Thinking about attention in games:

Backward and forward induction[1]

Colin Camerer
Division HSS 228-77, Caltech
camerer@hss.caltech.edu

Eric Johnson

Department of Marketing, Columbia University
ejj3@columbia.edu

1

Behavioral economics improves economic analysis by using psychological regularity to suggest limits on rationality and self-interest (e.g., Camerer and Loewenstein, 2003). Expressing these regularities in formal terms permits productive theorizing, suggests new experiments, can contribute to psychology, and can be used to shape economic policies which make normal people better off.

The crucial leap in behavioral economics is taking the details of thinking seriously, rather than settling for the conventional apology that models are only "as if" representations of the output of a neural "black box". Thinking seriously about cognitive detail, and formalizing those ideas, has become more fashionable in economics recently. Models of intertemporal choice seek to understand precisely how people think about the future (see Frederick, Loewenstein, and O'Donoghue, 2002). Theories of social preference formalize how people feel about allocations of payoffs to themselves and others (Fehr and Gachter, 2000; Camerer, 2003, chapter 2). Another major innovation in economics is the observation that people think differently about gains and losses, which implies that the way outcomes are described or "framed" can affect which choices are made (Rabin, 2000; Tversky & Kahneman, 1991).

Strategic interactions among many players, or "games", are of special interest for behavioral economics because even games which are very simple representations of naturally-occurring situations can give rise to equilibria that strain computational ability. As a result, games are a domain in which models of limited thinking are quite likely to provide better predictions than equilibrium theories. Indeed, behavioral game theory researchers have developed precise models of how deeply people are thinking about games which are as general as equilibrium theories, typically more accurate (and often more precise as well; Nagel, 1995; Stahl and Wilson, 1995; Costa-Gomes, Crawford, & Broseta, 2001; Camerer, Ho and Chong, 2003a,b; Camerer, 2003, chapter 5).

This chapter discusses how recording the information people seek in games can permit sharp inference about what goes on inside the black box (the brain) and can therefore inform theories about why people behave the way they do. Until recently, experimental economics was strongly influenced by the tradition of identifying unobserved (latent) utilities with choices, or revealed preferences. This tradition focuses attention on only one type of data from a decision--the observed choice. Looking only at choices is scientifically efficient because many other aspects of cognition *can* be measured (as we show below); and these other measures can be used to distinguish between alternative theories which can account for observed choices equally well. Indeed, as economists have become more interested in evolutionary and behavioral approaches, the number of theoretical alternative accounts has increased. The "process data" described in this chapter help by providing a richer set of data which allows us to sort through alternative explanations more quickly.

Measuring information search improves scientific efficiency in two ways:

- *Process data facilitates the evaluation of multiple theories.* In early experiments in game theory, the fundamental question was whether players chose equilibrium strategies (and when alternative equilibrium concepts made different predictions, which were most accurate). From this relatively well proscribed set of alternatives, a game theorist must now consider many more possibilities,

including concerns about fairness and limited cognition. Studying thinking allows us to explore this space of alternative explanations more efficiently by providing additional measures and constraints within the same experimental session. For example, in bargaining games fairness theories require that players look at the payoffs of others, even if those payoffs are strategically irrelevant. Theories of limited cognition can therefore be tested by looking at what information players **do not** examine.

- *Process data can help us understand heterogeneity among individuals*. While a single experiment might provide a dozen observed choices for an individual, augmenting that study with process data can greatly increase our ability to understand how individuals differ, when quite different processes lead to the same observed choices. As a result, we can identify which 'type' an individual is more quickly if we can watch the inputs to thinking. Understanding heterogeneity is important because it is well-established that how different types interact is important in determining system-wide behavior.[2]

This chapter gives two concrete examples of the advantages of using process data to study cognition in games, and also mention the caveats of using process data. We start by offering a quick overview of process measures, then provide a quick review of several studies that have applied process data to games. Finally, we close by discussing the future of process data in economics.

Visiting the Thinking Factory.

There are many ways to study the psychological processes underlying cognition. We lump different methods together under the term "process data", meaning the flow of information used in service to cognition. All the methods share one characteristic: They observe how people acquire information. By observing the arrival and use of information, we can infer something about the nature of underlying cognition. Note well that we are <u>not</u> talking about peoples' self-reports about their own thinking. While self-reports are sometimes insightful, there are well-known problems with introspective access to cognition which limit what can be learned from just asking people what they were thinking (e.g., Nisbett and Wilson, 1977).

Johnson et al (2002) introduced the analogy of the human brain as a "thinking factory." One can make educated guesses about a factory's production process by observing the flow of inputs into the factory, the length of processing time before new inputs enter the factory, and the factory's final output without ever looking inside the factory. Lawyers studying jury decisions (whose deliberations cannot be directly observed) and intelligence agencies make inferences like this routinely. Similarly, the order in which subjects gather information, and how long they use information before getting more information, can be used to make educated guesses about how subjects are thinking (i.e., the production process in the neural "thinking factory"). If information is not acquired, it cannot be used, just as a car factory which never receives shipments of rustproofing chemicals cannot possibly be rustproofing its cars.

Comment: Colin: I might prefer process data, what do you think? SWITCHED

---

[2] The most familiar examples is models of repeated games, in which a small percentage of types with "unusual" behavior can profoundly affect what other players do.

An analogy can be made to revealed preference. Choices between objects (such as commodity bundles or gambles) reveal unobservable preferences. Similarly, studies of industrial output reveal unobserved production functions. Asking subjects to `choose' information is asking them to reveal their preference for information, which can provide clues to unobservable thinking patterns.

Psychologists and consumer researchers have employed many process tracing techniques, including recording people 'talking aloud' while thinking (Ericsson & Simon, 1980), watching the physical retrieval of information required to solve a problem (Payne, 1976 and Jacoby et al, 1985), or asking people what information they would use (Wason, Johnson-Laird).

 The basic idea behind all process tracing techniques is illustrated by recording eye fixations, because it underlies most visual information acquisition   It surprises most people to learn that our visual world is, in part, illusory. While we believe that we see a continuous stream of visual input, the eye is more like a still camera than a movie camera. For almost all vision that requires significant acuity, the eye makes a series of jumps or *saccades* between pauses, called *fixations*.   During the 100 milliseconds or so required for a saccadic movement, vision is actually suppressed. One tends not see a brief light that is flashed during an eye movement. It is the brain that takes the sensory output of these discrete snapshots each lasting around .2 to 5 seconds and creates the perception that we see a continuous, uninterrupted view.   The beauty of eye fixation recording is that it is exactly what goes on in the real world: All high level cognition, such as reading, uses eye fixations, and eye fixations have been quite useful in studying reading and similar tasks (Just & Carpenter, 1976). Eye fixations have also been used to study choice (Russo & Rosen, 1974). But recording eye fixations (typically by filming the direction of eye gaze relative to stimuli spaced sufficiently apart on a screen) is expensive and the method is somewhat uncomfortable. It is also not easily adapted to group settings.

