

Imperfect choice or imperfect attention?

Understanding strategic thinking in private information games *

Isabelle Brocas

Juan D. Carrillo

University of Southern California and CEPR

Stephanie W. Wang

Colin F. Camerer

California Institute of Technology

July 12, 2010

Abstract

In experiments, people do not always appear to infer the information of other players from their choices. To understand this thinking process further, we use “Mousetracking” to record which game payoffs subjects look at, and for how long, in games of private information with three information states, which vary in strategic complexity. Subjects often deviate from Nash equilibrium choices, converge only modestly toward equilibrium across 40 trials, and often fail to look at payoffs which they need to in order to compute an equilibrium response. When cluster analysis is used to group subjects according to lookup patterns and choices, the clusters appear to correspond approximately to level-3, level-2 and level-1 thinking in level-k cognitive hierarchy models. Deviations from Nash play are associated with failure to look at the necessary payoffs. The connection between looking and choices is strong enough that the time durations of looking at key payoffs can predict choices, to some extent, at the individual level and at the trial-by-trial level.

*Support of LUSK Center (IB), the Office of the Provost at USC and the Microsoft Corporation (JDC), HFSP, NSF and Moore Foundation grants (CFC) and Moore Foundation (SWW) is gratefully acknowledged. Yi Zhu provided excellent research assistance. Mousetracking was developed by Chris Crabbe and Walter Yuan as an extension to their Multistage program. We are very grateful for their remarkable combination of enthusiasm and speed. Helpful comments were received from audiences at USC Law, Moore Foundation Retreat (March 09), ESA Washington 2009, Stanford SITE 2009, Southampton and Edinburgh. Send correspondence to: <brocas@usc.edu>.

1 Introduction

Equilibrium analysis in game theory is a powerful tool in social and biological sciences. However, empirical choices often deviate from equilibrium in ways that are consistent with limited strategic thinking and imperfect response (Camerer, 2003). Private information games are especially interesting because limited strategic thinking in those games could mean that strategically naïve agents do not fully understand the link between the information other agents have and the choices those agents will make—e.g., they will not fully anticipate adverse selection. Many types of lab and field evidence (described just below) suggest some players do not make rational inferences in private information games.

Understanding strategic thinking in private information games is important because they are such popular tools for modeling contracting, bargaining, consumer and financial markets, and political interactions. If there is widespread strategic naivete, then distortions due to hidden information can be larger in these markets than predicted by equilibrium analysis, and ideal policy responses may be different (e.g., Crawford, 2003).

The experimental games we use to measure strategic thinking are two-person betting games with three states and two-sided private information. Players privately observe a state-partition (either one or two of the three states) and must choose whether to bet or not bet. If either player does not bet, both players earn a known sure payoff.

In some of the games, the no-trade theorem applies (Milgrom and Stokey, 1982). This theorem shows that if an initial allocation is Pareto-efficient (or there are initial trades to efficiency), and rationality is common knowledge, further trading cannot occur for informational reasons alone. In field data—many financial markets, prediction markets, sports events—there appears to be substantial zero-sum trading based on perceived information advantages. But it is difficult to rule out the hypothesis that this betting observed in the field is rational and occurs because of risk preference, hedging motives, or a consumption complementarity between betting on events and watching them. However, these alternative explanations for betting can largely be controlled in laboratory experiments.

In fact, three previous experimental studies have reported behavior in private information betting games similar to ours. Mutual betting is never a Nash equilibrium in these games, but in all of them betting is common. In the pioneering study, Sonsino et al. (2002) examine games with 4 and 8 states played 250 times. Sovik (2004) replicated their 4-state game with fewer repetitions, higher stakes, and using the strategy method. Rogers et al. (2010) also replicate the 4-state game with some small design changes, and compare the ability of different behavioral models to explain the data (cognitive hierarchy and QRE models fit equally well). All these papers find substantial rates of betting, with surprisingly little learning over many trials.

Two other papers explore behavior in simple experimental assets markets in which no trade is predicted to occur. Carrillo and Palfrey (2008) consider a bilateral bargaining game where an asset of uncertain but pure common value can be traded using either a seller take-it-or-leave price or a double auction mechanism. They find substantial trade. Angrisani, Guarino, Huck and Larson (2008) observe a substantial amount of trade (about as much as in comparable environments where trade should occur). Trading declines substantially over 30 periods, particularly if information is not too noisy or the information gap is particularly salient.

Besides the overbetting and overtrading just described, lab and field evidence suggest there are strategic thinking limits in many other different games with private information. Examples include the winner's curse in common value auctions (Kagel and Levin, 2002), overbidding in private value auctions (Crawford and Iriberry, 2007), the lemons problem in adverse selection markets (Bazerman and Samuelson, 1983; Charness and Levin, 2009); settlements in zero-sum games (Carrillo and Palfrey, 2009); and overcommunication in sender-receiver games (Cai and Wang, 2006; Wang et al., 2010). Field examples include bidding consistent with an unanticipated winner's curse in offshore oil lease auctions (Capen et al., 1971), consumer reactions to disclosure of product quality (Mathios, 2000; Jin and Leslie, 2003; Brown et al., 2009), and evidence that individual stock investors trade too much (Odean, 1999).¹ The empirical evidence could conceivably be explained by two types of theories: (1) *Imperfect choice*: Players attend to all the information in the game and analyze it fully, but either choose imperfectly or believe others will choose imperfectly (e.g., Weiszacker, 2003); or (2) *Imperfect attention*: players do not fully analyze the game, and hence choose imperfectly compared to a benchmark theory which assumes full analysis.

Different theories have emerged recently of how imperfect choice and imperfect attention are combined in different specific ways.² In level-k or cognitive hierarchy (CH) theories there is a hierarchy of agent reasoning in which higher-level thinkers respond to what they expect lower-level thinkers to do (see e.g. Camerer et al., 2004; Crawford, 2008). Quantal response equilibrium (QRE) assumes players' beliefs are statistically accurate but players respond noisily to expected payoffs (McKelvey and Palfrey, 1995). In cursed equilibrium, players underestimate the link between the private information and strategies of other players but otherwise optimize (Eyster and Rabin, 2005). In analogy-based expectation equilibrium (ABEE), players draw analogies between similar games and use only coarse

¹Ross (1989) writes: "It is difficult to imagine that the volume of trade in security markets has very much to do with the modest amount of trading required to accomplish the continual and gradual portfolio rebalancing inherent in our current intertemporal models" (p. 94)

²See, e.g., Carrillo and Palfrey (2009); Rogers et al. (2010).

statistics within analogy classes to form beliefs (see Jehiel, 2005; Jehiel and Koessler, 2008; see Huck, Jehiel and Rutter (2007) for experimental evidence).

These different theories can all potentially explain overbetting choices. However, they assume fundamentally different cognitive processes. For example, in CH and level-k theories some players bet because they don't realize others are better informed than they are. In ABEE players form analogy classes which can hide differences in other players' information.³ In QRE theories, players know that others have better information but they bet and expect to lose small amounts, because they choose imperfectly.

Theories that have different underlying cognitive processes will make different predictions in other games, predict different responses to changes in payoffs,⁴ and can predict different responses to variables like payoff levels, time pressure and information displays. Therefore, it is useful to know, for making predictions about other games and economics conditions, which theories are most accurate about *both* cognition and choices. Doing so benefits greatly from having an experimental paradigm which measures both choices and attention to game structure information.

We use the simple class of two-person betting games in this paper because they capture the essence of many economic situations with adverse selection, and they also allow measures of both choice and visual attention to payoffs which distinguish different theories. By examining the decision-making process in one class of games, we hope to shed some light on what might be driving apparent partial ignorance of adverse selection in other economic situations.

The games we use have some novel properties for the study of behavior (independent of using them to measure attention). The variations test robustness across payoff structures and help identify behavioral strategies. Nash and level-k models predict a mixture of betting and no-betting across the games, so there is a rich statistical variety of behavior. The different information states also create strategic situations which increase in complexity of processing required, denoted D1 (simple) and D2 (complex) games (which is very useful for distinguishing levels of reasoning). However, even the most complex games are much less complicated than most games analyzed in theory (and in naturally-occurring settings). The presumption in doing such experiments is that if strategic thinking is limited in these relatively simple games, it is likely to be limited in more complex games as well.⁵

³Note that the ideal experiment to test ABEE in our setting is to not make the information partitions common knowledge, but instead see if players search for structure information or do so only imperfectly (supporting coarse "analogies"). This is an interesting experiment, but is quite different than ours, and might also benefit from using a combination of choice and attention measurement to infer what the analogy classes are.

⁴For example, QRE predictions are affected statistically by changing *any* of the payoffs in the games we study. Level-k and cursed equilibrium predictions are only affected by *some* payoff changes.

⁵In pilot experiments we used 4-state games which were already much more complex, especially in terms

The game structures also permit close association of visual attention and choice. Information acquisition is measured by hiding payoff information in opaque boxes and using a Mousetracking system. The payoffs are revealed when a mouse is moved into the box (and the left button is held down, so players must consciously exert effort to pay attention).

