion. Throughout the paper I assume that subjects are rational, risk-neutral decision-makers; but I allow "social" preferences that reflect altruism, spite, fairness, or reciprocity when they seem important, as indicated below.

2. Early Experiments that Studied Cognition in Games by Monitoring Information Search

In this section I review CJ's and CGCB's experimental designs and results. Their and CGC's experiments randomly and anonymously paired subjects to play series of different but related two-person games, with different partners each play and no feedback between plays. The goal was to suppress learning and repeated-game effects in order to elicit subjects' responses, game by game, to each game as if played in isolation, and so to reveal strategic thinking as clearly as possible.⁸

The structure of the games was publicly announced except for hidden, varying payoff parameters, to which subjects were given free access, game by game, one at a time, before making their decisions.⁹ With low search costs, free access made the entire structure effectively public knowledge, allowing the results to be used to test theories of behavior in complete-information versions of the games.¹⁰ Varying the payoff parameters makes it impossible for subjects to

⁷ The proposed explanation differs greatly from classical search theory in purpose, but only slightly in methods.

⁸ "Eureka!" learning remains possible, but it can be tested for and seems to be rare. Initial responses yield insights into cognition that also help us think about how to model learning from experience, but that is another story.

⁹ Access was via a MouseLab interface that automatically records a sequence of parameter opening and closing times, which makes it possible to test models of the order and/or duration of parameter look-ups. Subjects were not allowed to write, and the frequencies with which they looked up the parameters made clear that they did not memorize them. Subjects were taught the mechanics of looking up parameters and entering decisions, but not information-search strategies. MouseLab is an automated way to track search as in eye-movement studies of individual decisions (Payne, Bettman, and Johnson [1993], <http://www.cebiz.org/mouselab.htm>). Wang, Spezio, and Camerer [2006] illustrate the use of a modern, more powerful eye-tracking method.

¹⁰ A partial exception is that CJ's experiments evoked nonpecuniary social preferences like those in ultimatum experiments, and these and subjects' risk aversion are uncontrolled and privately known. Privately known social preferences are easily accommodated in the analysis of CJ's results, and risk aversion was probably insignificant.

remember the current game's parameters from previous plays, and so gives them incentives to search for the information their decision rules require. It also allows stronger separation of the decisions implied by equilibrium and leading alternative decision rules than in designs such as Nagel's or HCW's, in which subjects play the same game over and over again.

2a. Camerer, Johnson, Rymon, and Sen's alternating-offers bargaining experiments

CJ [1993, 2002] pioneered the use of search for hidden payoff parameters to study cognition in games, eliciting subjects' initial responses to series of three-period alternating-offers bargaining games.¹¹ Previous experiments yielded large, systematic deviations from the subgame-perfect equilibrium offer and acceptance decisions when players have pecuniary preferences, like those observed in ultimatum experiments. The deviations were attributed to cognitive limitations preventing subjects from doing the required backward induction, or believing that their partners would; to subjects having social preferences that modify their pecuniary payoffs; or both. Most researchers now agree that both factors are important, but in the early 1990s this was less clear.

CJ addressed the cognitive aspect of this question more directly by creating a design to study cognition via search and by deriving cognitive implications of alternative models of behavior and using them to analyze the search data. Within a publicly announced structure, they presented each bargaining game to subjects in extensive form as in Figure 1, as a sequence of three pies and the associated offer and acceptance decisions. Discounting was simulated by shrinking the pies over time, from roughly \$5.00 in round 1 to roughly \$2.50 in round 2 and \$1.25 in round 3; but the pies were varied slightly from game to game, to preserve subjects' incentives to search.

The pies were normally hidden in "boxes" as for rounds 2 and 3 in Figure 1, but subjects were allowed to look them up as often as desired, one at a time. In Figure 1 the subject has opened the box to look up the \$5.00 round-1 pie.¹² Subjects' knowledge of the structure of the games and their free access to the pies allowed them to evaluate their own and their partners' pecuniary payoffs for any combination of offer and acceptance decisions.

If free access to the pies induces public knowledge of pecuniary payoffs, and if it is also public knowledge that subjects maximize their own expected pecuniary payoffs, then the results

¹¹ CJ's [1993] and CJ's [2002] designs differed in some ways, for example framing in losses versus in gains, that are not important for my purposes and are not discussed here. At roughly the same time in the early 1990s, CJ [2004] did a MouseLab study of forward induction in extensive-form games. Algaze (Croson) [1990] reported a very brief study of search for hidden payoff information in matrix games. Neither of the latter papers is discussed here.

¹² CJ used a "rollover" option in MouseLab, in which subjects could open the box that concealed a pie by moving the cursor into it, revealing the pie for as long as the cursor was in the box. Subjects could also use the interface to look up their roles in each round; but these were known and those look-ups were not reported or analyzed.

can be used to test theories of behavior in complete-information versions of the game, which has a unique subgame-perfect equilibrium whose offer and acceptance decisions are easily computed by backward induction. Even if players have privately observed social preferences, the incomplete-information version of the game has a generically unique sequential equilibrium whose strategies are easily computed by backward induction. In each case, the subgame-perfect or sequential equilibrium initial offer depends on both the second- and third-round pies, so that the search requirements of equilibrium are mostly independent of preferences.¹³ From now on I use “subgame-perfect equilibrium” to include pecuniary payoff-maximization.

In CJ’s baseline treatment, in which subjects were rewarded according to their payoffs playing the games against each other, subjects’ decisions were far from the subgame-perfect equilibrium, replicating the results of previous studies and suggesting that requiring subjects to look up the pies did not significantly affect their decisions.

CJ took the analysis a step further by using a model of cognition and search to analyze the search data. They first noted that 10% of their baseline subjects never looked at the third-round pie and 19% never looked at the second-round pie. Thus, even if those subjects’ decisions conform to equilibrium (given some specification of preferences, with or without a social component), they cannot possibly be making equilibrium decisions for the reasons the theory assumes. In a non-magical world, their compliance with equilibrium cannot be expected to persist beyond sample.

This observation motivates a basic general restriction on how cognition drives search, which—anticipating CGCB’s term for it—I call “Occurrence”: If a subject’s decision rule depends on a piece of hidden payoff information, then that piece must appear in her/his look-up sequence. Occurrence, as a cognitive restriction, goes against GP’s proposal, but it is still uncontroversial enough to be widely accepted by theorists. In this case at least, the epistemic foundations of equilibrium have implications for the interpretation of decisions it is hard to justify ignoring.

If Occurrence were the whole story, there would be little to gain from studying cognition via search. Because CJ’s subjects who never looked at the second- or third-round pies tended to make decisions far from subgame-perfect equilibrium, there is little risk of misinterpreting them; even so, Occurrence helps by ruling out explanations in which subjects’ decisions are in sequential

¹³ Only “mostly” because with only pecuniary preferences, the first-round pie, as long as it is large enough, does not affect the equilibrium initial offer. With social preferences the first-round pie may be relevant because it may influence the responder’s acceptance decision.

equilibrium for extreme distributions of social preferences. Inferences based on Occurrence are sometimes useful in CGCB's and, as we will see, CGC's analyses as well, but the full power of monitoring search depends on analyzing the order, and perhaps the duration, of subjects' look-ups.

CJ's analysis of order and duration is based on the argument that in their design the backward induction that is the easiest way to compute sequential or subgame-perfect equilibrium decisions has a characteristic search pattern, in which subjects first look up the third-round pie, then the second-round pie (possibly re-checking the third), and so on, with most transitions from adjacent later- to earlier-round pies. Their argument rests on the empirical generalization that most subjects use the interface as a computational aid, making the comparisons or other operations on which their decisions are based via adjacent look-ups and relying on repeated look-ups rather than memory. This observation motivates another basic restriction, which—again anticipating CGCB's term—I call “Adjacency”: The hidden parameters associated with the simplest of the operations on which a subject's decision rule depends will appear as adjacent look-ups in his look-up sequence.¹⁴

Adjacency, unlike Occurrence, requires assumptions that not all theorists find compelling. It is theoretically possible for a subject to scan the pies in any order, memorize them, and then “go into his brain” to figure out what to do, in which case the order and duration of his look-ups will reveal nothing about cognition. (Here, brain imaging has a potential advantage over monitoring search because involuntary correlates of such a subject's thinking may still be observable.)