Many researchers now employ computer-based information acquisition as an analog for recording eye fixations. These methods display information hidden in boxes on a computer screen and observe a person as they "acquire" or "look up" information using some kind of pointing device like a mouse. People navigate through a task by moving the mouse and clicking on boxes which reveal information they want to know.

Tracking "clicks" is now a major activity for internet web sites, but computer-based information acquisition has a long history that predates the web. It has been used extensively to study individual decision-making and consumer choice.   Studies of decision making include the choice of gambles (Payne & Braunstein, 1978), the setting of reservation prices for gambles using different response modes (Johnson and Schkade), the effects of time pressure on choice ((Payne, Bettman, & Johnson, 1993; Reiskamp and Hoffrage, 2000) and the study of preference reversals. (Schkade & Johnson, 1989). Studies of consumer choice includeJacoby, Jaccard, Kuss, Troutman, & Mazursk (1985).

Computerized methods offer several advantages over verbal self-reports, eye movements, having subjects request information displayed on file cards, and other older methods. Using a mouse reduces the impact of the measurement procedure upon the underlying process. Computer-based methods also easily adapted to group situations, unlike verbal protocols and eye movements. The recording of information acquisition is automatic, so experiments can be run on networked computers without obtrusive intervention by experimenters. An important question is whether recording process data in these ways changes the process (a la the Heisenberg uncertainty principle in physics—does the act of experimental

4

observation change the process being observed?). The answer is generally "No" but we defer further discussion to the end of this chapter.

Thinking Backwards

Consider an alternating offer sequential bargaining game (Stahl, 1972; Rubinstein, 1982). This study (Camerer et al, 1993; Johnson, et al 2002) illustrates the use of process data to (1) distinguish among alternative theories and (2) to examine individual differences.

In this game, two players, 1 and 2, bargain over a pie that shrinks in value in each of three periods (reflecting discounting due to impatience). The pie is worth about $5 in the first period, $2.50 in the second period, $1.25 in the third period, and nothing after that. Starting with player 1, the players alternate making offers, which the responding player can accept or reject. If an offer is rejected, the pie shrinks and the player who rejected it then makes a counteroffer.

Play is not well predicted by equilibrium: If players are purely self-interested (and believe others are too[3]), perfect equilibrium divisions can be derived by backward induction, starting with the third-period division and working backward. The perfect equilibrium is (approximately[4]) for player 1 to offer $1.25 and keep the rest for himself. In experiments, subjects typically offered something between the $1.25 equilibrium and $2.50 (which is an equal split of the first-round pie). Offers average around $2.00. Lower offers, including offers near the equilibrium prediction of $1.25, are frequently rejected.

At least two theories offer possible explanations. The first, limited computation (LC), explain observed departures from perfect equilibrium by suggesting that players do not reason game-theoretically, and hence do not initially understand how the structure of the game conveys bargaining power. The second class of theories uses equilibrium social preferences to account for departures from perfect equilibrium. They posit a "social utility" for others' payoffs or differences in payoffs (Loewenstein, Thompson and Bazerman (1989); Fehr and Schmidt (1999); Bolton and Ockenfels (2000); Charness and Rabin (2002); Camerer (2003), chapter 2), a preference for reciprocating fairness and unfairness (Rabin (1993)), or an unobserved component of payoffs which appears like noise or a mistake to an outside observer (McKelvey and Palfrey (1995, 1998); Goeree and Holt (2001); Weizsacker (in press)).

Of course, these two theories are not mutually exclusive, and people may well differ in the degree they believe in fairness or have limited computational abilities. Clean predictions

---

[3] Of course, the perfect equilibrium prediction in an n-stage game requires n levels of iterated belief in self-interest. For example, in our game player 1 offers $1.25 only if she is self-interested, believes 2 is (which guarantees a minimal ultimatum offer in the third round), and believes 2 believes she (1) is (which guarantees a $1.25 offer to 1 in round 2, leading player 1 to offer $1.26 in round 1),

[4] In the third round, player 1 offers $.01 or nothing (depending on whether player 2 will accept nothing) and earns $1.24-$1.25. In the second round, player 2 must offer $1.24-$1.26 to player 1, leaving $1.24-$1.26 to herself. In the first round, player 1 must therefore offer $1.24-$1.27 to player 2. For simplicity, we refer to this range of perfect equilibria as "the" equilibrium at $1.25, recognizing that the equilibrium is not unique, and is subgame perfect rather than only Nash (every offer is consistent with Nash). Note that a short cut to calculating equilibria in games of this form is that the first player earns the sum of the pie sizes in odd-numbered rounds ($5.00 + $1.25), minus the pies values in even-numbered rounds (- $2.50). Alternatively, the first player should offer the third pie minus the second pie to the second player, which does not even require her to know the first-period pie size.

based on choices alone are ambiguous. Table 1 shows four theories which combine elements of fairness and limited computation. We describe each combination briefly.

        If players care only about their own payoffs, and can compute perfect equilibrium, then GT predicts is that players will offer $1.25. Now suppose players can compute perfect equilibrium, but they either care about fairness or believe others do. This combination of "game theoretic with fairness" (GT-F). suggests that offers will lie in an interval between the GT prediction and the equal-split point, [$1.25,$2.50]. If players may do not reason game-theoretically and not initially understand how the structure of the game creates bargaining power. Such a "limited computation" (LC) theory, predicted offers could lie anywhere; if we assume an upper bound at the equal split point $2.50, we predict an offer p in the range [0,2.50]. If players have limited computation <u>and</u> are concerned with fairness, we predict offers will be more generous than in the LC case, $p_1 \geq p$. Note that under these conditions, only the GT prediction seems inconsistent with the existing data, and our results.

Table 1: Model Predictions

Theory	Treatments: Bargaining with...		
	...other subjects.		....robots and instruction.
GT	$1.25	$1.25	$1.25
GT-Fairness	[1.25,2.50]	1.25	1.25
LC	p in [0,2.50]	p	1.25
LC-Fairness	$p_1 \geq p$	$p_2 < p_1$	1.25
Mean offer	$2.11	$1.84	$1.22

        Clearly additional help might be useful in understanding the relative role of these two accounts. In Johnson et al (2002) we gathered process tracing data using the MouseLab computer-based information monitoring system (Johnson, Payne, Schkade, & /Bettman, 1991). This system displayed the parameters of the problem behind the labeled boxes.   Players acquired the information by moving the cursor into the box.   When the cursor left the box, it reverted to a blank square.