Several studies have used similar attention measures to understand cognitive processes in economic decisions.⁶ These techniques were first applied to games by Camerer et al. (1993), who found that a failure to look ahead to future payoffs was linked to non-equilibrium offers in alternating-offer bargaining (Johnson et al., 2002). Information acquisition measures were then used to study forward induction (Johnson and Camerer, 2004), level-k models in matrix games (Costa-Gomes et al., 2001), two-person “beauty contest” games designed to have good type identification (Costa-Gomes and Crawford, 2006), learning in normal-form games (Knoepfle et al., 2009) and sender-receiver games with strategic information transmission (Wang et al., 2010).⁷ Crawford (2008) summarizes these findings and argues for the methodological value of information acquisition measures.⁸

Note that our paper is the first to apply the Mousetracking technique to two-sided private information games. This is a substantial expansion of previous literature for two reasons. First, there is an empirical puzzle that needs to be resolved— why do people seem to ignore adverse selection, by betting and trading so much in experiments and appearing to exhibit strategic naivete in field data? Since theories of imperfect choice and imperfect attention could both explain these data, measuring information use directly is potentially an efficient way to contribute to resolving the empirical puzzle.

Second, equilibrium play in our private information games requires very special kinds

of measuring information acquisition and using those data to distinguish different choice theories. Sonsino et al. (2002) use games with up to 8 states.

⁶The earliest studies focus on basic cognitive processes like reading and visual perception. The pioneering work in decision making is on multi-attribute choice (Payne et al., 1993) and advertising (Lohse, 1997). Recent economic studies propose and test ‘directed cognition’ (Gabaix et al., 2006), sequential search under time pressure (Reutskaja et al., in press), calibrate accumulator race-to-barrier models (Armel and Rangel, 2008), and look for within-attribute versus within-choice comparisons (Arieli et al., 2009).

⁷Wang et al. (2010) were the first to combine lookup information and pupil dilation measures, and to use those measures to see how well private information could be predicted from information measures. They found that a hypothetical subjects using out-of-sample forecasts based on another player’s choices and observed information lookup and pupil dilation could earn more than the average subject, and could earn about 50% of the incremental value from average earnings to ideal earnings (based on empirical best responses).

⁸Many other variables associated with strategic choice could be measured along with information lookups, such as response times (Rubinstein, 2006), physiological measures, and neural activity (e.g., Camerer, 2008; Coricelli and Nagel, 2009). The combination of belief elicitation and choices (Costa-Gomes and Weizsäcker, 2008; Tingley and Wang, 2010) can also help in understanding the underlying cognitive processes. Attentional data is, in that respect, complementary to these other techniques.

of lookups. To correctly anticipate what other players will do, players must look at the numerical value of payoffs which they know they will never receive (because those payoffs occur in states the players know are impossible, but which they also know others think are possible). This kind of attention-to-the-impossible is not required by any of the games previously studied using Mouselab or eyetracking. Paying this kind of attention is quite counterintuitive and therefore could be highly diagnostic of sophisticated strategic thinking.

Our design has another simple advantage. There are only seven payoff boxes but it is not necessary to look at all of them to discover an equilibrium choice. Players could deviate from rational information acquisition by looking at too few boxes or by looking at too many. In other experimental designs rationality requires looking at all the boxes, so any noise in lookups is interpreted as too little rationality. This censoring-induced bias is not present in our design.

Finally, note that other private information games which are behaviorally interesting are not well suited to inferring cognition from attention. For example, in a typical winner's curse auction design, all players get a signal of value distributed around the true value, and they know that others get different signals. In private value auctions all players get private values from a known prior distribution. In these auctions, there is no obvious way to display information about what other players see or know, so that strategic thinking can be inferred from measured attention.⁹ Our design therefore has the advantage of being able to clearly connect visual attention of one player to the potential payoffs of other players, in order to make inferences about steps of strategic thinking as hypothesized in different theories.

Here is a summary of the basic results: Subjects play Nash equilibrium about half the time in simple D1 games and a quarter of the time in more complex D2 games. The combination of choices and lookup analysis suggest that some subjects look at the payoffs necessary to make an equilibrium choice, but sometimes these lookups do not translate into making the equilibrium choice. In other cases, they do not look at all the necessary payoffs and do not make equilibrium choices.

There is heterogeneity in both lookups and choices. Subjects can be classified into clusters using both measures. In one cluster, subjects usually look at the necessary payoffs and play Nash (corresponding to a level-3 type), but this cluster is small (about 15% of the subjects). Two other clusters look at the necessary payoffs and usually play Nash in the simpler D1 games, but look and choose less strategically in the more complex D2 games. In a fourth cluster– the most common– subjects spend less total time making

⁹For example, you cannot put other players' signals in boxes and see if a player opens them, because that design would change the information structure of the game.

choices, look at necessary payoffs less often, and rarely play Nash. Differences in lookup patterns can be used to predict choice with some reliability at the individual level, and on a trial-by-trial basis.

Actual earnings in the four clusters are interesting. Playing Nash is an empirical best response in the simple D1 games but not in complex D2 games. The first cluster does not earn the most money because they look the most strategically and play Nash often in both types of games, but they act as if they overestimate the rationality of their opponents in D2 games (which is an earnings mistake). The middle clusters earn the most. The fourth cluster, who “underlook” compared to theory, earns the least.

The particular heterogeneity of information processing and subsequent choices we observed across players suggested a decent fit with the level-k model (with some interesting exceptions). It is possible that QRE, cursed equilibrium, and ABEE fit the behavioral choices reasonably well, but they do not make any clear predictions about information processing requirements, or the link between what information subjects look up and what strategies they choose. Indeed, if those theories are taken seriously as algorithms describing what information people attend to and process to make choices according to the theory, they require full attention to all payoffs. It is possible that next-generation theories could emerge in which imperfect attention is associated with imperfect choice, but the theories as currently understood are not of that type. Therefore, since the primary goal is to learn why there is so much overbetting in previous experiments, by measuring both strategy choices and information lookup, we will concentrate only on basic regularities in the data and a comparison with levels-of-cognition models.

The paper is organized like so. Section 2 contains the theory and design and discusses some features of Mousetracking. Section 3 presents an aggregate analysis of the experimental data. Section 4 provides a cluster analysis, sorting subjects into groups based on attentional information and choices of subjects. Section 5 selects attentional variables necessary to predict behavioral types and eventual patterns of choices across games. Section 6 concludes.

2 Theory and Design

2.1 Theory: Equilibrium

We consider the following game. Nature draws a state, A , B or C , and communicates partitioned information to two players: ‘ C ’ or ‘not C ’ (i.e., $\{A,B\}$) to player 1 and ‘ A ’ or ‘not A ’ (i.e., $\{B,C\}$) to player 2. Players choose ‘Bet on the state’ or ‘Sure payoff’ (S). If either of them chooses S , they each earn the number in the box under S . If both choose to bet, the payoffs for the two players (1 and 2) depend on the state (A , B , C) and are

shown in the top and bottom rows of the matrix. *All* three-state, two-player games with the information partition described above fall into one of three classes, characterized by the type of information search required for equilibrium choice. We call them F, D1 and D2. Figure 1 illustrates each class using the payoffs of one of the games in the experiment.

<table border="1" style="border-collapse: collapse; width: 100%;"> <tr> <td style="width: 10%;"></td> <td style="width: 20%;">A</td> <td style="width: 20%;">B</td> <td style="width: 20%;">C</td> <td style="width: 20%;">S</td> </tr> <tr> <td>1</td> <td>25</td> <td>5</td> <td>20</td> <td>10</td> </tr> <tr> <td>2</td> <td>0</td> <td>30</td> <td>5</td> <td></td> </tr> </table> <p>F choice: 1 in $\{C\}$</p>		A	B	C	S	1	25	5	20	10	2	0	30	5		<table border="1" style="border-collapse: collapse; width: 100%;"> <tr> <td style="width: 10%;"></td> <td style="width: 20%;">A</td> <td style="width: 20%;">B</td> <td style="width: 20%;">C</td> <td style="width: 20%;">S</td> </tr> <tr> <td>1</td> <td>25</td> <td>5</td> <td>20</td> <td>10</td> </tr> <tr> <td>2</td> <td><u>0</u></td> <td>30</td> <td>5</td> <td></td> </tr> </table> <p>D1 choice: 1 in $\{A, B\}$</p>		A	B	C	S	1	25	5	20	10	2	<u>0</u>	30	5		<table border="1" style="border-collapse: collapse; width: 100%;"> <tr> <td style="width: 10%;"></td> <td style="width: 20%;">A</td> <td style="width: 20%;">B</td> <td style="width: 20%;">C</td> <td style="width: 20%;">S</td> </tr> <tr> <td>1</td> <td><u>25</u></td> <td><u>5</u></td> <td>20</td> <td>10</td> </tr> <tr> <td>2</td> <td><u>0</u></td> <td>30</td> <td>5</td> <td></td> </tr> </table> <p>D2 choice: 2 in $\{B, C\}$</p>		A	B	C	S	1	<u>25</u>	<u>5</u>	20	10	2	<u>0</u>	30	5	
	A	B	C	S																																											
1	25	5	20	10																																											
2	0	30	5																																												
	A	B	C	S																																											
1	25	5	20	10																																											
2	<u>0</u>	30	5																																												
	A	B	C	S																																											
1	<u>25</u>	<u>5</u>	20	10																																											
2	<u>0</u>	30	5																																												

Figure 1 Different classes of situations for a given game. **Possible payoffs** in the information set are in bold. Payoffs in the MIN set are underlined.