Fortunately, subjects' searches in designs like CJ's, CGCB's, and CGC's exhibit strong regularities that make Adjacency a reasonable working hypothesis. When challenged, CJ defended their Adjacency-based characterization of backward-induction search by running a “robot” treatment with the same games as their baseline, in which subjects were told that they were playing against a computer that simulated a rational, self-interested player. This was followed after four periods by a “robot/trained subjects” treatment in which the same subjects received training in backward induction (but not search) and continued to play against robots as before. The latter subjects' search patterns were close to the backward-induction pattern [CJ 2002, Figure 6]. Although the shift in search patterns was small prior to training, these results provide support for CJ's characterization, Adjacency, and of course Occurrence. As illustrated below, further (and

¹⁴ This informal definition, like the one for Occurrence, is intentionally vague regarding how often look-ups or operations appear to accommodate variations in CJ's, CGCB's and CGC's use of Occurrence and Adjacency. The notions are made more precise in CGCB's analysis and, as explained below, CGC's. Note that both are general restrictions on how cognition drives search, which can be applied across a variety of games and decision rules.

sometimes stronger) support for Adjacency is provided by CGCB's trained subjects, CGC's robot/trained subjects with high compliance with their assigned type's guesses (Table 2), and CGC's baseline subjects with high compliance with their apparent rule's guesses (Table 3).

CJ's robot subjects' offer and acceptance decisions were shifted away from the baseline patterns toward subgame-perfect equilibrium, but were still far from it. Their robot/trained subjects' decisions were approximately in subgame-perfect equilibrium [CJ 2002, Table II]. These shifts can be attributed to the robot treatment's "turning off" social preferences, assuming subjects don't think of experimenters or their funding agencies as "people"; the robot treatment's eliminating strategic uncertainty; and/or cognition. CJ suggest that the deviations from equilibrium in the baseline are due to a combination of social preferences and cognition, with both important.

Returning to cognition and search, CJ's baseline subjects' searches were nearly the opposite of the searches of robot/trained subjects and CJ's characterization of backward induction search: baseline subjects spent 60-75% of the time looking up the first-round pie and only 20-30% looking up the second-round pie and 5-10% looking up the third-round pie, with most transitions forward, from earlier to later rounds. Importantly, subjects who looked up the second- and third-round pies more often, or had more backward transitions, also had a weak tendency to make, or accept, offers closer to the subgame-perfect equilibrium [CJ 2002, Figures 4-5]. Thus, CJ's baseline subjects' deviations from backward induction search were correlated with their deviations from subgame-perfect equilibrium decisions, in the direction that an epistemic, procedural view of subjects' decision-making would suggest. Although the correlation is weak, this result is an exciting first indication that subjects' search patterns might reveal something about their strategic thinking.

2b. Costa-Gomes, Crawford, and Broseta's matrix-game experiments

CGCB adapted CJ's methods, building on SW's [1994, 1995] designs, to study cognition via search in a series of 18 2×2 , 2×3 , or 2×4 matrix games with unique pure-strategy equilibria, some of which can be identified by iterated dominance and some without pure-strategy dominance. The games were designed to turn off social preferences, and CGCB's results show little evidence of them. I therefore assume that CGCB's subjects maximized their own expected pecuniary payoffs.

Within a publicly announced structure, CGCB presented each game to subjects via MouseLab, as a matrix with players' payoffs spatially separated to ease cognition and clarify inferences from search. The payoffs were hidden but subjects were allowed to look them up as often as desired. In the 2×2 game in Figure 2 the subject, framed as the row player, has opened the

box with his own payoff, 42, when he chooses decision # and his partner chooses @.¹⁵ If free access induces public knowledge of the payoffs and it is public knowledge that subjects maximize their expectations, then the structure is public knowledge and the results can be used to test theories of behavior in complete-information versions of the games.

Although there are close connections between epistemic analyses of equilibrium decisions in extensive- and normal-form games, their cognitive foundations are very different. The different presentation of payoff information in CGCB's matrix games allows them to explore aspects of strategic thinking that do not come into play in CJ's bargaining games. Moreover, although CGCB's games have small strategy spaces, their sequence of 18 games creates a large space of possible decision histories, which allows their design to separate the implications of leading normal-form theories of decisions more strongly than in previous designs in which subjects play series of different matrix games with small strategy spaces as in SW [1994, 1995], or in which they repeatedly play the same normal-form game with large strategy spaces as in Nagel and HCW.

Finally, and most importantly here, the 8-16 hidden payoffs in CGCB's design create a large space of possible information searches, which allows the design to separate leading theories' implications for search as well as decisions. In CJ's design, a subject's searches can vary in only one important dimension: backward or forward in the pies. Measuring a subject's searches in this dimension can convey a limited amount of information about his strategic thinking—though as we have seen, this information can be quite revealing. In CGCB's games, by contrast, a subject's searches can vary in three important dimensions: up-down (or not) in his own payoffs, left-right (or not) in his partner's payoffs, and the frequency of transitions from his own to his partner's payoffs. With the subject framed as the row player in Figure 2, it is clear that, assuming Adjacency, the first of these dimensions is naturally associated with decision-theoretic rationality, the second with using others' incentives to predict their decisions, and the third with interpersonal payoff comparisons. It would be difficult to imagine an empirically successful theory of initial responses to this kind of game in which those three traits were not independently variable and important. Only a design with a search space as rich as CGCB's can separate the implications of alternative theories for both search and decisions strongly enough to identify their relationships.

¹⁵ Instead of the rollover option CJ used, CGCB used a "click" option, in which subjects could open a box by moving the cursor into it and left-clicking the mouse. Before he could continue, a subject had to close the box by right-clicking, which could be done from anywhere in the display.

In addition to a baseline treatment that paired subjects to play the 18 games with other subjects, CGCB conducted a trained subjects treatment, identical to the baseline except that each subject was trained and rewarded for identifying equilibrium decisions. This treatment confirms that subjects trained and motivated to find equilibrium guesses could do so; and provides data on *Equilibrium* search behavior that are helpful in evaluating CGCB's model of cognition and search.

CGCB's games have unique equilibria that are easily identified by direct checking, best-response dynamics (which always converges in their games), or (in most of their games) iterated pure-strategy dominance. Yet, as in previous studies of initial responses to matrix games, CGCB found systematic patterns of deviation from equilibrium, with high equilibrium compliance in games solvable by one or two rounds of iterated dominance but much lower compliance in games solvable by three rounds or the circular logic of equilibrium without dominance [CGCB, Table II]. These patterns are not well explained by noisy generalizations of equilibrium such as QRE. CGCB explained them via a structural non-equilibrium model of initial responses in the spirit of SW's, Nagel's, and HCW's models, in which each subject's decisions are determined, in all games, by one of a small set of types, which determines his decisions, with error, in each game. The possible types were restricted to general principles of decision-making, so that the theory's predictions would be potentially as portable to new games as equilibrium predictions.

The leading types in CGCB's analysis include *L1* (for Level 1, as named by SW), called *Naive* in CGCB and *L1* here from now on, which best responds to a uniform random *L0* "anchoring type"; *L2*, which best responds to *L1*; *Equilibrium*, which makes its equilibrium decision; *D1* (Dominance 1), which does one round of deletion of dominated decisions and then best responds to a uniform prior over the other's remaining decisions; *D2*, which does two rounds of iterated deletion and then best responds to a uniform prior over the other's remaining decisions; and *Sophisticated*, which best responds to the probabilities of other's decisions, as estimated from subjects' observed frequencies, included to test whether subjects have prior understanding of others' decisions that transcends simple rules.¹⁶ Because CGCB gave first priority to separating strategic from nonstrategic types, *L1*'s decisions were perfectly confounded with those of a maximax type CGCB called *Optimistic*. CGCB's econometric analysis of decisions alone estimated high frequencies of *L1*, *L2*, and *D1*. Because those types mimic equilibrium in simple

¹⁶ Lk 's and $Dk-1$'s decisions both survive k rounds of iterated elimination of dominated decisions and so in two-person games are k -rationalizable [Bernheim 1984]. Although $Dk-1$ types are closer to how theorists analyze games, Lk types seem more natural and predominate in applications.

games but deviate systematically in more complex games, this estimated type distribution allows the model to explain the aggregate relationship between complexity and equilibrium compliance.

Turning to CGCB's analysis of search, the main difficulty was imposing enough structure on the enormous spaces of possible decision and search histories to describe subjects' behavior in a comprehensible way. Although CJ identified a correlation (and a "right" direction for it) between subjects' decision and search deviations from subgame-perfect equilibrium in their alternating-offers bargaining games, their analysis does not show how to define or identify such a relationship in the higher-dimensional spaces of possible decisions and searches created by CGCB's design.

CGCB addressed this issue by using the types as models of cognition and search as well as decisions. They took an explicitly procedural view of decision-making, in which a subject's type and the associated cognitive process determine his search, and his type and search then determine his decision, game by game.¹⁷ They characterized the link between cognition and search via the Occurrence and Adjacency restrictions described above, which generalize the ideas behind CJ's characterization of backward-induction search to a much wider class of games, patterns of hidden payoff information, and types. With these restrictions on cognition and search, the types provide a kind of basis for the spaces of possible decision and search histories, imposing enough structure to make it meaningful to ask whether subjects' decisions and searches are related in a coherent way.