6

Figure 1: Computer Display, Alternating Offer Game

Software records the entrance and exit time for each box, typically to the 1/60[th] of a second.   Typically these data are characterized by a dense series of acquisitions with durations
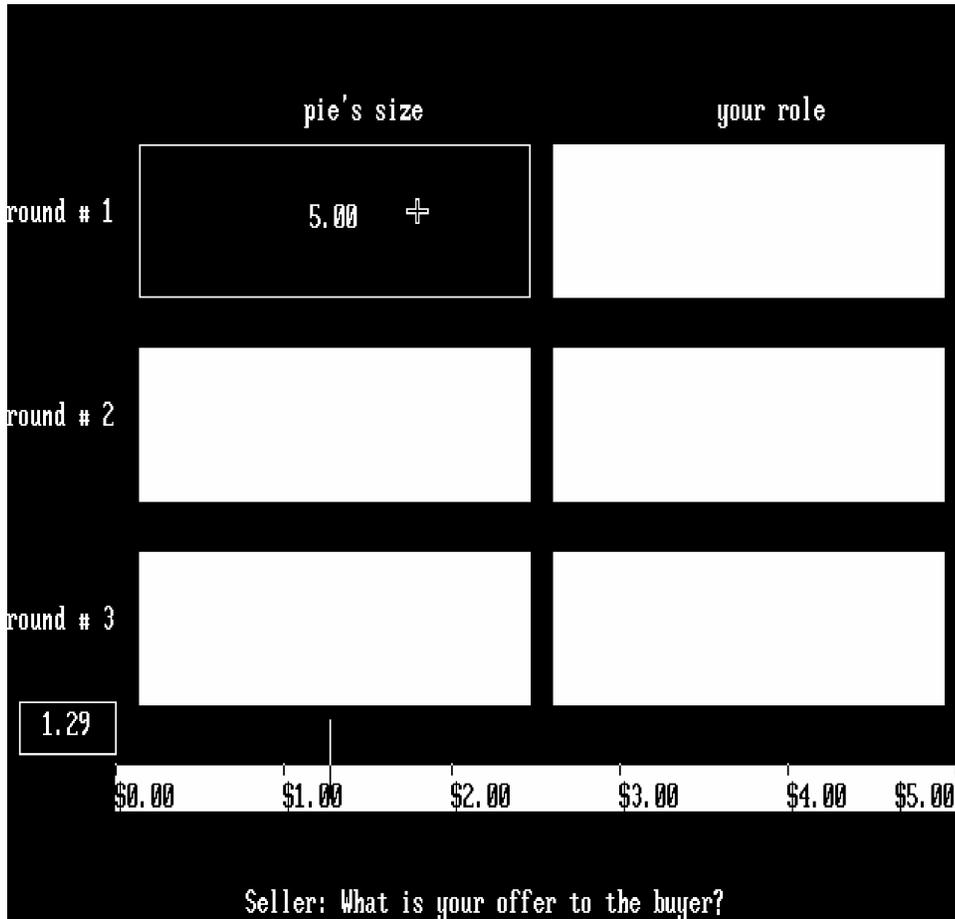

much like those generated by eye fixation recording, ranging in length from 200 milliseconds to a few seconds in length. Consistent with the notion that this is a relatively effortless and costless response, there are many occasions when subjects open the same box more than once, showing that it is easier to move the mouse than memorize something.[5]   Figure 2 shows the data, both in terms of number of acquisitions and in terms of the distribution of total time.   Figure 2 also

---

[5] There are exceptions. We had the opportunity to demonstrate this design to a well known game theorist, who used the mouse (with two hands) and wrote the numbers down on a piece of paper.   Fortunately for us, such behavior is exceedingly rare: To date it has been confined to Nobel Prize winners.

demonstrates that using process tracing data represents a challenge in data presentation: How can we show mean looking times and numbers of look-ups, as well as transitions from cell to cell? The "icon graph" at the left of Figure 2 is one method to show all these data: For each box (i.e., each amount to be divided in each round), the width of the bars is proportional to the number of acquisitions, the height is proportional to the total looking time, and the black arrows to the left of the figure show the frequency of transitions (bolder = more frequent, with fewer than one transition on average omitted).[6]

Figure 2: Distribution of Acquisitions, and Transitions by Round.
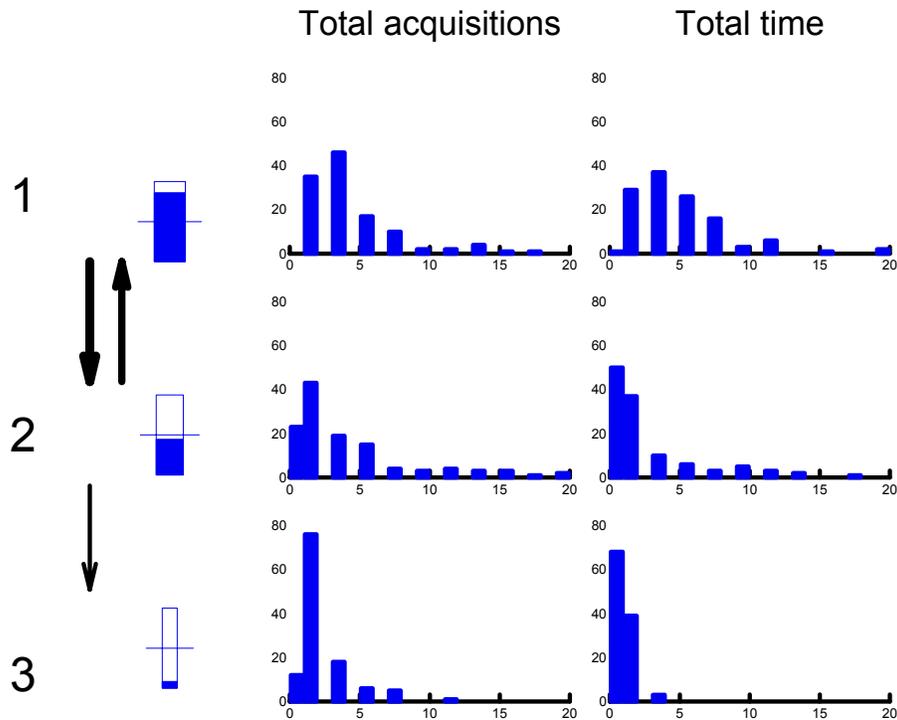
Like most studies, we found a fairly broad range of offers in the first round from around the equilibrium of $1.25 to some just above the midpoint of $2.50. Based upon choice data alone, we were unable to distinguish between theories. The process tracing data, shown in Figure 2 seems more compatible with a limited cognition view: Many subjects do not look at the

---

[6] The skewness of these data is typical. For analysis, it is often useful to either use nonparametric methods, which often lack statistical power, or to analyse the data with parametric techniques after a log transform $x' - log(x+c)$ where c is a small constant, typically .5, used to prevent taking the log of 0. Means are reported after an inverse transformation.

subsequent rounds: 19% do not open the second box, and 10% do not open the third. Both pieces of information would be necessary to calculate the equilibrium.