To imagine practical applications, think of the information partitions as resulting from coarse categorization or imperfect perception. For example, if the states are (Bad, Medium, Good), some player 1 agents may have expertise in detecting Good states but cannot distinguish Bad from Medium, whereas some player 2 agents have expertise in detecting Bad states but cannot distinguish between Good and Medium. Importantly, while player 1 agents cannot distinguish Bad from Medium, they know that other player 2 agents can (and vice versa).

The left table illustrates the full information F class. In the F class, player 1 has a singleton information set $\{C\}$. In order to choose the optimal strategy, she only needs to compare the two sure payoffs, $[1C]$ and $[S]$. In this example, she knows that if she bets she will obtain 20 for sure (if the other player bets), which is higher than the sure payoff $S = 10$. We call the set of information lookups that is necessary to choose the optimal strategy the “Minimum Information Necessary” (MIN). (In this F game, the MIN set is $[1C, S]$).

The table in the center illustrates the D1 class. Player 1 has information set $\{A, B\}$. She must see whether Player 2 will bet in $\{A\}$ (by looking at the $[2A]$ payoff of 0) and look at her payoffs in $[1A]$ and $[1B]$ (as well as $[S]$, of course). This process is strategic because it requires (i) looking at another player’s payoff (which the F case did not), (ii) realizing the different information partition of the opponent (which is explained as part of the instructions and made clear visually), and (iii) making an inference assuming the simplest level of rationality (dominance) of the opponent.

The third table, on the right, illustrates the D2 class. Player 2 has information set $\{B, C\}$. Whether or not Player 1 will bet in $\{C\}$ does not fully determine her decision. She must deduce whether Player 1 will bet in $\{A, B\}$. That deduction requires looking at

her own $[2A]$ payoff. Thus, she must look at all of her own payoffs $[2A, 2B, 2C]$, as well as $[1A, 1B]$ (and $[S]$).¹⁰ Notice also that looking at $[2A]$ means looking at a possible payoff that she knows with certainty will not occur (an impossible counterfactual). This is quite counterintuitive and is therefore a clear hallmark of strategic reasoning.

The situations can be categorized from simplest to most complex, based on the number of lookups in the MIN set, and the degree of strategic thinking required. In F choices the information required to play correctly is minimal and no strategic thinking is involved. D1 choices require a more subtle reasoning and paying attention to the rival's payoff. D2 requires a further step in reasoning: the subject must anticipate how a rival will behave in a D1 situation and best respond to it. Considering those three classes of situations allows us to judge comparative statics on behavior as a function of complexity. Notice also that subjects never need to look at all the payoffs to reach a decision. Therefore, it is possible to determine whether subjects look too much or too little, compared to the MIN benchmarks, and to study the effects of these tendencies.

Figure 2 shows the five payoff variants used in the experiment. (Game 1 is the example that was used above for illustration.) An important feature of the design is that betting is not always an equilibrium strategy. Among the 10 two-state information sets, 6 are in the D1 class (3 predict betting and 3 no betting) and 4 are in the D2 class (2 predict betting and 2 no betting). The games were also constructed to minimize the number of two-state information sets in which naively comparing the average of the two payoffs to $[S]$ gives the same choice as Nash analysis (the only ones are $\{B, C\}$ in game 3 and $\{A, B\}$ in game 5).

Game 1				Game 2				Game 3						
	A	B	C	S		A	B	C	S		A	B	C	S
1	25	5	20	10	1	0	20	5	15	1	5	20	20	15
2	0	30	5		2	0	20	5		2	10	25	10	
Game 4				Game 5										
	A	B	C	S		A	B	C	S					
1	15	0	5	10	1	10	25	25	20					
2	15	5	25		2	25	40	15						

¹⁰In fact, it is not strictly necessary to look at $[1A]$ in D1 or D2. Player 1 anticipates that Player 2 will never bet in $\{A\}$, and therefore does not need to check her own payoff in that state. For analogous reasons, Player 2 in D2 does not need to look at $[2B]$ either. Overall, strictly speaking, MIN in D1 is $[1B, 2A]$ and MIN in D2 is $[1B, 2A, 2C]$. This more refined definition would make little difference in the data analysis since subjects virtually always look at all the payoffs in their information set.

Figure 2 Payoff-variations of the betting game.

Two features of this set of games are worth noting. First, in games 2 and 3 it is mutually-beneficial to agree to bet in state B . Therefore, it is *not* always the case that betting is a mistake (as in some previous experiments).¹¹ Second, suppose that players have social preferences over the payoffs of others, in the sense that they are willing to sacrifice money to reduce inequality (Fehr and Schmidt, 1999), benefit the worst-off player, or increase the total payoff (Charness and Rabin, 2002). Depending on the model and parameters of social preferences, we could predict in which games and for which class of situations these subjects would deviate from Nash. Empirically, however, such sacrificial non-Nash behavior is extremely rare in the simplest F class, when some social preferences might predict Nash deviations most clearly.¹² The use of Mousetracking is also helpful in ruling out the influence of social preferences in these games. Socially-regarding players always need to look at payoffs of other players, even in the possible singleton information sets, to decide whether it is worth deviating from Nash in order to express their social preference. But in F games players generally look at the possible payoff of the other player only rarely and quickly.

Theoretical predictions from Nash equilibrium are shown in Figure 3 for the two-state information sets. ‘Y’ and ‘N’ denote the equilibrium choice to ‘bet’ (Yes) and ‘not bet’ (No) respectively. We include the class of the situations in parentheses. The asterisk * shows when the equilibrium choice coincides with a naïve strategy that would consist of betting if the average of the payoffs in the information set exceeds the sure value S . Predictions in the full information sets are omitted as they simply consist in betting if and only if the state payoff exceeds the sure payoff. These predictions are all made under the assumption of risk neutrality. Sufficiently risk averse or risk loving attitudes could lead to different equilibrium predictions in some of the games. The players in our experiments, however, participate in all the games so within-subject choice patterns allow us to reject risk attitudes as the primary factor behind equilibrium deviations. Furthermore, Rogers

¹¹Earlier studies have been criticized for predicting “inaction” in all trials (never bet, always go for the sure payoff). With boundary predictions of this sort, any deviation from equilibrium will look like systematic bias—overbetting, rather than betting deviated from the predicted rate. It is also possible that subjects would presume that experimenters would not create a game in which they always take the same action (not betting) which will naturally introduce some betting, and hence overbetting. However, in our setting, they could just as easily underbet as overbet.

¹²Consider player 2 who knows that it is state A in game 5. While the Nash strategy is to bet because the A payoff is higher than the sure payoff, an inequity-averse player might sacrifice and not bet so that both players could have 20 instead of the vector (10, 25). Similarly, a player 1 who knows it is state C in game 4 might bet contrary to the Nash solution due to efficiency concerns because player 2 would receive 25 rather than 10 and the total earnings for both would be 30 rather than 20. However, Non-Nash behavior in these two scenarios rarely occurs (about 3% of the time).

et al. (2009) controlled for risk attitudes by using gambles as outside options rather than sure payoffs and still found substantial non-equilibrium betting.¹³

	Game 1			Game 2			Game 3			Game 4			Game 5		
	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
1	N(D1)			Y(D1)			Y(D1)			Y(D2)			N(D1)*		
2				Y(D1)			Y(D2)*			N(D1)			N(D2)		

Figure 3 Nash equilibrium and class of situation in each game.

2.2 Theory: Level-k

This section describes one version of how ‘level-k subjects’ will act in this game, including both lookup occurrence and choice.¹⁴

The first step is to define the behavior of a level-0 subject. We assume that a level-0 randomizes between ‘betting’ and ‘no betting’.¹⁵

A level-1 player best responds to level-0. Because she believes others randomize, she only needs to look at her information set, average the payoffs, compare this average to $[S]$, and choose the best option. If we restrict the analysis to games in which this naïve strategy differs from Nash play, level-1 agents should never play Nash in D1 or in D2 (see Appendix section 7.5 for an analysis of the other cases).

Level-2 players best respond to level-1. Subjects must therefore open the boxes a level-1 would open to deduce the level-1 strategy, and then determine her own best response. A level-1 opponent with information set $\{A\}$ will open $[2A, S]$ and a level-1 opponent with information set $\{B, C\}$ will open $[2B, 2C, S]$. A level-2 subject then needs to open these

¹³Our equilibrium corresponds to trembling-hand perfection. There are other Bayesian Nash equilibria which are not trembling-hand perfect and which we therefore neglect. For example, in game 1 a player 1 in $\{C\}$ could choose not to bet because she is indifferent between betting and not betting (if player 2 doesn’t bet in $\{B, C\}$). See Sovik (2004) for a sophisticated discussion of other solution concepts and their implications for betting patterns. (Note, however, that player 1 in game 1 with partition $\{C\}$ virtually always bet in our experiment).