Incorporating search into the econometric analysis yields a somewhat different view of subjects' deviations from equilibrium than previous analyses of decisions. It shifts CGCB's estimated type distribution toward *L1* at the expense of *Optimistic* and *D1*, leaving *L1* and *L2* as the only empirically important types. Part of this shift occurs because *L1*'s searches, unlike *L1*'s decisions, are clearly separated from *Optimistic*'s, and *L1*'s search implications explain more of the variation in subjects' searches and decisions than *Optimistic*'s, which are too unrestrictive to be useful. Another part of the shift occurs because *L1*'s search implications explain more of the variation in subjects' searches and decisions than *D1*'s, which are much more restrictive than *Optimistic*'s but too weakly correlated with subjects' observed decisions. *D1* loses frequency to *L2* as well, even though their decisions are only weakly separated in CGCB's design, because *L2*'s search implications explain more of the variation in subjects' searches and decisions. Thus,

¹⁷ Because a type's search implications depend not only on what decisions it specifies, but why, something like a types-based model seems necessary here. In CJ [1993] types are implicit in the discussion and limited to two, which might be called "subgame-perfect equilibrium" and "other". CJ [2002] adapted CGCB's analysis by defining extensive-form "types" modeled after CGCB's and SW's normal-form types, using them to construct a more structured data analysis than CJ's [1993].

analyzing search not only yields more precise estimates of subjects' types, it can correct distortions in type estimates based on decisions alone that stem from a design's failure to fully separate types.

These shifts illustrate an important principle. Because the number of experimental treatments and subjects that can be run is limited, data are scarce relative to the plausible theories of behavior, and tradeoffs in discriminating among theories are inevitable. Gathering search (or other non-decision data) as well as decision data can make such tradeoffs less stringent. Although gathering and analyzing non-decision data has its own costs, the optimal amount to gather is not always zero.

Overall, CGCB's analysis of decisions and search gives a strikingly simple view of behavior, with $L1$ and $L2$ making up 90% of the population. This type distribution and the clear relationships between subjects' cognition as revealed by search and their decisions support my claim that their deviations from equilibrium in these games are due mainly to how they think about others.

3. Costa-Gomes and Crawford's Two-Person Guessing Game Experiments

CGC [2006, 2007] adapted CGCB's methods to elicit subjects' initial responses to a series of 16 dominance-solvable two-person guessing games, cousins of Nagel's and HCW's n -person guessing games. In this section I review CGC's design and their results for decisions, which provide even stronger evidence that the deviations from equilibrium in initial responses to games are due mainly to strategic thinking. In Section 4 I review CGC's analysis of cognition and search.

3a. CGC's design

In CGC's games, newly designed for the purpose of studying cognition via decisions and search, two players make simultaneous guesses. Each player has his own lower and upper limit, both strictly positive, as in some of HCW's games, to ensure finite dominance-solvability. Unlike in previous designs, however, players are not required to guess between their limits: To enhance the separation of types via search, guesses outside the limits are automatically adjusted up to the lower limit or down to the upper limit as necessary. Thus, the only thing about a guess that affects the outcome is the adjusted guess it leads to. Each player also has his own target, and (unlike in Nagel's and HCW's "winner-take-all" games) his payoff is higher, the closer his adjusted guess is to his target times his partner's adjusted guess.

In the most important departure from previous guessing designs, the targets and limits vary independently across players and games, with the targets either both less than one, both greater

than one, or (unlike in previous designs) mixed.¹⁸ The resulting games are asymmetric and dominance-solvable in 3 to 52 rounds, with essentially unique equilibria determined (but not always directly) by players' lower limits when the product of the targets is less than one or their upper limits when the product is greater than one. In game 13 in Figures 4-7, for instance, player i has limits 300 and 500 and target 0.7, and player j has limits 100 and 900 and target 1.5 [CGC, Table 3]. The product of targets is $1.05 > 1$, player i 's equilibrium guess is at his upper limit 500, and player j 's equilibrium guess is at his best response to 500 of 750, below his upper limit.

From the point of view of studying decisions, CGC's design combines the main strengths of SW's and CGCB's designs, with subjects playing sequences of different but related games, and the main strengths of Nagel's and HCW's designs, games with very large strategy spaces. This combination greatly enhances the separation of *Equilibrium* and other leading types' decisions.

CGC's games explore different aspects of strategic thinking than CJ's, CGCB's, or Nagel's and HCW's games. Of particular note is the subtle way in which the location of the equilibrium is determined by the product of players' targets, which adds greatly to the power of the design to distinguish equilibrium from boundedly rational strategic thinking. The only important difference between some of CGC's games is whether the product of targets is slightly greater or slightly less than one. *Equilibrium* responds very strongly to this difference, but low-level Lk or Dk types, whose guesses vary continuously with the targets, respond much less. Also noteworthy is the strong separation of Lk 's and $Dk-1$'s decisions, which are perfectly confounded in most of Nagel's and HCW's treatments and only weakly separated in their other treatments and in CGCB's design.

In addition to a baseline treatment that paired subjects to play the 16 games with other subjects, CGC conducted six different robot/trained subjects treatments, identical to the baseline except that each subject was trained and rewarded as a type: $L1$, $L2$, $L3$, $D1$, $D2$, or *Equilibrium*. These treatments assess the types' cognitive demands, confirming for instance that subjects trained and motivated to make equilibrium guesses could do so; and provide data on the search behavior of subjects of known types that are helpful in evaluating the model of cognition and search.

In all treatments, within a publicly announced structure, CGC presented each game to subjects as an array of targets and limits, with those payoff parameters hidden but subjects allowed to look them up as often as desired, one at a time, using MouseLab's click option as in CGCB. In Figure 3 the subject has opened the box to look up his own ("Your") lower limit, 100.

¹⁸ In previous designs the targets and limits were the same for both players, and varied only across treatments.

3b. CGC's analysis of decisions

The strong separation of types' implications for guesses [CGC 2006, Figure 5] and the clarity of CGC's baseline subjects' responses allow many of their types to be confidently identified from guesses alone. Of 88 subjects, 43 have clear "fingerprints" in that they made guesses that complied *exactly* (within 0.5) with one type's guesses in 7-16 of the games (20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*).¹⁹ Figure 4 [CGC 2006, Figure 2] shows the fingerprints of the 12 whose apparent types were *L2*. Of their 192 (= 12×16) guesses, 138 (72%) were exact, which means they tracked the complex pattern of the games' *L2* guesses with a remarkable degree of accuracy. I stress that these baseline subjects, unlike the robot/trained subjects, were taught nothing about strategic thinking: The models of others' guesses implicit in their apparent types were self-generated.

Given how strongly CGC's design separates types' guesses, and that guesses could take 200-800 different rounded values, these 43 subjects' compliance is far higher than could occur by chance. Further, because the types specify precise, well-separated guess sequences in a very large space of possibilities, their compliance rules out alternative interpretations of their guesses.²⁰ In particular, because the types build in risk-neutral, self-interested rationality and perfect models of the game, the deviations from equilibrium of the 35 whose apparent types are *L1*, *L2*, or *L3* can be attributed to non-equilibrium beliefs, not irrationality, risk aversion, altruism, spite, or confusion.

CGC's other 45 subjects' types are less apparent from their guesses; but *L1*, *L2*, and hybrids of *L3* and/or *Equilibrium* are still the only types that show up in econometric estimates.²¹ The fact that most subjects follow low-level *Lk* types, which mimic equilibrium in games that are dominance-solvable in small numbers of rounds but deviate systematically in some more complex

¹⁹ 11 of these subjects were from an "open boxes" treatment, not discussed here, identical to the baseline but with the parameters continually visible. The results of this treatment (and analogous treatments in CJ and CGCB) confirm that making subjects look up the parameters does not significantly affect their decisions, so that the data can be pooled with baseline decision data, as here. CGC's open boxes subjects have numbers that begin with a 5.

²⁰ By contrast, in SW's or CGCB's matrix-game designs, even a perfect fit does not distinguish a subject's best-fitting type from nearby omitted types; and in Nagel's and HCW's guessing-game designs, with large strategy spaces but with each subject playing only one game repeatedly, the ambiguity is worse.

²¹ Nagel's results are often viewed as evidence that subjects perform finitely iterated dominance, as in *Dk-1*. But *Lk*'s and *Dk-1*'s decisions are perfectly confounded in Nagel's main treatments and weakly separated in Nagel's and HCW's other treatments and in CGCB's design. CGC's clear separation of *Lk* from *Dk-1* allows them to conclude that *Dk* types don't exist in significant numbers, at least in this setting, and thus that subjects respect low levels of iterated dominance as a by-product of following *Lk* types, not because they explicitly perform it. *Sophisticated*, which is clearly separated from *Equilibrium*, also doesn't exist in significant numbers. CGC's [2006, Section II.D] specification test rules out significant numbers of other types omitted from the specification.

games, also explains the inverse relationship between strategic complexity and equilibrium compliance observed in CGCB and previous experiments [CGCB, Table II].