Yet there seems to be great heterogeneity in offers and search. Without process data, the range of offers might be duly noted and remain unexplained. To help us understand this heterogeneity, however we posited three types of information acquisition, roughly corresponding to levels of look ahead: Level 0 considered only the current payoff in Round 1, Level 1 looked at this round payoff and that in the next round, while Equilibrium players must consider all three rounds.
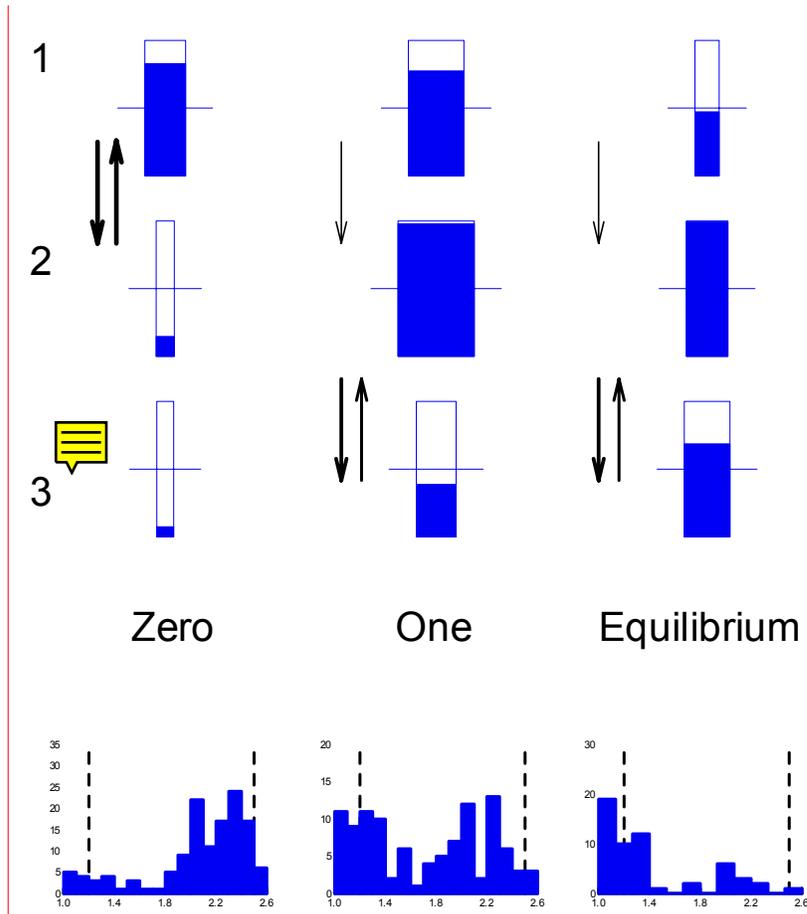

Figure 3: Search and Offers by Type

Types are classified by identifying which box was looked at longest in the first bargaining round. This is simple, and allows us to test if the offers are indeed related to search. Figure 3 displays the results for the process tracing data and offers, aggregating across our first two studies. The top of the figure displays an icon graph that is very much as predicted: Those

9

trials which are characterized as exhibiting "Level 0 look ahead" spend most of their time with the first round payoff, spend very little time looking at other payoffs, and make very few comparisons between subsequent payoffs. Their mean offers are $2.07. Level 1 cognition is characterized by increased time spent on the second round payoff, and have mean offers of $1.71. Finally, those trials which are characterized by the most acquisitions of the last round payoff, look much like we expect from backward induction: Minimal examination of the first round payoff, concentration on the second and third round payoff with many comparisons, and a mean offer of $1.44. While there is substantial dispersion in the offers produced within each information-lookup type category, the differences in the offer distributions across the three columns show that it is possible to statistically predict what a player will offer from what information they look at most often. The big loser among the alternative theories is the idea that players make nonequilibrium offers because they do equilibrium computations but adjust for fairness concerns of others ("GT-fairness"), since the most fair offers come from subjects who simply do not look ahead. Furthermore, in one experimental condition subjects bargain against a computerized opponent who is programmed to maximize its own payoff (and to expect that other subjects will too). In that condition, must subjects do not offer $1.25, which they would if they were computing money-maximizing offers optimally; although they learn to do so quite rapidly with a little instruction in backward induction.

This study shows that process data allows us to both (1) increase the set of strategies considered (to include steps of thinking) and (2) provide a better account of heterogeneity in these data.

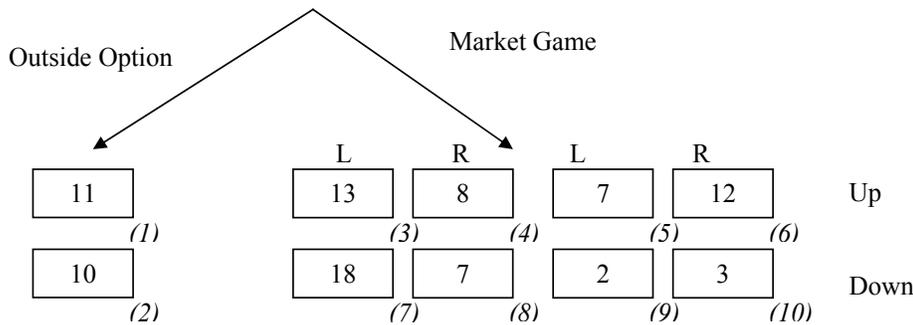## Thinking Forwards

Figure 4: Forward Induction Game

Figure 4 shows a game we used to study forward induction. Player 1 can choose the outside option and guarantee a payoff of 11, or play a simultaneous-move "battle of the sexes" subgame, choosing up or down. What would you do?

A rational Player A plays the market subgame and then plays Down,[7] expecting player 2 to reason as follows:

> Player 1 could have had a payoff of 11. Instead, she entered the game, signaling that she intends to get more than 11. The only equilibrium which yields more is (Down, Left). Since I have deduced that she will move Down I should move Left.