¹⁴As in many games, the level-k assumption that players doing k steps of reasoning best respond to the assumption that others are doing k-1 is easier to work with theoretically than the cognitive hierarchy assumption that level-k’s best respond to players at all lower levels.

¹⁵Other specifications are possible. The most natural alternative is a dominance-satisfying level 0 who makes Nash choices in F games. Since level-1 players always satisfy (strict) dominance, level-2 players in our specification act like level-1 players in this alternative specification in D1 games, so if our specification is wrong relative to this alternative, it simply “misnumbers” levels (inferring level 2 when it should be level 1). This is a minor difference in terms of basic conclusions.

boxes and also open $[1A, 1B]$ to figure out her best response. Given this type of reasoning, a level-2 subject will inevitably miss $[1C]$, her payoff in the information set that she knows for sure will not be realized. This means that in D1 games she will look at MIN, but in D2 games she will not. In our games, payoffs are designed in such a way that level-2 agents always play Nash in D1. They do not play Nash in D2 except in Game 5.

A level-3 subject best responds to a level-2. She will open all the boxes, that is, she will always look at MIN. She may or may not play Nash given that she best-responds to a player who does not necessarily play Nash. However, in our games, payoffs are such that a level-3 subject always plays Nash. Any subject with level- k (> 3) will also open all boxes. Because she best responds to a Nash player, she will also play Nash. Overall, all levels 3 and above look and play at equilibrium and therefore cannot be distinguished.

Figure 4 summarizes the patterns of lookup and choices of each relevant level- k (X denotes a box the subject opens). We use the level- k specification rather than the cognitive hierarchy one because the former allows us to make crisp parameter-free predictions about lookup and choices for each level. Under cognitive hierarchy, in contrast, the predictions depend on the parameter τ . For example, if τ is sufficiently low, then all levels above 0 could exhibit the same behavior. Quantal Response Equilibrium does not make explicit predictions about lookup patterns and Rogers et al. (2009) have already tested how well QRE and CH fit choices in similar betting games. We always take the perspective of a player 1 with the $\{A, B\}$ partition in the graphical representations for parsimony.

	Level-1	Level-2	Level-3 (and above)												
	A B C S	A B C S	A B C S												
1	<table border="1" style="display: inline-table; border-collapse: collapse;"> <tr><td>X</td><td>X</td><td></td></tr> </table> <table border="1" style="display: inline-table; border-collapse: collapse; margin-left: 10px;"> <tr><td>X</td></tr> </table>	X	X		X	<table border="1" style="display: inline-table; border-collapse: collapse;"> <tr><td>X</td><td>X</td><td></td></tr> </table> <table border="1" style="display: inline-table; border-collapse: collapse; margin-left: 10px;"> <tr><td>X</td></tr> </table>	X	X		X	<table border="1" style="display: inline-table; border-collapse: collapse;"> <tr><td>X</td><td>X</td><td>X</td></tr> </table> <table border="1" style="display: inline-table; border-collapse: collapse; margin-left: 10px;"> <tr><td>X</td></tr> </table>	X	X	X	X
X	X														
X															
X	X														
X															
X	X	X													
X															
2	<table border="1" style="display: inline-table; border-collapse: collapse;"> <tr><td></td><td></td><td></td></tr> </table>				<table border="1" style="display: inline-table; border-collapse: collapse;"> <tr><td>X</td><td>X</td><td>X</td></tr> </table>	X	X	X	<table border="1" style="display: inline-table; border-collapse: collapse;"> <tr><td>X</td><td>X</td><td>X</td></tr> </table>	X	X	X			
X	X	X													
X	X	X													
	D1: notMIN-notNash D2: notMIN-notNash	D1: MIN-Nash D2: notMIN-notNash*	D1: MIN-Nash D2: MIN-Nash												

Figure 4 Lookup behavior and choice by level- k , $k = 1, 2, 3$. Subject is Player 1 in $\{A, B\}$ (* with the exception of game 5).

2.3 Design and procedures

Our experiments use a “Mousetracking” technique which implements a small change from previous methods. Information is hidden behind blank boxes. The information can be revealed by moving a mouse into the box and clicking-and-holding the left button down.

If subjects release the button the information disappears.¹⁶ Mousetracking is a cheap way to measure what information people might be paying attention to. It also scales up cheaply because it can be used on many computers at the same time.¹⁷ Mousetracking is a tool for a testing theories that make explicit the assumptions made about the underlying cognitive process. Each possible process corresponds to certain patterns of visual attention and information lookup that can be easily tracked with this tool.

There were five sessions in SSEL at Caltech and CASSEL at UCLA. All interactions between subjects were computerized, using a Mousetracking extension of the open source software package ‘Multistage Games’ developed at Caltech.¹⁸ Table 1 summarizes the five sessions in which a total of 58 subjects participated. No subject participated in more than one session. Subjects earned \$18-\$23 plus a show-up fee.¹⁹

Session #	Location	Subjects
1	Caltech	12
2	Caltech	12
3	UCLA	10
4	UCLA	10
5	UCLA	14

Table 1: Summary of Experimental Sessions.

In each session, subjects played betting games with the five different sets of payoffs described in Figure 2. Furthermore, each of the five games were flipped with respect to player role and states to create five more games where Player 1 in the original game had the mirror image payoffs of Player 2 in the other version and vice versa. Subjects played this set of ten games four times for a total of 40 trials in each session. Subjects were

¹⁶In earlier versions of Mouselab, subjects do not have to hold down a button to keep information visible because information is revealed instantly (< 10 msec) when the mouse moves into the box (in the original Camerer et al. (1993)) or when a mouse click is made (Costa-Gomes et al., 2001). This small innovation of requiring click-and-hold is intended to help ensure that subjects are actually attending to the information. It is unlikely to make a big difference compared to other techniques. Nonetheless, it eliminates the necessity to filter out boxes opened for a “short” time (where the minimal duration is defined by the experimenter). It also ensures that an unusually high time spent in a box reflects a long fixation and not to the subject having left the mouse in that position.

¹⁷Cheap scaling-up is a substantial advantage for studying multi-person games and markets compared to single-subject eye-tracking using Tobii or Mobile eyelink camera-based systems, which cost about \$30,000 each.

¹⁸Documentation and instructions for downloading the software can be found at <http://multistage.ssel.caltech.edu>.

¹⁹Camerer et al. (2004) suggests that higher stakes could raise the subjects’ level of cognition and more work can be done to examine how changing the stakes might change the information lookup patterns.

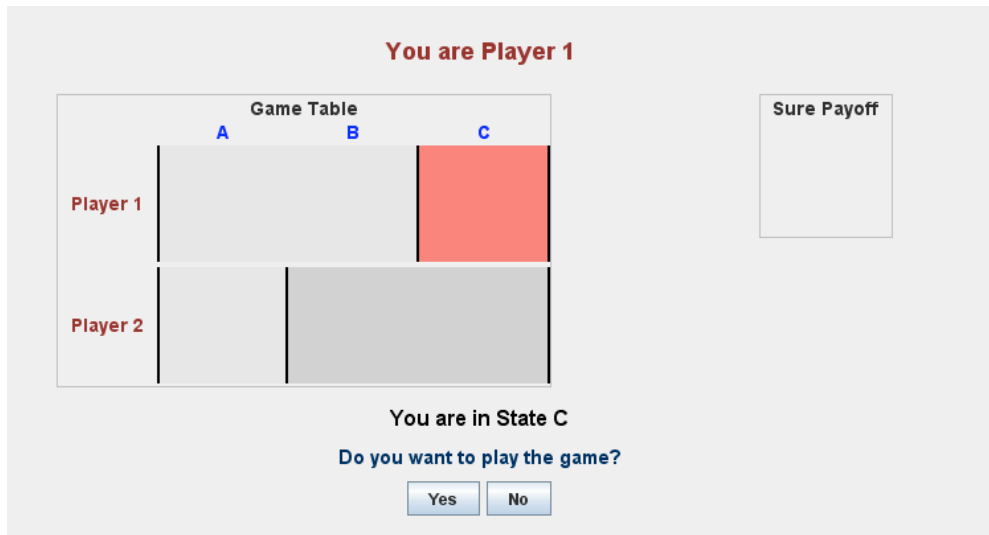


Figure 5 Screenshot of computer interface with all payoffs hidden.

randomly re-matched with another subject and randomly assigned to be either Player 1 or Player 2 in each trial.

Figure 5 displays a sample screenshot where all the payoffs are hidden in the game matrix just as it was in each trial of the sessions. Subjects could move the mouse over any of the cells and click-and-hold to reveal the payoff.²⁰

Subjects had to pass a short comprehension quiz as well as go through a practice trial to ensure that they understood the rules before proceeding to the paid trials. A survey including demographic questions, questions about experience with game theory, poker and bridge, as well as the three-question Cognitive Reflection Test (CRT: Frederick, 2005) was administered at the end of each session.