CGC's results for decisions provide very strong evidence that subjects' deviations from equilibrium in initial responses to games are due mainly to non-equilibrium strategic thinking, not preferences or irrationality. As noted in the Introduction, one could still use subjects' guesses alone to model their behavior via revealed preference, but such a model would misattribute the cause of the deviations, and so would predict well beyond sample only by coincidence.

4. Costa-Gomes and Crawford's Analysis of Cognition and Search

CGC's [2006, Section II.E; 2007] model of cognition and search refines CGCB's model, adapting their Occurrence and Adjacency restrictions to give a tractable characterization of each type's search requirements. With regard to search, CGC's design combines the strengths of CJ's presentation of games as functions of a small number of hidden parameters within an intuitive common structure, which allows subjects to focus on predicting others' responses without getting lost in the details of the structure; and CGCB's high-dimensional search spaces, which make search more informative and allow greater separation via search. CGC's design strongly and independently separates the implications of leading types for search and decisions, which makes it easier to identify relationships between them and multiplies the power of the design. Finally, it makes each type's implications search independent of the game, which simplifies the analysis.²²

This section begins with a discussion of the issues that arise in specifying a model of cognition and search. It then presents CGC's leading types' search requirements and illustrates how they are derived. Finally, it presents sample search data for some of CGC's robot/trained and baseline subjects. As these data will be used to show, CGC's design and characterization of types' search implications make it possible to read the algorithms that a large minority of subjects used to choose their guesses directly from their search sequences. Other subjects' cognition is not apparent from their searches, but CGC's (2006) measures of their compliance with leading types' search implications have considerable discriminatory power in the econometric analysis, often allowing those subjects' types to be reliably estimated from searches alone, without regard to guesses.

²² By contrast, the lack of a simple common structure in CGCB's design makes rules' search implications vary from game to game in ways so complex you need a "codebook" to identify them.

4a. Specification issues

Studying cognition via search requires a model of how cognition determines subjects' look-up sequences. Previous papers have taken quite different positions on this issue. CJ's analysis gave roughly equal weight to look-up durations and total numbers of look-ups ("acquisitions") of each pie and to the numbers of transitions between look-ups of adjacent pies. Rubinstein's [2007] analysis considered only durations. Gabaix et al.'s [2006] focused on total numbers of look-ups rather than durations, but also considered some aspects of the order of look-ups. CGCB's and CGC's analyses focused instead on which look-ups subjects make, in the sense of Occurrence, and on the order of look-ups in the sense of Adjacency, relegating durations to a secondary role.²³

On another dimension, CJ's and Rubinstein's analyses and most of Gabaix et al.'s aggregated search data across subjects and over time, while CGCB and CGC took the position that cognition and search are so heterogeneous that it is essential to study them at the individual level.

CGCB's and CGC's focus on Occurrence and Adjacency follows naturally from a procedural view of decision-making and the empirical tendency, now confirmed by a large body of MouseLab data, of most subjects to perform the operations on hidden parameters on which their decisions are based via adjacent look-ups, relying on repeated look-ups rather than memory. In this view—perhaps too extreme—duration is unimportant because the information content of a look-up is independent of its length as long as it suffices for cognition; look-ups too short for comprehension, less than 0.18 seconds, were filtered out in the analyses discussed here. Although duration might still be correlated with time spent thinking about a particular parameter, which might be important in a more refined model of cognition, search, and decisions, a procedural view does not suggest such a correlation, and CGCB's and CGC's subjects sometimes left boxes open for long periods while staring out the window, etc., which would weaken any such correlations.²⁴ Total numbers of look-ups are important, but are captured indirectly through CGC's notion of search compliance.

4b. CGC's model of cognition and search

In CGC's model of cognition and search, each leading type implies a generically unique, pure adjusted guess in each game, which maximizes its expected payoff given the beliefs regarding others' guesses implicit in the type. (The leading types all specify best responses to some beliefs.) Each type is thereby naturally associated with algorithms that process hidden payoff information

²³ CGCB and CGC made no claim that durations are irrelevant, just that durations don't deserve priority. CGCB [Table IV] present some results on durations under the heading of "gaze times".

²⁴ Spezio, Wang, and Camerer's [2006] eye-tracking methods have an advantage in avoiding this ambiguity.

into decisions, which CGC used as models of cognition. Given the need to go beyond Occurrence and the lack of an accepted theory of cognition and search, the goal was to add enough restrictions to extract the signal from subjects' search sequences but not so many that they distort its meaning. CGC derived types' minimal search implications under conservative assumptions, based on Occurrence and Adjacency, about how cognition determines search [CGC 2006, Section I.B].

The leading role in these derivations is played by a type's *ideal guesses*, those that would be optimal given the type's beliefs, ignoring its limits. Given the quasiconcavity of CGC's payoff functions, a subject can enter his ideal guess and know that his adjusted guess will be optimal without checking his own limits. Thus, a type's ideal guess not only determines its adjusted guess and the resulting outcome, it also determines the type's minimal search implications.

The left-hand side of Table 1 [CGC 2006, Table 4] lists the formulas for the leading types' ideal guesses in CGC's games, which are easily derived as in CGC [2006, Section I.B], using CGC's notation for the limits and targets, a^i for the player's own lower limit, b^i for the player's own upper limit, and p^i for the player's own lower target, with analogous notation using superscript j 's for the player's partner's limits and target. The right-hand side of Table 1 lists the leading types' minimal search implications expressed as sequences of parameter look-ups, first in CGC's notation and then in terms of the associated box numbers (1 for a^i , 2 for p^i , 3 for b^i , 4 for a^j , 5 for p^j , 6 for b^j) in which MouseLab records subjects' look-up sequences in our design. Table 1 shows look-ups in the order that seems most natural; but that order is not required in the analysis.²⁵

The search implications are derived as follows. Evaluating a formula for a type's ideal guess requires a series of arithmetic *operations*, some of which—the innermost operations, whose parameters are in square brackets in the right-hand side of Table 1, like $[a^j, b^j]$ for $L1$ —are *basic* in that they logically precede other operations. Like CGCB, CGC assumed that subjects perform basic operations via adjacent look-ups, remembering their results, and otherwise relying on repeated look-ups rather than memory. Basic operations are then represented by adjacent look-ups that can appear in either order but cannot be separated by other look-ups. The look-ups of other

²⁵ In CGC's design, unlike in CGCB's, *Equilibrium*'s minimal search implications are simpler than any boundedly rational type's implications. This makes it harder to explain deviations from equilibrium by cognitive complexity. But we will see that high-compliance *Equilibrium* robot/trained subjects search more than high-compliance robot/trained subjects of other types, so CGC's *Equilibrium* search implications may not reflect its complexity.

operations can appear in any order and are (conservatively) allowed to be separated. In Table 1 such operations are represented by look-ups within curly brackets or parentheses.²⁶

An *L1* player *i*, for instance, best responds to the belief that player *j*'s guess is uniformly distributed between his limits. This yields a guess for *j* that is never adjusted, and that averages $[a^j+b^j]/2$. CGC (2006, Section I.B) shows via a certainty equivalence property of CGC's games (Observation 2) that *L1*'s ideal guess is $p^i[a^j+b^j]/2$, which will be automatically adjusted, if necessary, to $R(a^i, b^i; p^i[a^j+b^j]/2) \equiv \min\{b^i, \max\{a^i, p^i[a^j+b^j]/2\}\}$. The only basic operation is $[a^j+b^j]$. An *L1* player *i* therefore has minimal look-up sequence: $\{[a^j, b^j]$ (to compute *j*'s average guess), p^i (to identify *i*'s ideal guess) $\} \equiv \{[4, 6], 2\}$, of which $[4, 6]$ cannot be separated.

An *L2* player *i* best responds to the belief that player *j* is *L1*, taking the adjustment of *j*'s guess into account. An *L1* player *j*'s adjusted guess is $R(a^j, b^j; p^j[a^i+b^i]/2)$, so an *L2* player *i*'s ideal guess is $p^i R(a^j, b^j; p^j[a^i+b^i]/2)$, which will be automatically adjusted to $R(a^i, b^i; p^i R(a^j, b^j; p^j[a^i+b^i]/2))$. An *L2* player *i* therefore has look-up sequence: $\{([a^j, b^j], p^j)$ (to predict *j*'s *L1* ideal guess), a^j, b^j (to predict *j*'s *L1* adjusted guess), p^i (to identify *i*'s ideal guess) $\} = \{([1, 3], 5), 4, 6, 2\}$.²⁷ This illustrates the fact that CGC's design separates the search implications of different types as strongly as it separates their implications for guesses [CGC 2006, Figure 5].