This clever argument is called "forward induction" (Kohlberg & Mertens, 1986; van Damme, 1989): choices at current nodes have implications for play in future nodes. To do forward induction, player 2 must look player 1's choices, ones which will not be played and which are not choices for player 2. Player 2 must reason that those payoffs may offer clues about what player 1 intends to do.

We ran three sessions of games similar to that in Figure 4, with 14 subjects each. (The games differed in Player A's option payoff, either 14 or 11, and Player B's option payoff, either 6, 10, or 16.) Contrary to forward induction, player 1 subjects chose the outside option more than half the time (55%) of the time. (In a structurally similar game, Schotter et al, 1994, observed 70% choice of the outside option; Camerer, 2003, chapters 5 and 7, summarizes other studies showing weak support for forward induction). When Player A's did choose to play the matrix game, they chose the bottom row 95% of the time; so the Player A's who anticipated that Player B's would do forward induction were betting heavily by choosing the bottom row almost all the time. In addition, B's played left about 75% of the time. Furthermore, there was little change across twelve repetitions, indicating no detectable learning.

Process data provide a partial answer to why the forward induction was rare. Figure 5 presents an icon graph of Player B's looking time. Note that there are very few lookups of acquisition of either outside option payoff— the box S11 icon is skinny and the associated looking time is low (little shaded area)-- although player B has to look at player A's option payoff (box S11) to make a forward induction inference. But note that Player B's do not think that Player A's playoff are irrelevant; in fact, they spend considerable time looking at Player A's payoffs in the matrix subgame. It is just the outside option payoffs of both players that are neglected by Player B's. This is particularly interesting because both players alternated roles. It is also interesting because Player A's **are** affected by the outside option.

Comparing behavior when payoffs change also helps clarify what players are doing. Recall that Player A's option payoff is either 11 or 14, in different treatments. In both cases the option payoff is above the "bad" pure equilibrum payoff for Player A (which is 7) and also above the expected mixed-strategy equilibrium payoff (which is 10). Forward induction point to the "market, (bottom,left)" equilibrium in both cases. But when the option payoff is 14, the "market, (bottom, left)" equilibrium can also be reached by three steps of deletion of dominated strategies: If B's think that A's will never violate dominance, they think the A's will never play "market, top"; eliminating that dominated strategy, B's will always choose "left"; if A's believe B's will play that strategy, A's will play "market, bottom". Thus, if A's think that B's think that A's will

---

[7.] More formally: The subgame has two pure-strategy Nash equilibria, (T,r) and (B,l), and a mixed-strategy equilibrium (.5T+.5B,.5r+.5l) (with expected utilities (10,7.5)). The entire game has Nash equilibria in which player 1 moves L because she thinks subgame play will give her less than 14. Of those, the equilibria (LT,r) and (L(.5T+.5b),.5r+.5l) are subgame perfect too. The perfect, Nash equilibrium (RB,l) is the unique "stable" equilibrium which satisfies forward induction.
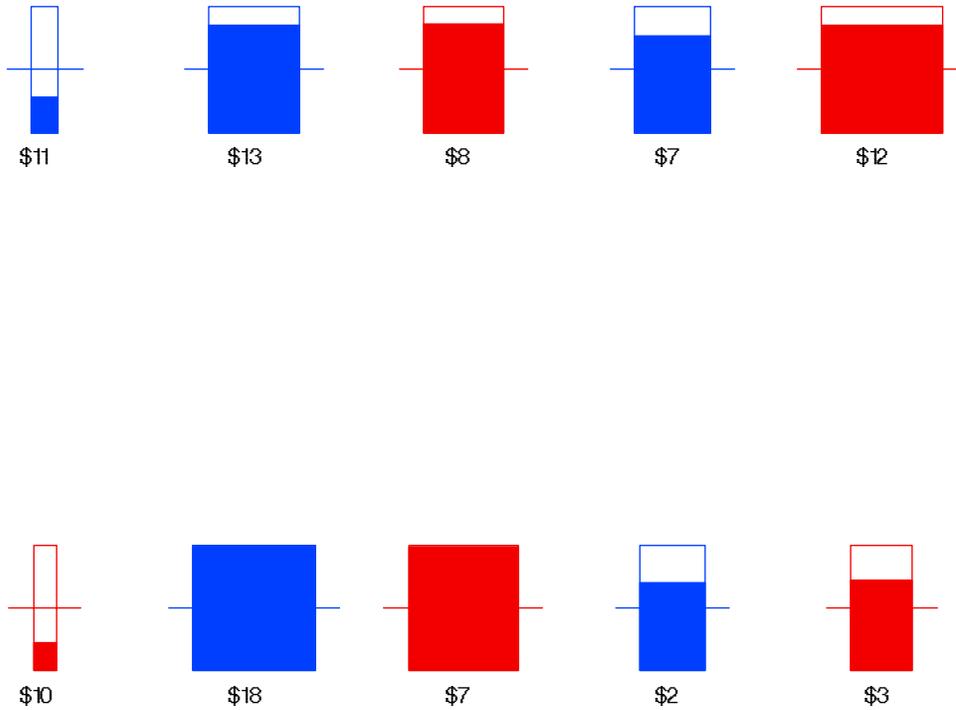
Figure 5: Icon Graph, Player B

obey dominance, A's should play "market, bottom". If B's think A's will obey dominance the B's will play "left".

Thus, when A's option payoff is 14, applying dominance iteratively leads to the "market, (bottom, left)" outcome even without forward induction reasoning. One might therefore expect that the forward induction outcome is more frequent when A's outside option payoff is higher, 14 rather than 11 (since both forward induction and iterated reasoning predict that outcome when A's option payoff is 14). Intuitively, in rational analyses the higher option payoff for A gives the B players more reason to believe that A will play the market game, and more reason—if they are iteratedly rational—to infer that they should play "left", which in turn draws strategically-minded A's into the market game. In fact, the opposite is true: When the option payoff was 14, Players A entered the market game only 38% of the time; when the option payoff was 11 they entered the market game 57% of the time. So changing A's option payoff had the opposite effect of what a rational analysis would predict.

The fact that more Player A's chose the option when its payoff was higher (14 rather than 11) is consistent with a model in which players are limited in their strategic thinking, but are not

12

sure other are rational (e.g., Camerer, Ho and Chong, 2003b). Crawford (2003) put a similar analysis to work in a careful study of deception.