Plan for analysis: Since there are always many ways to look at data on both attention and choice, the analysis has several parts. The first type of analysis aggregates the data to look at general patterns of strategy choice (section 3). Occurrence of lookups (which payoffs did they look at most frequently?) and lookup duration (how long did they look at each payoff?) are also reported. We find that it is useful to combine lookup and choice

²⁰It is conceivable that subjects played differently than if all boxes were open. We did not perform control sessions but previous studies have shown that behavior is very similar in sessions with open and closed boxes (Costa-Gomes et al., 2001; Johnson et al., 2002). The overall betting probabilities are also quite similar between our study and Rogers et al. (2010) and Sonsino et al. (2002). So there is no reason to think closing boxes makes a substantial difference in general, and specifically in these games.

information too. In this aggregate analysis, lookup and choice statistics are averaged across players. We then create clusters of subjects based on a combination of lookup occurrences and strategy choice (section 4). Cluster analysis is often a sensible middle ground between treating each subject as if they were the same (as in aggregate analysis) and treating each subject as if they were unique (in subject-specific analysis). Finally, we look at the relationship between choice and lookup variables at the individual and trial-by-trial levels (section 5).

3 Aggregate analysis

3.1 Occurrence of lookups and equilibrium play

We first report the frequency of lookup patterns and Nash play of the two populations of subjects (Caltech and UCLA) in the three game classes (F, D1 and D2). The *occurrence of lookups* is the simplest measure of attention, and the most conservative one.²¹ Occurrence is a binary variable that takes value 1 if the box under scrutiny has been opened at least once during the game for at least one millisecond.²²

In the F class, subjects look at MIN 97% (Caltech) and 95% (UCLA) of the time. These subjects then play the equilibrium strategy with 99% (Caltech) and 96% (UCLA) frequencies. This suggests that both populations understand the fundamentals of the game: they compare the payoff of betting with the sure payoff and choose according to the largest of the two. More interestingly, the subjects who look at MIN and play Nash spend on average 57% (Caltech) and 60% (UCLA) of the time looking at MIN. But they look at all 7 payoffs of the game only 18% (Caltech) and 22% (UCLA) of the time. This suggests that subjects look at the payoffs strategically and succinctly. As mentioned before, near perfect Nash behavior and frequent disregard of other's payoff also provides support for the hypothesis of selfish maximization in this game.

Now we turn to the more interesting and challenging D1 and D2 games. Proportions of Nash equilibrium play are reported in Figure 6.

²¹Occurrence is a conservative measure in the sense that it does not imply any relation between the length of time people look at a number, and the extent to which the number is processed or used algorithmically. The only implied relation is that a number which is not seen at all cannot be used in decision making (except to the unobservable extent that people infer or imagine a value, rather than look it up.)

²²Given our software configuration, the minimal duration of a click is around 50ms. Although this may be too short for the eye and brain to perceive and process the information, we decided not to filter the data at all in an attempt to provide an upper bound on attention. (In contrast, filtering is unavoidable with eye-tracking because some durations are extremely short).

Caltech	Game 1			Game 2			Game 3			Game 4			Game 5		
	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
1	.66			.65			.76			.29			.72		
2		.35			.73			.79			.52				.13

UCLA	Game 1			Game 2			Game 3			Game 4			Game 5		
	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
1	.58			.39			.60			.32			.59		
2		.33			.46			.88			.39				.03

Figure 6 Empirical frequency of equilibrium play in each game. Bold **frequencies** are in cells where Nash predicts betting. Equilibrium play in class F is close to 1 and is omitted to aid visual clarity.

The proportion of equilibrium behavior is generally quite low, given that there are only two choices. Decisions differ across classes of situation and, to a lesser extent, across populations. The frequency of equilibrium behavior is quite different in each game, but there are no systematic differences between games in which subjects either are or are not supposed to bet.²³ For most of the analysis below, we will pool together all information sets within the same class. In order to distinguish luck from strategic reasoning, we separate the * cases where Nash equilibrium coincides with the naïve averaging strategy from the cases where equilibrium and averaging differ.

Table 2 provides basic statistics of the frequencies of equilibrium behavior (Nash) and lookups in the relevant payoff boxes (MIN). We categorize decisions in the attention and choice dimensions and report the fraction of observations of each type: MIN lookup and Nash play, MIN lookup but non-Nash play, no MIN lookup but Nash play and no MIN lookup and non-Nash play. We also add the conditional probability of equilibrium choice as a function of lookups (standard errors clustered at the individual level are reported in parentheses).

When averaging and equilibrium choices coincide (D1* and D2*), subjects make choices that are the Nash choices (though many likely reached these choices through averaging only) significantly more often than when they do not (rows 1 + 3), but it is poorly or even

²³This counters the suggestion that excessive betting in previous betting game experiments have been driven by the excitement of having a payoff that depends on some resolution of uncertainty. The symmetry in the lack of betting or no-betting when either is the equilibrium choice suggests other factors are driving the non-equilibrium choices.

% of observations	Caltech			
	D1	D2	D1*	D2*
MIN-Nash	.61 (.065)	.21 (.044)	.61 (.077)	.42 (.082)
MIN-notNash	.25 (.039)	.42 (.052)	.24 (.065)	.12 (.044)
notMIN-Nash	.02 (.008)	.06 (.020)	.10 (.052)	.37 (.090)
notMIN-notNash	.12 (.036)	.31 (.060)	.04 (.033)	.09 (.040)
Pr[Nash MIN]	.71 (.052)	.33 (.056)	.72 (.075)	.78 (.075)
Pr[Nash not MIN]	.15 (.060)	.15 (.049)	.70 (.208)	.81 (.086)
# observations	330	181	67	67

% of observations	UCLA			
	D1	D2	D1*	D2*
MIN-Nash	.43 (.056)	.16 (.038)	.44 (.069)	.51 (.072)
MIN-notNash	.39 (.048)	.37 (.051)	.39 (.055)	.07 (.036)
notMIN-Nash	.05 (.014)	.06 (.017)	.15 (.057)	.37 (.064)
notMIN-notNash	.13 (.044)	.41 (.061)	.02 (.017)	.05 (.027)
Pr[Nash MIN]	.53 (.057)	.30 (.057)	.53 (.065)	.88 (.062)
Pr[Nash not MIN]	.27 (.083)	.13 (.033)	.86 (.095)	.87 (.058)
# observations	458	278	82	91

Table 2: Occurrence of lookups and equilibrium play.

negatively related to whether they look at MIN (rows 5 and 6). It suggests that this class of situations mixes sophisticated (or cognitive) and naïve (or lucky) subjects.

From now on and unless otherwise stated, we will focus our attention on the more interesting case where the Nash and averaging strategies differ (D1 and D2). For those cases, the average probability of playing Nash is low in D1 (.63 and .48) and even lower in D2 (.27 and .22), but not far from previous results on the betting game (Rogers et al., 2010). Subjects look more often at MIN in D1 than in D2. This is quite natural, since the set of MIN payoffs is substantially smaller in the former (3 boxes) than in the latter case (5 boxes) and does not involve looking at a payoff the subject knows for sure will not realize. More interestingly, the frequency of equilibrium choice is 2-5 times higher when subjects look at the correct boxes than when they do not. This is in stark contrast to the D1* and D2* games where frequency of equilibrium choice does not differ depending on whether the MIN boxes are looked at or not. The frequency was even high for trials where the

subject did not look at the MIN boxes because just looking at own payoff and averaging leads to choices that coincide with equilibrium choices. Thus these differences in frequency of equilibrium choices condition on whether the MIN boxes were opened between D1/D2 games and D1*/D2* games suggest that lookups can be predictive of equilibrium choice.

Since there are few notMIN-Nash observations, all notMIN observations will henceforth be pooled together into one category. We can then determine whether the patterns of lookup and play is the same at Caltech and UCLA for each class of situations. A two-sample Wilcoxon rank-sum test shows that the difference in the distributions across populations are statistically significant both in D1 (at the 1% level) and in D2 (at the 5% level). Subjects in Caltech play closer to the theory than subjects in UCLA. However, the reasons are different in the two classes of situations (see Table 2). In D1, both populations look at MIN equally often but, conditionally on a correct lookup, Caltech subjects play Nash with substantially higher probability. In D2, Caltech subjects look correctly more often and, conditional on a correct lookup, they all play Nash roughly equally often.

Finally, we conduct a simple study of learning by dividing the sample into early play (first 20 trials of the sessions) and late play (last 20 trials), so that each subsample contains the set of 10 games played twice. Recall that at the end of each trial the history table contains the player’s choice, the other players choice, the realized state and the payoffs for both players. We find a consistent but modest increase in the frequency of MIN-Nash play and a modest decrease in the total lookup duration. Together they suggest that subjects develop some skills in determining which boxes are relevant, and that this improvement results in choices that are marginally closer to equilibrium (see Appendix section 7.2.1 for details). However, consistent with Sonsino et al. where the subjects did not converge to equilibrium play after hundreds of trials, the learning is also quite limited in our study and non-existent for some subjects.