In CGC's [2006] econometric analysis of search, not discussed here, search compliance for a given subject, type, and game is measured by the density of the type's complete minimal search sequence in the subject's look-up sequence for the game, allowing for the heterogeneity of search behavior.²⁸ CGC's measure is a significant advance on CGCB's measure, which is based on the percentages of a type's Occurrence and Adjacency requirements satisfied by the entire sequence.

²⁶ *L1*'s search implications illustrate an important advantage of the automatic adjustment feature of CGC's design. *L1*'s ideal guess depends on its own target but only its partner's limits, while *L2*'s and *DI*'s depend on both players' targets and limits and *Equilibrium*'s depends on both players' targets and a combination of its own and its partner's lower or upper limits. In other designs, such as CGCB's, *L1*'s decisions almost inevitably depend only on its own payoff parameters, and more sophisticated types' decisions depend on both own and other's parameters. Thus the automatic adjustment feature allows CGC to separate solipsism from the strategic naivete of *L1*. CGC's data give no evidence of solipsism, but a great deal of evidence of naivete. CGC's data also show that most subjects understood and relied upon automatic adjustment, which was carefully explained to them.

²⁷ With automatic adjustment, an *L2* player *i* does not need to know his own limits to play the game or think about the effects of his own guess being adjusted, only to predict *j*'s *L1* guess. By contrast, an *L1* player *i* doesn't need to know his own limits, only *j*'s. Because the possible values of the limits are not public knowledge, an *L2* player *i* cannot infer that adjustment of player *j*'s ideal guess can occur only at his upper (lower) limit when $p^j > 1$ ($p^j < 1$). An *L2* subject who incorrectly infers this may omit $a^j = 4$ ($b^j = 6$) when $p^j > 1$ ($p^j < 1$).

²⁸ As is evident from Tables 2 and 3, subjects' look-up sequences vary widely in what CGC called "style": Most robot/trained and baseline subjects with high exact compliance consistently look first at their type's minimal search sequence and then continue looking, apparently randomly, or stop and enter their guess (for example *L2*

4c. Sample search data

Table 2 gives a sample of the information search data for CGC's robot/trained subjects and Table 3 gives an analogous sample for baseline subjects of various assigned or apparent types, with Table 1's search implications repeated for convenience. In each case the subjects were chosen for high exact compliance with their types' guesses, not for compliance with any theory of search; subjects' frequencies of exact guesses are in parentheses after their types. Only the orders of look-ups are shown; and only from the first two or three games, but those games are representative.

Recalling that the theory allows any order of look-ups grouped within square or curly brackets or parentheses, the searches of high-guess-compliance robot/trained or baseline subjects conform closely to CGC's theory, with a subject's assigned or apparent type's minimal sequence unusually dense in his observed sequence.²⁹ The only exception is the *Equilibrium* subjects, who search far longer and in more complex patterns than CGC's theory suggests, perhaps because its minimal *Equilibrium* search requirements allow more luck than these subjects enjoyed.³⁰ Baseline *L1*, *L2*, and perhaps *L3* and *Equilibrium* subjects' searches are very close to those of their robot/trained counterparts, suggesting that (unlike in CJ) training had little effect on their search behavior.³¹ Perhaps *Equilibrium* search in normal-form games is less unnatural than backward-induction search in CJ's extensive-form games. For the simpler types *L1*, *L2*, and perhaps *L3*, the algorithms subjects use to identify their types' guesses can be directly read from their searches.

robot/trained subject 910, *L3* subject 1008, and *DI* subject 1501 in Table 2; and *L2* baseline subjects 108 and 206 in Table 3). But some such subjects look randomly first and turn to the relevant sequence at the end (*L1* robot/trained subject 904). CGC's [2006, Section II.E] econometric analysis uses a binary nuisance parameter to distinguish these "early" and "late" styles and filter them out to obtain a better measure of search compliance.

²⁹ CGC's specification analysis turned up only one clear violation of their proposed characterization of types' search implications, which is instructive. Baseline subject 415 (not shown in Table 3), whose apparent type was *L1* with 9 exact guesses, had zero *L1* search compliance in 9 of the 16 games because he had no adjacent $[a^i, b^j]$ pairs. His look-up sequences, however, were rich in (a^i, p^i, b^j) and (b^j, p^i, a^i) triples, in those orders; but not in such triples with other superscripts. This strongly suggests that 415 was an *L1* who happened to be more comfortable with three numbers in working memory than CGC's characterization of search assumes, or than their other *L1* subjects were. But because this violated CGC's assumptions on search, this subject was "officially" estimated to be *DI*.

³⁰ One of the methods CGC allow for identifying *Equilibrium* guesses is equilibrium checking, which has the least search requirements among all methods. Equilibrium checking can identify the *Equilibrium* guess very quickly if the player has the luck to check the equilibrium first [CGC 2006, Appendix H; CGC 2007]. Allowing this is unavoidable without risking incorrectly concluding that a subject has violated *Equilibrium*'s search implications.

³¹ CGC's baseline subjects with high compliance for some type are like robot/untrained subjects, which don't usually exist because one cannot tell robot subjects how they will be paid without training them in how the robot works. These "naturally occurring" baseline robot subjects provide an unusual opportunity to separate the effects of training and strategic uncertainty, by comparing their behavior with robot/trained subjects' behavior.

CGC's [2006, Section II.E, Table 7] econometric analysis shows that such inferences are usually consistent with estimates based on guesses alone, and that search compliance as measured here is also useful in identifying the types of subjects whose types are not apparent from their searches. For some subjects, econometric estimates based on guesses and search together resolve tensions between guesses-only and search-only estimates in favor of a type other than the guesses-only estimate. Those estimates confirm the presence of significant numbers of subjects of types *L1*, *L2*, *Equilibrium*, or hybrids of *L3* and/or *Equilibrium* in the population, and the absence of significant numbers of subjects of other types. Once again, subjects' deviations from equilibrium can be attributed mostly to non-equilibrium strategic thinking, not preferences or irrationality.

For some subjects search is an important check on type inferences based on guesses. Baseline subject 309, whose 16 exact *L2* guesses seem overwhelming evidence that his type is *L2*, violated *L2* Occurrence by missing one of its required look-ups in each of games 1-5 (Table 3 shows his look-ups for games 1-3). Just as for CJ's subjects who never looked at the second- or third-round pie, in games 1-5 this subject could not have been making *L2* guesses for the reason the theory assumes, and his compliance could not be expected to persist beyond sample. Fortunately, 309 had a Eureka! moment after game 5, and from then on complied almost perfectly with *L2*'s search requirements as well as its guess requirements.³²

5. Further Questions Search Analysis Might Answer

To illustrate some of the further possibilities for search analysis, this section discusses two questions raised by CGC's [2006] analysis of decisions that resist analysis via decisions alone. These questions will be addressed in CGC [2007].

5a. What are CGC's baseline apparent *Equilibrium* subjects really doing?

Figure 5 [CGC 2006, Figure 4] graphs the guesses of CGC's 8 baseline subjects with 7 or more exact *Equilibrium* guesses. The 16 games are ordered by strategic structure as in CGC [2006, Table 3] (not in the randomized order in which subjects played them), with the 8 games with mixed targets (one greater and one less than one) in the right half of the figure. Of these subjects' 128 guesses in the 16 games, 69 (54%) were exact *Equilibrium* guesses. In CGC's [2006]

³² Subject 309 omitted look-ups 4 and 6 (his partner's lower and upper limits) in game 1 and look-up 4 in games 2-5. This suggests that he did not yet understand the need to check his partner's lower limit to be sure of his *L2* guess even when his own target, or the product of targets, was greater than 1. However, he omitted look-up 4 even in game 4 where both targets were less than 1, showing that his error was probably more complex. That these omissions did not lead to non-*L2* guesses in games 1-5 is an accident of our design with no greater significance.

likelihood-based econometrics, given their a priori specification of possible types and the large strategy spaces of CGC's games, this is overwhelming evidence that their types are *Equilibrium*. But as Figure 5 makes clear, their *Equilibrium* compliance was far higher for games without mixed targets (55 out of 64 possible exact *Equilibrium* guesses, or 86%) than for games with mixed targets (14 out of 64, or 22%). Thus it is (even nonparametrically) clear that these subjects, despite *Equilibrium* compliance that is off the scale by normal standards, are actually following a rule that only mimics *Equilibrium*, and that only in games without mixed targets.