About ten years ago, we created some behavioral game theory around this concept—especially the idea that players treated choices of other players as gambles, rather than as the result of a delicate (iterated) reasoning process. At that time, we did not know precisely how to model this idea, and rationality-based game theory was so roaringly productive that buyers for a cognitively-plausible alternative were scarce. It is heartening that theorists eventually accepted the mathematical challenge of modelling the limits on rationality that we envisioned, and are now relatively hungry for data to guide their theorizing.

Thinking Normally

The studies above explore whether players look backward and forward in "extensive-form" games that take place in several stages.. Inspired by our earlier studies, Costa-Gomes, Crawford and Broseta (2001,2002) used Mouselab to answer a more basic question in game theory: How strategically do players think in "normal form" (matrix) games?

Their ambitious study used 18 different games, with from 2 to 4 strategies for each player. The games were carefully designed so that players using different decision rules would pick different strategies, and different rules imply different patterns of information acquisition. Assuming each player has a single strategy or "type", they then use a sophisticated Bayesian procedure (cf. Harless and Camerer, 1994; El-Gamal and Grether, 1995;) to infer which type a player is most likely to be, given the player's choices and information acquisition across the 18 games.

Some of the decision rules they consider do not use sophisticated strategic thinking about what others will do. These rules are: "naïve" or level 1 (L1) (choose strategies with the highest average payoff, averaging all payoffs equally; "optimistic" or maximax (choose the strategy with the highest possible payoff); "pessimistic" or maximin (choose the strategy with the highest possible minimal payoff); and "altruistic" (maximize the sum of the two players' payoffs).

There are also five strategic types: "L2" (best-respond to L1); "D1" which does one round of deleting dominated decisions, then best-responds to a uniform prior on the remaining decisions; "D2", which does two rounds of deletion then best-responds to remaining strategies; "sophisticated" who guess accurately the proportion of the population that chooses each nonequilbrium strategy[8]; and "Nash" (equilibrium).

As in our work, subjects participated in baseline and "open boxes" conditions, to see whether hiding payoffs behind boxes limited strategic thinking (it did, but only a little). They also trained subjects to use the various decision rules (and rewarded them for picking the strategies those decision rules should pick) to calibrate what information acquisition looks like for various decision rules.

The percentages of untrained subjects choosing equilibria were 90%, 65%, and 15% when equilibria require 1, 2, and 3 levels of iterated dominance, respectively, which is consistent with much other evidence that people do only one or two steps of strategic reasoning (e.g.,

---

[8] Importantly, in the first version of their paper the L2, D1 and D2 types were not included in the analysis, and their procedure classified nearly half of the subjects as "sophisticated". When these types are included the fraction of estimated "sophisticates" falls almost to zero, which shows that specifying a rich space of decision rules is important to draw the right conclusions.

Nagel, 1999; Camerer, 2003, chapter 5; Camerer, Ho and Chong, 2003). However, trained subjects find equilibria almost perfectly (90-100% of the time). Like our finding that human subjects exposed to backward induction quickly make equilibrium offers to robot subjects, the powerful effects of training show that strategic thinking is not "computationally difficult" per se; it just requires decision rules which are not unintuitive, but easily taught. Strategic thinking is not like weight lifting or dunking a basketball, where performance is constrained by physical limits. Instead, strategic thinking is more like learning to windsurf or ski, which require people to learn skills which are unnatural and overcome natural instincts (overcoming the reluctance to lean forward going downhill and backward toward the water). Learning backward induction requires overcoming the intuition that future subgames which are unlikely to be played can be ignored; and learning iterated reasoning requires looking as much at another player's payoffs as one's own.

Since different decision rules point to different strategies in each game, by looking at a player's strategy choices the player can be classified according to which decision rule they use most often (cf. Harless and Camerer 1994, 1995; El-Gamal and Grether, 1995).

Table 2 shows the fraction of subjects estimated to be using each of the decision rules, when the estimation is doing using either the decisions alone, *or* the decisions plus information search data. The table makes three important points:

First, about 90% of the subjects appear to use simple rules (L1, L2) which are only minimally strategic (i.e., a player A does not assume that B is reasoning about A). Hardly any consistently choose equilibrium.

Second, if CGCB had done their study using only decision data, without process data, they would have drawn the wrong conclusion. Here's why: When only decisions are used to classify players, about 20% of players are assumed to use each of the decision rules L1 and D1. But when information search is used as well, the 19% originally classified as D1 shrinks to zero,

Table 2: Estimated frequencies of various decision rules and their expected payoffs (Costa-Gomes et al, 2001)

Decision rule	Information used to estimate frequencies		expected payoff
	Decisions	Decisions and search	
Altruistic	.089	.022	17.11
Pessimistic	.000	.045	20.93
Naïve (L1)	.227	.448	21.38
Optimistic	.000	.022	21.38
L2	.442	.441	24.87
D1	.195	.000	24.13
D2	.000	.000	23.95
Equilibrium	.052	.000	24.19
Sophisticated	.000	.022	24.93

and the percentage classified as L1 roughly doubles. The reason for this is that the D1 and L1 types often make the same decisions, but the D1 types must look at other players' payoffs (to see whether other players have dominated strategies that can logically be eliminated), which L1

14

types do not. Most of these players look at others' payoffs rarely, so the process data show it is statistically much more likely that they are L1's than D1's (despite the fact that the predicted choices of those types often coincide). The process data therefore overturns an incorrect conclusion drawn from the decision-only data.

Third, there appears to be a sensible cost-benefit tradeoff between complexity of rules and their payoffs which explains why most subjects use simple rules (cf. Johnson and Payne, 1985). The second column of Table 2 shows the expected dollar payoff to subjects using each type of rule (if they were randomly matched with every other subject, given how subjects actually behaved). The L1 and optimistic rules make a little more than $21 and don't require looking at the other player's payoffs at all. The L2 and D* rules earn about 10% more, about $24, and L2 actually earns <u>more</u> than equilibrium! This might seem anomalous since it is easy to confuse "equilibrium" and "optimal". But equilibrium strategies are only universally optimal— i.e., better than any other strategy—if they are dominant strategies. But generally, an equilibrium strategy can be a bad response to disequilibrium moves by others, and simpler rules like L2 appear to be more robust. Note also that since the sophisticated types' best respond to the actual mixture of decision rules being used, by definition they have the highest expected earnings. However, they earn only <u>$.06</u> less than the L2's, which means that being fully sophisticated is only slightly more profitable than doing two steps of strategic thinking.

Conclusions

In this chapter we presented three examples of how process data can be used to understand how people play sequential and matrix games. In all three studies, motivated, intelligent subjects behave sensibly, but do not exhibit the extent of strategic reasoning which is commonly assumed when game theory is applied to understand auctions, industrial pricing, political maneuvering, incentive design, and so forth.