3.2 Duration of lookups

This section associates choices with *duration of lookups*. Figure 7 summarizes the average time spent by subjects opening all boxes (in seconds), the average number of clicks, and the proportion of the total time in each box.²⁴ Subjects are classified by their lookup and choice in the three types mentioned above (MIN-Nash, MIN-notNash, notMIN). Because duration is quite similar across populations, we pool Caltech and UCLA subjects

²⁴The data is very similar if we record instead the percentage of clicks in each box. The average duration of a click varies a little between individuals but is relatively constant within individuals.

together.²⁵ Finally, the data is separated into D1 and D2 classes. As before, we express results as if the subject is always Player 1 in information set $\{A, B\}$.²⁶

D1	MIN-Nash [399]	MIN-notNash [259]	notMIN [130]
	A B C S	A B C S	A B C S
1	<u>.17</u> <u>.19</u> .06	<u>.19</u> <u>.18</u> .05	<u>.29</u> <u>.27</u> .09
2	<u>.20</u> .14 .09	<u>.16</u> .18 .10	<u>.03</u> .06 .06
	Total duration: 6.4s # clicks: 15.7	Total duration: 5.1s # clicks: 14.4	Total duration: 2.2s # clicks: 5.7
D2	MIN-Nash [82]	MIN-notNash [178]	notMIN [199]
	A B C S	A B C S	A B C S
1	<u>.13</u> <u>.12</u> <u>.13</u>	<u>.16</u> <u>.15</u> <u>.12</u>	<u>.22</u> <u>.20</u> <u>.01</u>
2	.16 <u>.20</u> <u>.17</u>	.15 <u>.16</u> <u>.14</u>	.17 <u>.12</u> <u>.09</u>
	Total duration: 11.4s # clicks: 23.6	Total duration: 7.9s # clicks: 20.7	Total duration: 3.7s # clicks: 9.3

Figure 7 Duration of lookups in D1 and D2 [# of observations in brackets].
Subject is Player 1 in $\{A, B\}$.

Let us focus first on the observations where subjects look at MIN (MIN-Nash and MIN-notNash). The relative time spent in the different boxes is remarkably similar when we compare types MIN-Nash and MIN-notNash within each class. The only difference is that subjects who play correctly spend marginally more time thinking throughout the game. Lookups, however, are vastly different between classes. Differences are expected since MIN is defined differently in the two classes. At the same time, it suggests that subjects do not look at all the payoffs and then think about how to play the game. Instead, they look economically, succinctly, and sequentially at the information they *think* will be most relevant for decision-making. More precisely, in D1 subjects barely look at payoffs in the state that cannot be realized ($[1C, 2C]$), whereas in D2 they spend almost as much time as in the other relevant boxes.²⁷ Total duration increases significantly from D1 to

²⁵For each subject and each class of games D1 and D2, there is no significant difference between the total duration distributions at Caltech and UCLA.

²⁶So, for example, box [2A] in game 1 is pooled together with box [1C] in game 4: in both cases it corresponds to a subject in a D1 situation and looking to the box of her rival in the full information set.

²⁷Because subjects in the D2 trials cannot know a priori that it is not a D1 game, it makes sense that they look at the [2A] box before realizing that this payoff is not sufficient for solving the game. The prediction is then that in D2 games, subjects will look at all 7 boxes.

D2, reflecting the fact that subjects realize the increased difficulty of the situation. Also and perhaps surprisingly, subjects do not spend more time looking at their own relevant payoffs than looking at their opponent’s relevant payoffs. Taken together, the results suggests that subjects who look at MIN do think carefully and strategically about the game. They do or do not play Nash, but this has more to do with their cognitive capacity to solve the equilibrium or to their expectations about others’ behavior than with an inability to understand that the game has some strategic elements.²⁸

Subjects who do not look at MIN (not MIN) spend very little time thinking about the game. In D1 they barely look at the payoffs of the other player, and spend very little time on their payoff in state C . From the lookup pattern, it seems that these individuals simply average their payoffs and compare it to the sure alternative (77% of the total looking time is spent on $[1A, 1B, S]$). As for D2, the distribution of lookups is quite similar to that of MIN types in D1. This means that some of these subjects realize the strategic component of the decision-making process. However, they fail to understand that their payoff in the state they know will not be realized is now crucial to unveil the equilibrium.

Transitions between boxes also can be used to judge the cognitive process (see Appendix section 7.1 for details). However, transitions are very closely linked to occurrences. The main feature is that for Nash play trials in D1 there are substantial transitions from player 1’s own payoff cells to the crucial $[2A]$ cell (which is consistent with the remarkable fact that MIN-Nash players in D1 actually look *more often* at $[2A]$ than at their own $[1A, 1B]$ cells).

3.3 Summary of aggregate analysis

The results obtained so far can be summarized as follows. (i) Nash equilibrium is chosen in 54% of D1 situations and 24% of D2 situations, and does not depend on whether the equilibrium prescribes ‘bet’ or ‘no bet.’ This result is broadly consistent with previous studies and reflects severe limits on strategic thinking. (ii) The likelihood of equilibrium behavior is two to almost five times greater among subjects who look at MIN, which suggests that lookup is an imperfect but reasonably good predictor of Nash choice. (iii) Individuals who look at MIN have similar duration and transition lookup patterns irrespectively of whether they end up playing Nash or not. It indicates that sorting out the correct information is important but not sufficient for equilibrium play. (iv) Some subjects spend little time on the games, exhibit non-strategic lookup behavior, and rarely play according to the theory.

²⁸Players do look at more than just MIN in a number of trials but this type of “overlooking” has been observed in other moustracking studies (Knoepfle et al., 2009; Wang et al., 2010). Since the looking is costless, such “overlooking” could be due to curiosity or limited memory.

(v) There is a modest but consistent increase in correct lookups and equilibrium behavior and a modest decrease in total lookup time over trials. However, behavior remains heterogeneous and far from equilibrium by the end of the sessions.

4 Cluster analysis

Research in experimental games have generally taken two approaches. One is the strategy followed in section 3, which consists in studying aggregate behavior. Another, which is more difficult and informative, is to do subject-by-subject type classification (Costa-Gomes et al. (2001) is a paradigmatic example).

In this section, we take an intermediate approach, which is to search for clusters of people (as in Camerer and Ho (1999)). The cluster approach includes the aggregate and subject-specific approaches as limiting cases. The aggregate approach implicitly assumes a single cluster is of interest. The subject-specific approach requires each subject to be in a singleton cluster.²⁹ One advantage of the clustering approach is that it creates statistical evidence on how well the extreme single-cluster and subject-specific approaches capture what is going on. This clustering technique is one of many ways to organize the data and allows us to see if the clusters correspond to particular types or rules specified by theory.

4.1 Clusters based on MIN lookup and choices

To find the clusters, we use the six aggregate statistics described in section 3 at the subject level: the three combinations of lookup and play (“MIN-Nash”, “MIN-notNash”, and “notMIN”) in the two classes of situations (D1 and D2). We then compute the percentage of trials in which each subject’s combination of lookup and choice is of either type in each class of situations. Each subject is thus measured by six percentages.

We group the 58 subjects of our experiment (see Table 1) in clusters based on the MIN-Nash, MIN-notNash, and notMIN percentages for both D1 and D2 games. There is a wide array of heuristic clustering methods that are commonly used but they usually require the number of clusters and the clustering criterion to be set ex-ante rather than endogenously optimized. Mixture models, on the other hand, treat each cluster as a component probability distribution. Thus, the choice between different numbers of clusters and different models can be made using Bayesian statistical methods (Fraley and Raftery, 2002). Popular heuristic approaches such as ‘k means clustering’ are equivalent to mixture models where a particular covariance structure is assumed.

²⁹Of course, these data are often combined in informal or formal cluster analyses, as we did in section 3 when we separated the analysis by type of lookup and play.

We implement model-based clustering analysis with the Mclust package in R (Fraley and Raftery, 2006). A maximum of nine components/clusters are considered for up to ten different models and the combination that yields the maximum Bayesian Information Criterion (BIC) is chosen. Specifically, hierarchical agglomeration first maximizes the classification likelihood and finds the classification for up to nine groups for each model. This classification then initializes the Expectation-Maximization (EM) algorithm which does maximum likelihood estimation for all possible model and number of clusters combinations. Finally, the BIC is calculated for all combinations with the EM-generated parameters.

For our multidimensional data, the model with ellipsoidal distribution, variable volume, equal shape, and variable orientation that yields *four clusters* maximizes the BIC at -1117 . Table 3 shows the frequencies of subjects from Caltech and UCLA in each cluster. Clusters are listed from high to low according to general frequency of Nash choice and MIN looking rates (see below for more details). A graph in Appendix section 7.4 shows a projection of the clusters into two of the six dimensions as a visual aid.