The puzzle is deepened by noting that all the ways game theorists teach people to identify equilibria (best-response dynamics, equilibrium checking, and iterated dominance) work equally well with and without mixed targets. Further, CGC's *Equilibrium* robot/trained subjects, who were taught these three ways to identify their equilibrium guesses, have roughly the same equilibrium compliance with and without mixed targets [Figure 6; CGC 2007]. Thus whatever the baseline apparent *Equilibrium* subjects were doing, it is not one of the first things a game theorist would think of. (Subjects' debriefing questionnaires did not reveal what it was.) Nonetheless, the rule or rules they follow have a decidedly non-random structure: All 44 of those subjects' deviations from *Equilibrium* (the solid line in Figure 5) when it is separated from *L3* (dotted line), with or without mixed targets, are in the direction of (and sometimes beyond) their *L3* guesses; though this could reflect the fact that in CGC's games, *L3* guesses are always less extreme than *Equilibrium* guesses.

CGC's [2007] analysis will try to resolve the puzzle by using the search data to answer the following questions:

- (1) How do the baseline apparent *Equilibrium* subjects find their equilibrium guesses in the games without mixed targets: best-response dynamics, equilibrium checking, iterated dominance, or something else that doesn't "work" with mixed targets? Refining CGC's (2006) characterization of *Equilibrium* search to separate the three methods and redoing the estimation with the refined compliance measures, separately for games with and without mixed targets, should be revealing. The absence of baseline *Dk* subjects suggests that iterated dominance, even finitely truncated, is unlikely. Best-response dynamics, perhaps truncated after 1-2 rounds, seems more likely.

(2) How do the baseline apparent *Equilibrium* subjects' search patterns differ in games with and without mixed targets? How do the differences compare to the differences for baseline *L1*, *L2*, or *L3* subjects? CGC's 20 apparent baseline *L1* subjects' compliance with *L1* guesses is almost the same with and without mixed targets [CGC, Figure 1], which is unsurprising because whether or not the targets are mixed is irrelevant to subjects who don't try to model others' responses to incentives. But the 12 apparent *L2* [Figure 4; CGC, Figure 2] and 3 apparent *L3* [CGC, Figure 3] subjects' compliance with their types' guesses is much lower with than without mixed targets. This is curious, because for *L2* and *L3*, unlike for *Equilibrium*, games with mixed targets require no deeper understanding.

(3) How do *Equilibrium* robot/trained subjects with high compliance find their equilibrium guesses even in games with mixed targets? How do their searches in those games differ from baseline apparent *Equilibrium* subjects' searches? CGC strove to make *Equilibrium* robot/trained subjects' training as neutral as possible; but something must come first, and they were taught equilibrium checking first, then best-response dynamics, then iterated dominance. To the extent that these subjects used one of those methods, it explains why they have equal compliance with and without mixed targets. But if some of them used something else that deviates from equilibrium mainly in games with mixed targets, it might provide important clues to what the baseline *Equilibrium* subjects did.

5b. Why are *Lk* the only types other than *Equilibrium* with nonnegligible frequencies?

CGC's [2006] analysis of decisions and search estimated significant numbers of subjects of types *L1*, *L2*, *Equilibrium*, or hybrids of *L3* and/or *Equilibrium*, and nothing else that does better than a random model of guesses for more than one subject. Why do these types predominate, out of the enormous number of possibilities? Why, for instance, are there no significant numbers of *Dk* types, which are closer to what game theorists teach?

CGC's [2007] analysis will try to answer this question by using search and other methods to look more deeply into the following phenomena:

(1) Most robot/trained subjects could reliably identify their type's guesses, even for types as difficult as *Equilibrium* or *D2*. Individual subjects' exact compliance with their type's guesses was usually bimodal within type, on very high and very low.

Even so, there are several signs of differences in difficulty across types:

(2) None of CGC's 70 robot/trained Lk subjects ever failed their type's Understanding Test, while 1/31 failed the $D1$ Test, 1/20 failed the $D2$ Test, and 7/36 failed the *Equilibrium* Test.

(3) For those who passed the test, compliance was highest for Lk types, then *Equilibrium*, then Dk types. This suggests that Dk is harder than *Equilibrium*, but more analysis is needed to tell if this was an artifact of the more stringent screening of the *Equilibrium* Test.

(4) Within the Lk and Dk type hierarchies, compliance was higher for lower k as one would expect, except that $L1$ compliance was lower than $L2$ or $L3$ compliance. This may be because $L1$ best responds to a random $L0$ robot, which some subjects think they can outguess; but $L2$ and $L3$ best respond to a deterministic $L1$ or $L2$ robot, which doesn't invite gambling.

(5) Remarkably, 7 of our 19 robot/trained $D1$ subjects who passed the $D1$ Understanding Test, in which $L2$ answers are wrong, then "morphed" into $L2$ s when making their guesses, significantly reducing their earnings (Figure 7 and subject 804 in Table 2; recall that $L2$ and $D1$ are cousins, both making 2-rationalizable guesses). This kind of morphing is the only kind that occurred, which seems compelling evidence that Dk types are unnatural. But a comparison of Lk 's and $Dk-1$'s search and storage requirements may have something to add.

6. A Rational-Choice Model of Optimal Search for Hidden Payoff Information

This section outlines a simple rational-choice analysis in support of the Occurrence and Adjacency assumptions that underlie CGCB's and CGC's models of cognition and search. The analysis is general in that it takes as given the formula that relates a type's decision to the hidden parameters. It views search for hidden payoff information as just another kind of rational decision, deriving subjects' demand for it from the benefits of making better decisions under plausible assumptions about the benefits and costs of search and storing numbers in working memory.

The model rests on two assumptions about cognition and search:

(1) The costs of look-ups are small. There is a great deal of evidence that subjects in experiments with hidden but freely accessible payoff parameters perceive the cost of looking them up as negligible, scarcely larger than the cost of reading them in a printed table. Having to look things up has small effects on their decisions (as shown in CGCB's and CGC's [2006] open boxes control treatments), subjects usually make many more look-ups than efficient search requires, and they usually make some motivated purely by curiosity.

(2) There is a flow cost of keeping numbers in working memory, which starts small for the first number but even then is larger than the cost of a look-up; and which increases with the number of stored numbers. Total memory cost is the time integral of the flow cost, and is therefore proportional, other things equal, to total storage time, and increasing in the number of stored numbers. (If working memory were free, nothing would prevent the scanning and memorization referred to in my discussions of CJ and CGCB, but this is plainly unrealistic.)

Occurrence follows immediately from assumption (1). A rational player looks up all costlessly available information that might affect his beliefs. When, as in these designs, information comes in discrete quanta with nonnegligible effects on beliefs and the optimal decision, this conclusion extends to information available at low cost.³³

Given Occurrence, Adjacency (in CGC's sense that the basic (innermost) operations in square brackets in the right-hand side of Table 1 are represented by adjacent look-ups) follows from assumption (2). Under (2), a player minimizes the total memory plus look-up cost by processing the basic operations needed to evaluate the expression for his ideal guess before other operations with whose results they are to be combined, storing the results (meanwhile "forgetting" the parameters they combine), and then combining them. Basic operations take precedence over other operations because "distributing" them increases memory cost.³⁴ For example, in evaluating

³³ Note that because MouseLab allows a subject to enter a tentative guess without confirming it (the *s in the data in Tables 2 and 3), thereby saving storage cost, the variations in search style noted in Section 4b, footnote 24, are consistent with optimality when look-up costs are negligible even if storage costs are not.

³⁴ This effect is related to the reason that backward induction is the most efficient way to solve a finite-horizon dynamic programming problem like those subjects faced in CJ's design: other ways are feasible, but wasteful of storage and computational capacity (though the latter is assumed to be freely available here).

the expression $p^i [a^i + b^i]/2$ for *LI*'s ideal guess, processing $[a^i + b^i]$ first, storing the result, and then combining it with p^i yield the following sequence of numbers of numbers in working memory: 1, 2, 1, 2, 1. The distributed alternative of processing $p^i a^i$, storing the result, then processing $p^i b^i$ and combining it with $p^i a^i$ yields the sequence: 1, 2, 1, 2, 3, 2, 1, which dominates the first sequence. The first method also saves the cost of looking up p^i a second time, but this is much less important.

Although Occurrence and Adjacency are only necessary conditions for optimal search, I stop with them because they have considerable empirical support, they make the main patterns of subjects' search behavior in *CJ*'s extensive-form and *CGCB*'s and *CGC*'s normal-form games intelligible, and they seem more transparent than other conditions for optimality and thus more likely to be descriptive of subjects' search behavior.

I close by noting that although this model supports *CJ*'s, *CGCB*'s, and *CGC*'s use of Occurrence and Adjacency, it says nothing directly about how to measure search compliance in an econometric analysis. *CGC*'s use of the density of a type's minimal search sequence in the part of the observed sequence where the subject tends to make his relevant look-ups (his search "style," in *CGC*'s terminology) is a judgment call, which seems to be well supported by inspecting the data.