In the sequential bargaining studies, subjects rarely look ahead to future nodes which they believe (correctly) are unlikely to be reached. In the forward induction game, most second-mover players do not look back at the first-mover's possible payoffs (even though she might use the foregone payoff to infer something about what the first-mover will hope to get in the future, and hence what she might do). In the matrix games, most players appear to do one or two steps of strategic thinking.

Caveats

It is sensible to wonder whether the experimental procedure measuring process, by requiring a player to move a mouse to open boxes on a screen display, biases the results in some way. These concerns have been raised many times and addressed by further experiments:

Does process tracing change the underlying process, is it reactive? Within decision research, there is a significant literature replicating many standard phenomena established with pencil-and-paper methods, such as preference reversals, using computer based techniques (Payne et al., 1993), so the method of measurement does not change the observed result in these cases. Researchers have also directly compared behavior in games when players *do not* have to open

boxes, because they appear on the screen uncovered, with Mouselab trials where a mouse is used to open boxes. Costa-Gomes et al. (2001) presented one session of their normal form games without covering up the payoffs and noted no difference in observed choices. In our work on sequential bargaining, we also ran a control session with open boxes, and the offers were within $.01 of offers when boxes were closed and a mouse was used. Lohse and Johnson (1996) have compared eye fixation recording with one particular computer-based system, MOUSELAB and found only small differences (typically when the display was large, probably due to limited peripheral vision).

Can process tracing data really tell us about underlying decision processes? Will it know a game theorist when it meets one? Costa-Gomes et al. (2002) and Johnson et al (2002) ran controls where subjects were trained to follow game theoretic principles, to see what information acquisition by "game theorists" would look like. In our work, we ran several groups who had been trained to calculate the backward induction equilibrium. Their process was quite different from untrained subjects bargaining with each other (or with computerized opponents)— viz., the trained subjects opened the 3^{rd} and 2^{nd} round payoff boxes much more often, and made many more backward transitions 3^{rd} round payoff box to earlier-round boxes. A similar control has been run for normal-form games by Costa-Gomes et al. with similar results. In the decision literature, Bettman et al. (1990) reports success training subjects to follow various choice strategies.

How important is the position of payoffs in the display? In most western languages, people read from left to right and top to bottom. One might well be concerned that this left-right, top-down bias is reflected in which boxes are opened first. To measure this bias, we ran a group of subjects where the order of the payoffs were reversed from top to bottom (i.e., the 3^{rd} round payoff was at the top and the 1^{st} round at the bottom). The results looked like an inverted version of Figure 2, and replicated the results we discuss below.

Together, these results indicate that forcing players to open boxes to gather information *does not* produce important changes in cognition or eventual choices.

Avoiding Pitfalls

There are two potential pitfalls in collecting and analyzing process data. One is that it is crucial to minimize memorization so that we can rely on subjects constantly reacquiring information as they need it (rather than remembering). This design desideratum favors complex displays and subjects who are facile in using mice or other technologies. he fact that subjects typically open the same box more than once in a trial-- often many times— suggests it just as easy for them to move a mouse as it is to remember.

The second pitfall is that observations of decisions trickle out, like a dripping faucet, while process data are like a firehose of information. Since there are so many data (and multiple measures), it is difficult to see what is going on without having a strong theory that tells you exactly what patterns to look for. A strong theory also tells you how to design the experiment. For example, Costa-Gomes et al carefully chose games in which different decision rules led to different choices and to different patterns of information acquisition.

16

The Future

Studying thinking is hard work. By definition, thinking is an internal event and is relatively unobservable.   There may be great promise in the new technologies in neuroscience such as fMRI (functional magnetic resonance imaging), but currently there are significant challenges in these endeavors: Gathering data is costly, the interpretation requires significant knowledge of the brain and its structures, and the data has mediocre limited time resolution. In fact, progress in neuroscience has been rapid largely because results are only accepted when different methods (animal studies, lesion patients, physiological measures, fMRI) corroborate one another (see Camerer, Loewenstein and Prelec, 2003).

> **Comment:** Colin: provide references?

To economic theorists, optimization tools are a familiar hammer which makes everything look like a nail. As a result, economists will find it irresistible to build models in which the choice of decision rules and information can be "reverse engineered" as optimal solutions to decision making under constraint (e.g., Gabaix and Laibson, 200). Optimizing theories have proved very useful, and "optimally suboptimal" theories of bounded rationality can be useful too, provided they stick close to the details of what is known about cognitive mechanisms, or make sharp empirical implications. It is easy for such exercises to degenerate into "just so" stories which posit a very special set of premises under which a behavior is proved to be optimal. However, it is usually difficult to prove that the posited set of premises are the only ones under which the behavior is optimal (i.e., it is easier to prove that the premises are sufficient than that they are necessary). As a result, it is legitimate to judge whether the premises are reasonable using other sources of data. To be taken seriously, such theories should therefore take empirical constraint seriously on either the front end of the modeling exercise (grounding the basic features of the model firmly in empirical regularity) or on the back end (requiring a sharp implication which would falsify the model and is worth testing), or both.

References

Bettman, J. R., Johnson, E. J., & Payne, J. W. (1990). A componential analysis of cognitive effort in choice. *Organizational Behavior & Human Decision Processes, 45*(1), 111-139.

Bolton, Gary E., and Axel Ockenfels. 2000. ERC: A theory of equity, reciprocity, and competition. American Economic Review, 90, 166--93.

Camerer, Colin F. 2003. Behavioral Game Theory : Experiments in Strategic Interaction. Princeton: Princeton University Press.

Camerer, C., Ho, T. H. & Chong, K. (2003a). Models of thinking, learning and teaching in games. American Economic Review Papers and Proceedings.
Camerer, Colin F.; Teck Ho, and Kuan Chong (2003b). A cognitive hierarchy analysis of thinking in one-shot games. http://www.hss.caltech.edu/~camerer/thinking2002.pdf

Camerer, Colin F., Eric Johnson, Talia Rymon and Sankar Sen. 1994. Cognition and framing in sequential bargaining for gains and losses. In. K. Binmore, A. Kirman and P. Tani (Eds.), Frontiers of Game Theory. Cambridge MA: MIT Press, 27-47.

Camerer, Colin F., and George Loewenstein. 2003. Behavioral economics: Past, present, future. In C. F. Camerer, G. Loewenstein, and M. Rabin (Eds.), <u>Advances in Behavioral Economics</u>. Russell Sage Foundation/Princeton University Press.

Camerer, Colin F., Loewenstein, George, & Prelec, Drazen. Neural economics: How neuroscience can inform economics. Unpublished, http://hss.caltech.edu/~camerer/camerer.htm

Costa-Gomes, M., Crawford, V., & Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*. 69, 1193-1235.