Cluster	Caltech	UCLA	Total
1	4 (17%)	3 (9%)	7 (12%)
2	6 (25%)	6 (18%)	12 (21%)
3	6 (25%)	12 (35%)	18 (31%)
4	8 (33%)	13 (38%)	21 (36%)
	$N = 24$	$N = 34$	$N = 58$

Table 3: Number of Caltech and UCLA subjects in each cluster.

As noted in section 3, the Caltech and UCLA populations appear to be distinct but not widely so. This is confirmed in the cluster composition of Table 3: there are subjects from each population in every cluster, but a greater percentage of Caltech subjects are in clusters 1 and 2 compared to UCLA subjects.

Table 4 gives frequencies of Nash play and MIN lookup across the clusters (standard errors clustered at the individual level are in parenthesis).

Clusters 1 and 2 almost always look at MIN. The difference lies in their choices: cluster 1 plays Nash most of the time whereas cluster 2 plays Nash in D1 but not in D2. Cluster 3 plays Nash slightly less than cluster 2 and looks at MIN substantially less in D2. Cluster 4, which is the largest, plays Nash infrequently and often does not look at MIN. In the next subsections, we explore in more detail these differences across clusters.

D1				
Cluster	MIN-Nash	MIN-notNash	notMIN	# obs.
1	.95 (.028)	.05 (.028)	.00 (n/a)	94
2	.73 (.069)	.23 (.069)	.04 (.014)	159
3	.57 (.048)	.38 (.043)	.05 (.014)	248
4	.18 (.039)	.43 (.059)	.39 (.061)	287
Level 3+	1	0	0	
Level 2	1	0	0	
Level 1	0	0	1	
D2				
Cluster	MIN-Nash	MIN-notNash	notMIN	# obs.
1	.58 (.071)	.27 (.031)	.15 (.070)	60
2	.21 (.051)	.71 (.056)	.08 (.026)	96
3	.15 (.034)	.44 (.032)	.41 (.058)	139
4	.04 (.012)	.20 (.058)	.76 (.058)	164
Level 3+	1	0	0	
Level 2	0	0	1	
Level 1	0	0	1	

Table 4: Proportion of types by cluster in D1 and D2.

Table 4 also gives the predicted proportions expected from the various level-k types (assuming best response). Some cluster percentages correspond somewhat closely to the level-k predictions, but others do not. The closest are cluster 1, which is closest to level 3, and cluster 4, which is closest to level 1.³⁰

Clusters 2 and 3 switch from playing Nash in D1 to not playing Nash in D2, as level 2 thinkers are predicted to do, but their looking patterns do not shift as sharply as predicted. Cluster 2 subjects look at payoffs like level-3 agents are supposed to do (usually looking at MIN) but they play more like level-2 agents (Nash in D1 and notNash in D2). Cluster 3 corresponds a little more closely to level-2: subjects consistently miss the payoff [1C] in D2 games and their overall pattern of look-up matches quite well what we described in Figure 4 as a level-2 player. However, conditional on looking at MIN, they choose Nash

³⁰Notice that level-1 predicts ‘not MIN’ but also ‘not Nash’ in both D1 and D2. Among the 41% of cluster 4 subjects who do not look at MIN in D1, 7% play Nash and 34% do not. Similarly, among the 77% of cluster 4 subjects who do not look at MIN in D2, 6% play Nash and 71% do not. This provides further support for identifying cluster 4 as level-1 players.

in D1 much less often than a level-2 player should. In Appendix section 7.5 we analyze strategic choice in games where Nash play coincides with naïve averaging of payoffs in the information partition. The proportion of Nash play in D1 and D2 by subjects in cluster 4 and in D2 by subjects in cluster 3 increases substantially, a result that lends further support to the level-1 and level-2 interpretation of these clusters.

There are interesting differences across clusters regarding the effect of experience on lookup and choices. (See Appendix section 7.2.2 for details). As before, we divide the sample into early trials (first 20) and late trials (last 20). We do not find evidence of learning by subjects in clusters 2 and 4. Subjects in cluster 1 show significant learning in complex D2 games (MIN-Nash rates increase from .40 to .71). Learning in simple D1 games does not occur simply because these subjects play close to equilibrium right from the outset. Subjects in cluster 3 learn in simple D1 games (MIN-Nash rates increase from .44 to .69) but not in complex ones.

There are also socio-demographic differences across clusters but the results are often weak in significance. (See Appendix section 7.3 for details). Subjects who are male, have some experience in game theory, or in poker or bridge, and who answer “cognitive reflection test” questions (Frederick, 2005) more accurately are more likely to be classified in cluster 1 and less likely to be in cluster 4. They ‘convert’ MIN looking compliance to Nash choices more often. These suggestive results are encouraging for further exploration of individual differences but are not strongly established in our data.

4.2 Lookup statistics within clusters

Lookup frequencies and durations within clusters offer some clues about the different choice patterns in those clusters. Some diagnostic summary statistics for D1 and D2 games are presented in Tables 5 and 6. The tables report average total duration (total looking in all cells) and average numbers of transitions. They also report ratios of the time spent looking at the crucial MIN cells containing other player’s payoffs (and their impossible own payoff [1C]), compared to the average time spent looking at one’s own possible payoff cells.

Subjects in cluster 1 look for a long total time (6166 msec) in D1 but, most importantly, increase their looking the most—almost doubling it—from D1 to D2 games. They also have other-own looking ratios greater than 1 in both types of games. For example, in D1 games they look for 1357 msec at the other’s payoff [2A], which is 1.28 times the average lookup duration at their own possible payoffs [1A, 1B]. The fact that they look *longer* at the MIN-set payoffs of other players than they look at their own payoffs (in game D2 as well as in game D1) is a clear informational marker of strategic thinking.

Recall that subjects in cluster 2 look at MIN payoffs often but only choose Nash 76% and 24% of time in D1 and D2 games. They do spend a lot of total time looking, and spend about a third more time on D2 games than on D1 games. They also make a lot of transitions (even more than cluster 1 players make), but the ratios of durations in other-payoff cells to own-payoff cells are around 1. Cluster 3 shows a similar pattern of looking to cluster 2 except they spend slightly less total time, make fewer transitions, and look at their impossible payoff [1C] only half as often as their own possible payoffs (that looking ratio is .42).

Finally, cluster 4 subjects looking patterns are entirely different than for the other three clusters. Their total durations and number of transitions are about half as large as for the other clusters, and they look in D2 games about as long as in D1 games. The ratios of looking time in other-payoff MIN cells compared to own-payoff cells are .57 and .56 in two cases, and a measly .19 for the impossible payoff cell [1C] in D2 games.

D1					
Cluster	avg[1A, 1B]	[2A]	Ratio $\frac{[2A]}{\text{avg}[1A,1B]}$	Total Duration	Number of Transitions
1	1060	1357	1.28	6166	11
2	1154	1113	.96	6603	15
3	1148	1218	1.06	6065	13
4	769	435	.57	3574	7

Table 5: Average lookup times in MIN cells (in milliseconds) and between-cell ratios by cluster in D1 games.

D2							
Cluster	avg[1A, 1B]	[2B]	Ratio $\frac{[2B]}{\text{avg}[1A,1B]}$	[1C]	Ratio $\frac{[1C]}{\text{avg}[1A,1B]}$	Total Duration	Number of Transitions
1	1419	2062	1.45	1691	1.19	11500	18
2	1226	1364	1.11	1081	.88	8390	18
3	1239	1281	1.03	524	.42	7615	15
4	842	468	.56	157	.19	3952	8

Table 6: Average lookup times in MIN cells (in milliseconds) and between-cell ratios by cluster in D2 games.

The data from F games can be also used to shed some light on the nature of information processing. Recall that in F games the search problem is trivial (players need only compare

one sure payoff $[1C]$ with the no-bet outside value $[S]$). Also, subjects in all clusters choose the Nash strategy virtually all the time (hence, very little variation in behavior). Finally, remember that the information about lookups in F games is *not* used for clustering subjects. The analysis of F games thus provides a nice template to study general properties and differences in search patterns across clusters.

Some lookup duration statistics are reported in Table 7. Subjects look at the MIN cells $[1C]$ and $[S]$ about 350-700msec each, which is much higher than the 100-150msec they look at non-MIN (“Other”) cells.³¹ The last column reports the fraction of total lookup that is fixated on the two MIN cells. This fraction is around .6, much higher than a random-looking baseline (which is $2/7 = .28$).

Cluster	$[1C]$	$[S]$	Other	Total	$\frac{[1C+S]}{\text{Total}}$
1	607	582	117	1773	0.67
2	659	502	278	2549	0.46
3	694	446	159	1937	0.59
4	642	351	106	1522	0.65

Table 7: Average lookup times in MIN cells by cluster in F games

Except for cluster 2, these results show that subjects are generally searching very economically in the simplest games. However, cluster 2 subjects do seem to look differently in these F games, “overlooking” at unnecessary non-MIN cells more frequently and longer. The mean duration at non-MIN (“Other”) cells is about twice as high (278 msec) and less than half of the total duration is directed to the two MIN cells.