7. Conclusion

CJ's, *CGCB*'s, and *CGC*'s analyses of cognition in games via monitoring subjects' searches for hidden but freely accessible payoff information bridge part of the gap between neuroeconomics and conventional economics because they rest on explicit models of cognition but search, unlike neural correlates of cognition, can be viewed as a rational choice. This paper has sought to use those analyses to make two points about the potential uses of neural data in economics.

First, standard assumptions of rational choice and equilibrium have yielded successful explanations of many phenomena, which as *GP* note can usefully be tested via revealed preference analysis of decision data. But there are other, equally important phenomena that appear to stem from failures of the implicit assumptions about cognition that underlie standard analyses, for which tests that don't take cognition explicitly into account are likely to be biased and misleading.

Second, with unbounded capacity for experimentation it might be possible to discover all we need to know about behavior by observing decisions alone. But this is an arbitrary constraint, and *CJ*'s, *CGCB*'s, and *CGC*'s analyses show that expanding the domain of analysis beyond decisions can yield a clearer view of behavior than is practically achievable by observing only decisions.

References

- Algaze (Croson), Rachel. 1990. "A Test of Presentation Effects on Strategy Choice," B. A. Honors Thesis, University of Pennsylvania.
- Bernheim, B. Douglas. 1984. "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007-1028.
- Camerer, Colin. 2007. "Neuroeconomics," this volume.
- Camerer, Colin, and Eric Johnson. 2004. "Thinking about Attention in Games: Backward and Forward Induction," in Isabel Brocas and Juan Carrillo (editors), *The Psychology of Economic Decisions, Volume Two: Reasons and Choices*, Oxford: Oxford University Press.
- Camerer, Colin, Eric Johnson, Talia Rymon, and Sankar Sen. 1993. "Cognition and Framing in Sequential Bargaining for Gains and Losses," 27-47 in Kenneth Binmore, Alan Kirman, and Piero Tani (editors), *Frontiers of Game Theory*. Cambridge: MIT Press.
- Caplin, Andrew. 2007. "Economic Theory and Psychological Data: Bridging the Divide," this volume.
- Costa-Gomes, Miguel, and Vincent Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review* 96, 1737-1768.
- Costa-Gomes, Miguel, and Vincent Crawford. 2007. "Studying Cognition via Information Search in Two-Person Guessing Game Experiments," manuscript in preparation.
- Costa-Gomes, Miguel, Vincent Crawford, and Bruno Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 69, 1193-1235.
- Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," *American Economic Review*, 96, 1043-1068.
- Gul, Faruk, and Wolfgang Pesendorfer. 2007. "The Case for Mindless Economics," this volume.
- Ho, Teck Hua, Colin Camerer, and Keith Weigelt. 1998. "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'," *American Economic Review*, 88, 947-969.
- Johnson, Eric, Colin Camerer, Sankar Sen, and Talia Rymon. 2002. "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining," *Journal of Economic Theory*, 104, 16-47.
- Koszegi, Botond, and Matthew Rabin. 2007. "Revealed Mistakes and Revealed Preferences," this volume.

- McKelvey, Richard, and Thomas Palfrey. 1995. "Quantal Response Equilibria for Normal-Form Games," *Games and Economic Behavior*, 10, 6-38.
- Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313-1326.
- Payne, John, James Bettman, and Eric Johnson. 1993. *The Adaptive Decision Maker*. Cambridge, U.K.: Cambridge University Press.
- Rubinstein, Ariel. 2007. "Instinctive and Cognitive Reasoning: A Study of Response Times," *Economic Journal*, 117, in press.
- Schotter, Andrew. 2007. "What's So Informative About Choices?," this volume.
- Stahl, Dale and Paul Wilson. 1994. "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior and Organization*, 25, 309-327.
- Stahl, Dale, and Paul Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10, 218-254.
- Wang, Joseph, Michael Spezio, and Colin Camerer. 2006. "Pinocchio's Pupil: Using Eyetracking and Pupil Dilation To Understand Truth-telling and Deception in Games," manuscript.




Figure 1. Display for CJ's Alternating-Offers Bargaining Experiments [CJ 2002, Figure 1]

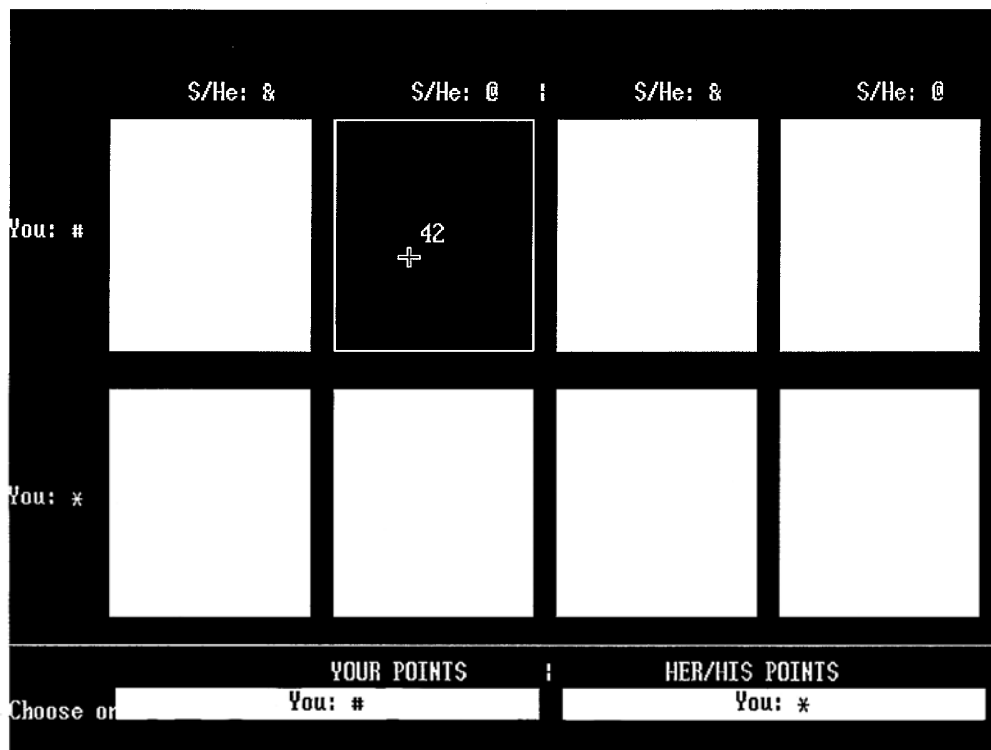
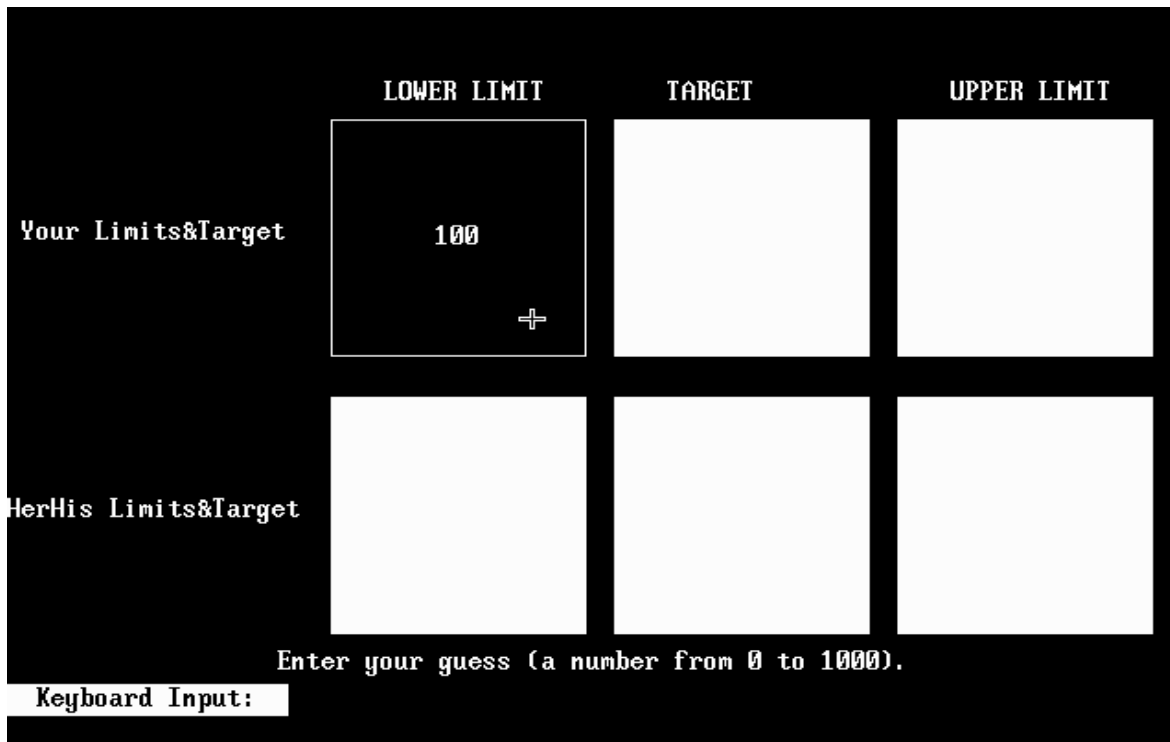


Figure 2. Display for a 2x2 Game in CGCB's Matrix-Game Experiments [CGCB, Figure 1]



Enter this box and click a mouse button when you are ready.