Costa-Gomes, Miguel A. and Vincent P. Crawford, "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study", 2002, working paper. http://weber.ucsd.edu/~vcrawfor/

Charness, Gary, and Mathew Rabin. 2002. Understanding social preferences with simple tests. <u>Quarterly Journal of Economics</u>, 117, 817-869.

Crawford, Vincent P. 2003. Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions, <u>American Economic Review</u>, 93, 133-149.

El-Gamal, Mahmoud, and David Grether. 1995. Are people Bayesian? Uncovering behavioral strategies. <u>Journal of the American Statistical Association</u>, 90, 1137-45.

Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. *37*, 215-251.

Fehr, Ernst and Simon Gachter. Fairness and retaliation: The economics of reciprocity, Journal of Economic Perspectives, 14, Summer 2000, 159-181.

Fehr, Ernst. & Schmidt, Klaus M. (1999). A theory of fairness, competition and cooperation. <u>Quarterly Journal of Economics</u>, <u>117</u>(3), 817-868.

Frederick, Shane; George Loewenstein and Ted O'Donoghue. 2002. Time Discounting and Time Preference: A Critical Review, <u>*Journal of Economic Literature*</u>, **40**(2), 351-401.

Gabaix, Xavier and David Laibson. 2000. A new challenge for economics: The "frame problem". In I. Brocas and J. Carillo eds, *Collected Essays in Psychology and Economics*, Oxford University Press.

Goeree, Jacob K. and Charles A. Holt. 2001. Ten little treasures of game theory and ten intuitive contradictions. <u>American Economic Review</u>, 91, 1402-22.

Harless, D. & Camerer. C. F. (1994). The predictive utility of generalized expected utility theories. Econometrica, 62, 1251-1290.

Harless, David W., and Colin F. Camerer. 1995. An error rate analysis of experimental data testing Nash refinements. European Economic Review, 39, 649--60.

Jacoby, J., Jaccard, J., Kuss, A., Troutman, T., & Mazursk. (1985). *New directions in behavioral process research*. New York U., NY, NY.

Johnson, E. J., Camerer, C. F., Sen, S., & Rymon, T. (2002). Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining. *Journal of Economic Theory*, 104, 16-47.

Johnson, E. J., Payne, J. W., Schkade, D. A., & /Bettman, J. R. (1991). Monitoring information processing and decisions: The MouseLab system Version 5.0.

Johnson, E. J. & Schkade, D. A. (1989). Bias in utility assessments: Further evidence and explanations. Management Science 35. 406-424.

Just, M. A., & Carpenter, P. A. (1976). Eye fixations and cognitive processes. *8*, 441-480.

Kohlberg, Elon, and Jean-Francois Mertens. 1986. On the strategic stability of equilibrium. Econometrica, 54, 1003--37.

Loewenstein, G., Thompson, L. & Bazerman, M. (1989). Social utility and decision making in interpersonal contexts. *Journal of Personality & Social Psychology, 57,* 426-441.

Lohse, G. L., & Johnson, E. J. (1996). A comparison of two process tracing methods for choice tasks. *Organizational Behavior and Human Decision Processes, 68*(1), 28-43.

McKelvey, R. D. & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. Games and Economic Behavior, 7, 6-38.

McKelvey, R. D. & Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. Experimental Economics, 1, 9-41.

Nagel, Rosemarie. 1999. A survey of research on experimental beauty-contest games. In D. Budescu, I. Erev, and R. Zwick (Eds.), Games and Human Behavior: Essays in Honor of Amnon Rapoport. Mahwah, NJ: Lawrence Erlbaum.

Nisbett, Richard C. and Timothy D. Wilson. 1977. Telling more than we can know: Verbal reports on mental processes, Psychological Review, 84, 231-259.

Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *16*, 366-387.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York, NY, US: Cambridge University Press.

19

Payne, J. W., & Braunstein, M. L. (1978). Risky choice: An examination of information acquisition behavior. *6*, 554-561.

Rabin, M. (1993). Incorporating fairness into game theory and economics. <u>American Economic Review,</u> <u>83</u>(5), 1281-1302.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica, 68*, 1281-1292.

Reiskamp, Jorg, and Hoffrage, U. 2000. When do people use simple heuristics and how can we tell? In G. Gigerenzer, P. Todd, ABC Research Group (Eds.), <u>Simple Heuristics that</u> <u>Make Us Smart.</u> New York: Oxford University Press.

Rubinstein, Ariel. 1982. Perfect equilibrium in a bargaining model. <u>Econometrica</u>, 50, 97--109. SAYS 1980 IN TEXT I BELIEVE

Rubinstein, Ariel. 1999. Experience from a course in game theory: Pre- and postclass problem sets as a didactic device. <u>Games and Economic Behavior</u>, 28, 155-70.

Russo, J. E. (1978). *Eye fixations can save the world: A critical evaluation and a comparison between eye fixations and other information processing methodologies* (J. K. Hunt, Trans. Vol. 5). Ann Arbor, MI: Association for Consumer Research.
Russo, J. E., & Rosen, L. D. (1974). An eye fixation analysis of multialternative choice. *Memory and Cognition, 3*, 267-276.
Schkade, D. A., & Johnson, E. J. (1989). Cognitive processes in preference reversals. *Organizational Behavior & Human Decision Processes, 44*(2), 203-231.

Schotter, Andrew, Keith Weigelt, and Charles Wilson. 1994. A laboratory investigation of multiperson rationality and presentation effects. <u>Games and Economic Behavior</u>, 6, 445--68.

Ståhl, Ingolf. 1972. <u>Bargaining Theory</u>. Stockholm: Economic Research Institute.

Stahl, D., & Wilson, P. (1995). On player's models of other players: Theory and experimental evidence. <u>Games and Economic Behavior</u>*, 7*, 218-254.
Tversky, A., & Kahneman, D. (1991). Loss Aversion in Riskless Choice - a Reference-Dependent Model. *Quarterly Journal of Economics, 106*(4), 1039-1061.

Van Damme, Eric, 1989. Stable equilibria and forward induction, *Journal of Economic Theory*, **48** (1989) 476-496.

Weizsacker, Georg. In press. Ignoring the rationality of others: Evidence from experimental normal-form games. Games and Economic Behavior.

Wason ?? I'D SAY TAKE THIS OUT

Johnson-Laird TAKE THIS OUT OF TEXT, NOT ON WEB yet ETC [I reviewed it for Psych
    Sci] ADD RUBINSTEIN 1999 INSTEAD, MAKES THE SAME POINT