Figure 3. Display for CGC's Two-Person Guessing Games [CGC 2006, Figure 6]




Figure 5. “Fingerprints” of 8 Apparent *Equilibrium* Baseline Subjects [CGC 2006, Figure 4]

Note: Only deviations from *Equilibrium*'s guesses are shown.
 69 (54%) of these subjects' 128 guesses were exact *Equilibrium* guesses.




Figure 6. “Fingerprints” of 10 UCSD *Equilibrium* Robot/Trained Subjects [CGC 2007]

Note: Only deviations from *Equilibrium*'s guesses are shown.
 92 (58%) of these subjects' 160 guesses were exact *Equilibrium* guesses.




Figure 7. “Fingerprints” of 6 York Robot/Trained Subjects Who “Morphed” from *D1* to *L2* [CGC 2007]

Note: Only deviations from *D1*'s guesses are shown.

28 (29%) of these subjects' 96 guesses were exact *D1* guesses and 72 (75%) were exact *L2* guesses.

Type	Ideal guess	Minimal search sequence
<i>L1</i>	$p^i [a^i + b^i] / 2$	$\{[a^i, b^i], p^i\} \equiv \{[4, 6], 2\}$
<i>L2</i>	$p^i R(a^i, b^i; p^i [a^i + b^i] / 2)$	$\{([a^i, b^i], p^i), a^i, b^i, p^i\} \equiv \{([1, 3], 5), 4, 6, 2\}$
<i>L3</i>	$p^i R(a^i, b^i; p^i R(a^i, b^i; p^i [a^i + b^i] / 2))$	$\{([a^i, b^i], p^i), a^i, b^i, p^i\} \equiv \{([4, 6], 2), 1, 3, 5\}$
<i>D1</i>	$p^i (\max\{a^i, p^i a^i\} + \min\{p^i b^i, b^i\}) / 2$	$\{(a^i, [p^i, a^i]), (b^i, [p^i, b^i]), p^i\} \equiv \{(4, [5, 1]), (6, [5, 3]), 2\}$
<i>D2</i>	$p^i [\max\{\max\{a^i, p^i a^i\}, p^i \max\{a^i, p^i a^i\}\} + \min\{p^i \min\{p^i b^i, b^i\}, \min\{p^i b^i, b^i\}\}] / 2$	$\{(a^i, [p^i, a^i]), (b^i, [p^i, b^i]), (a^i, [p^i, a^i]), (b^i, [p^i, b^i]), p^i, p^i\} \equiv \{(1, [2, 4]), (3, [2, 6]), (4, [5, 1]), (6, [5, 3]), 5, 2\}$
<i>Eq.</i>	$p^i a^i$ if $p^i p^j < 1$ or $p^i b^i$ if $p^i p^j > 1$	$\{[p^i, p^j], a^i\} \equiv \{[2, 5], 4\}$ if $p^i p^j < 1$ or $\{[p^i, p^j], b^i\} \equiv \{[2, 5], 6\}$ if $p^i p^j > 1$
<i>Soph.</i>	[no closed-form expression, but CGC took <i>Soph.</i> 's search implications to be the same as <i>D2</i> 's]	$\{(a^i, [p^i, a^i]), (b^i, [p^i, b^i]), (a^i, [p^i, a^i]), (b^i, [p^i, b^i]), p^i, p^i\} \equiv \{(1, [2, 4]), (3, [2, 6]), (4, [5, 1]), (6, [5, 3]), 5, 2\}$

Table 1. Types' Ideal Guesses and Minimal Search Sequences [CGC 2006, Table 4]

Note: $R(a, b; x) \equiv \min\{b, \max\{a, x\}\} \equiv \max\{a, \min\{b, x\}\}$ denotes x 's adjusted guess with limits a and b .

Table 2. Selected Robot/Trained Subjects' Information Searches and Assigned Types' Search Implications

		MouseLab box number			Types' Search Implications															
		<i>a</i>	<i>p</i>	<i>b</i>																
You (<i>i</i>)		1	2	3	<i>L1</i>	{[4,6],2}														
S/he (<i>j</i>)		4	5	6	<i>L2</i>	{([1,3],5),4,6,2}														
					<i>L3</i>	{([4,6],2),1,3,5}														
					<i>D1</i>	{(4,[5,1], (6,[5,3]),2}														
					<i>D2</i>	{(1,[2,4]),(3,[2,6]),(4,[5,1]),(6,[5,3]),5,2}														
					<i>Eq</i>	{[2,5],4} if pr. tar.<1, {[2,5],6} if > 1														

Subject	904	1716	1807	1607	1811	2008	1001	1412	805	1601	804	1110	1202	704	1205	1408	2002
Type(#ex.)	L1 (16)	L1 (16)	L1 (16)	L2 (16)	L2 (16)	L2 (16)	L3 (16)	L3 (16)	D1 (16)	D1 (16)	D1 (3)	D2 (14)	D2 (15)	Eq (16)	Eq (16)	Eq (15)	Eq (16)
Alt.(#ex.)																	
Game																	
1	123456	146462	462513	135462	134446	111313	462135	146231	154356	254514	154346	135464	246466	123456	123456	123123	142536
	4623	134646		1313	5213*4	131313	21364*	564623	423213	36231	5213	2646*1	135464	363256	424652	456445	125365
		23		6	5423		246231	1	2642			313	641321	565365	562525	632132	253616
							52						342462	626365	6352*4	11	361454
													422646	652651	65		613451
													124625	452262			213452
													5*1224	6526			63
													654646				
2	123456	462462	462132	135461	134653	131313	462135	462462	514535	514653	515135	135134	123645	123456	123456	123456	143625
	4231	13	25	354621	125642	566622	642562	546231	615364	6213	365462	642163	132462	525123	244565	456123	361425
				3	313562	333	223146	546231	23		3	451463	426262	652625	565263	643524	142523
				52			2562*6					211136	241356	635256	212554	1	625656
							2					414262	462*13	262365	146662		3
												135362	524242	456	654251		
												14654	466135		44526		
												6	6462		31		

Notes: The subjects' frequencies of making their assigned types' (and when relevant, alternate types') exact guesses are in parentheses after the assigned type. A * in a subject's look-up sequence means that the subject entered a guess there without immediately confirming it.

Table 3. Selected Baseline Subjects' Information Searches and Apparent Types' Search Implications

	MouseLab box number		
	<i>a</i>	<i>P</i>	<i>b</i>
You (<i>i</i>)	1	2	3
S/he (<i>j</i>)	4	5	6

Types' Search Implications	
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

Subject	101	118	413	108	206	309	405	210	302	318	417	404	202	310	315
Type(#ex.)	L1 (15)	L1 (15)	L1 (14)	L2 (13)	L2 (15)	L2 (16)	L2 (16)	L3 (9)	L3 (7)	L1 (7)	Eq (8)	Eq (9)	Eq (8)	Eq (11)	Eq (11)
Alt.(#ex.)								Eq (9)	Eq (7)	D1 (5)	L3 (7)	L2 (6)	D2 (7)		
Alt.(#ex.)								D2 (8)			L2 (5)		L3 (7)		
Game															
1	146246 213	246134 626241 32*135	123456 545612 3463*	135642	533146 213	1352	144652 313312 546232 12512	123456 123456 213456 254213 654	221135 123456 213213 45456*	132456 465252 13242*	252531 464656 446531 641252	462135 464655 645515 21354*	123456 254613 621342 *525	123126 544121 565421 254362	213465 624163 564121 325466
2	46213	246262 2131	123564 62213*	135642 3	531462 31	135263 1526*2	132456 253156 456545 463123 156562 62	123456 123456 231654 456*2 54123	213546 566213 545463 21*266 54123	132465 132*46 2	255236 62*365 243563	462461 352524 261315 463562	123456 445613 255462 513565 23	123546 216326 231456 *62 3	134652 124653 656121 3
3	462*46 466413 *426	246242 264231 466413 *426	264231 53	135642 2231	535164 2231	135263	312456 5231*1 236545 5233** 513	123455 645612 3 563214 563214 523*65 4123	265413 232145 563214 563214 523*65 4123	134652 1323*4	521363 641526 5263*6 52	462135 215634 *52 3	123456 123562 3	123655 463213	132465 544163 *3625