In Appendix section 7.4 we perform an in-depth comparative study of occurrence and duration of lookups in each cell by cluster (on the lines of the analysis performed in section 3.2). The analysis reinforces the main results of this section. Also, the behavior of subjects in cluster 1 is particularly revealing. Indeed, their looking patterns are a portrait of rationality: they look longer at $[2A]$ in D1 and at $[2B]$ and $[1C]$ in D2 than they look at their own payoff cells $[1A, 1B]$.

4.3 Regression analysis: Predicting choice from specific payoff durations

The previous analyses cluster observations into groups based on MIN lookup compliance and choices. In this section we perform another type of analysis, namely we use lookup durations to predict choices (cf. Wang et al. (in press)).

³¹In fact, 14 out of the 15 median cell durations are 0 in clusters 1, 3 and 4.

Our first set of regressions are done at the subject level. For each subject, we use the following statistics: the percentage of Nash play, her average lookup duration in each box, her total lookup duration, and her total number of transitions. Table 8 presents OLS regressions using average lookup durations (in seconds) for those payoffs in the MIN sets which are most likely to be predictive of the tendency to make Nash choices. D1 and D2 games are analyzed separately since the relevant cells (and therefore those which are likely to be predictive) are different.³² Looking at [2A] is associated with an increase in the frequency of Nash play in D1. Looking at [2B] and [1C] is associated with an increase in the frequency of Nash play in D2. A one-second increase in looking at any of those payoffs correlates with an increase in the Nash play percentage of about .02. Those findings are consistent with the previous analysis: spending more time and attention on MIN results in a behavior closer to equilibrium play. For both classes D1 and D2, we tried alternative models and found no effect of individual difference measures.

	D1		D2	
	coefficient	st. error	coefficient	st. error
Total duration	-.0014	.0022	-.0056***	.0018
Duration in [2A]	.0256**	.0100	—	—
Duration in [1C]	—	—	.0186***	.0052
Duration in [2B]	—	—	.0282***	.0081
# of transitions	.7247	.6871	-.3490	.3291
Constant	29.57***	9.35	23.5***	6.12
# Observations	58		58	
R^2	0.20		0.34	

Table 8: OLS regression on Percentage of Nash play (* = significant at the 10% level, ** = significant at the 5% level, *** = significant at the 1% level).

The second set of regressions treats each trial as a separate observation. Probit regressions are used to predict whether the choice is Nash (= 1) or not (= 0). The results are summarized in Table 9. For D1, the time spent in [2A] marginally affects the probability of playing Nash ($p = .055$). For D2, increasing attention in [2B] made Nash play more likely. In D1 there is a significant positive effect of experience (match number) but there is only a weak effect in D2 games. These trial-specific results are consistent with the OLS regression, although the R^2 values are significantly smaller. The only difference is that looking at [1C] is not significant (as it is in the subject-level OLS regression). There are

³²The sign of the coefficient of total duration in D2 is not consistently negative over the OLS regressions we ran. The other significant coefficient reported in Table 8 keep the same signs in the alternative models.

no effects of individual difference measures.

	D1		D2	
	coefficient	st. error	coefficient	st. error
Total duration	1.2×10^{-5}	1.8×10^{-5}	-2.0×10^{-5}	2.1×10^{-5}
Duration in [2A]	$2.7 \times 10^{-4*}$	1.4×10^{-4}	—	—
Duration in [1C]	—	—	1.1×10^{-4}	0.7×10^{-4}
Duration in [2B]	—	—	$2.8 \times 10^{-4***}$	0.8×10^{-4}
Experience (trial #)	.0134***	.0047	.0054	.0067
Constant	-.4535***	.1682	-1.0979***	.1936
# Observations	788		459	
R^2	0.06		0.07	

Table 9: Probit regression on probability of Nash play (st. errors are clustered at the subject level) (* = significant at 10% level, ** = significant at 5% level, *** = significant at 1% level)

The probit regressions predicting choices on a trial-by-trial basis can also be done separately for all the clusters. Table 10 reports these regressions. The results are generally consistent with the signs and magnitudes of the aggregate analysis, but are much weaker in significance (perhaps due to modest sample sizes in each cluster). In D1 games, longer lookups at [2A] predicts a higher likelihood of Nash play in three of the four clusters. In D2 games, longer lookups at [2B] predicts a higher likelihood of Nash play only in cluster 3. In the Appendix section 7.4 some other regressions are presented where subsets of clusters are pooled together.

4.4 Are players optimally inattentive? Earning and learning in the four clusters

One important question we have postponed until now is whether, given the behavior of other subjects, it is actually optimal to play Nash strategies or to search for information widely. If others are unlikely to play Nash, then it clearly does not pay to play the Nash strategy in some cases (e.g., in game 1 non-Nash betting could be rational for player 2 in $\{B, C\}$ if she thinks player 1 will mistakenly bet in $\{A, B\}$). Whether this is true is an empirical question, which the data will answer. Furthermore, if it is not always optimal to play Nash, it raises the possibility that subjects who are trying to economize on search costs need not pay attention to all the payoffs in the MIN set. Indeed, attending to MIN payoffs could actually be harmful (in the sense of being correlated with lower earnings).

D1	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Total duration	$-2.4 \times 10^{-4***}$	-2.1×10^{-5}	$4.4 \times 10^{-5*}$	5.3×10^{-6}
Duration in [2A]	$9.2 \times 10^{-4***}$	$3.0 \times 10^{-4*}$	-1.9×10^{-5}	$4.8 \times 10^{-4**}$
Match #	-.020*	.016	.035***	-.002
Constant	2.6***	.24	-.67**	-.89***
# Observations	94	159	248	287
R^2	0.26	0.04	0.07	0.05
D2	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Total duration	-2.0×10^{-6}	-4.3×10^{-5}	-4.3×10^{-5}	9.6×10^{-5}
Duration in [1C]	2.9×10^{-4}	9.3×10^{-5}	-9.5×10^{-5}	-.0014**
Duration in [2B]	7.4×10^{-6}	2.9×10^{-4}	$3.7 \times 10^{-4***}$	2.3×10^{-4}
Match #	.030**	.0025	.0059	-.015
Constant	-.73*	-.89*	-1.0***	-1.3***
# Observations	60	96	139	164
R^2	.11	0.03	0.08	0.09

Table 10: Probit regression on Nash play by cluster (* = significant at 10% level, ** = significant at 5% level, *** = significant at 1% level).

The first empirical observation is that playing Nash is optimal in D1 games (averaging across the entire sample) but is *not* optimal in D2 games. One reason is that the average payoff in the information sets $\{A, B\}$ for player 1 and $\{B, C\}$ for player 2 are often well above the sure payoff $[S]$ when not betting is the Nash choice (and viceversa). So if you choose to bet and if others are betting often, then the average payoff can be above the $[S]$ payoff. Put differently, the payoffs could be designed so that deviations from Nash choice are more likely to be empirical mistakes (and consequently we would expect to see fewer of them). Because Nash play is optimal in D1 games but not in D2 games, it is possible that cluster 1 players, who look at MIN payoffs and play Nash most often in both types of games, may not earn the most money. Perhaps surprisingly, this is true.

Table 11 shows average expected normalized earnings for each of the four clusters. The normalization scales earnings such that a value 0 matches random choice and 1 matches empirical best response.³³ Standard errors are computed across all members of each cluster

³³To calculate expected earnings, we take the expectations over ex-ante probabilities of each state of the information set, to smooth out luck from ending up in a high paying state of a particular information set. We also take the expectation over all possible payoffs if paired with all subjects who are in the other player role with equal probability. This reduces variability introduced by the matching mechanism. The

so that every subject is treated as one data point. For comparison, the average expected normalized earnings of a subject who always played the Nash equilibrium strategy are shown in the row labelled "Nash".

	F		D1		D2	
	Earnings	st. error	Earnings	st. error	Earnings	st. error
Cluster 1	.96	.04	.79	.07	-.05	.13
Cluster 2	.95	.03	.48	.13	.45	.12
Cluster 3	.94	.02	.20	.11	.59	.08
Cluster 4	.95	.03	-.31	.07	.59	.08
Nash	1.0	.00	1.0	.00	-.28	.05

Table 11: Expected normalized earnings (st. errors by cluster and game type)

All clusters are close to optimal in F games, which is not surprising given the small proportion of mistakes in this class.

In D1 games the clusters are clearly ranked from 1 to 4 in earnings. Since the aggregate empirical best response is Nash, the clusters that choose Nash most often also earn the most (and a Nash player would obtain a perfect normalized earning of 1). Note that cluster 4 subjects do much worse than random, since they bet more than half the time and therefore lose a lot.

In D2 games, clusters 2, 3 and 4 earn similar amounts. From their earnings, they seem to have an understanding based on the game structure, and perhaps history, that sometimes it pays to play Nash and sometimes it does not. Perhaps surprisingly, cluster 1 subjects— who frequently play Nash— are worse than random, and worse than any other cluster including the clueless level-1 types in cluster 4 (although not as bad as a Nash player)! The reason is simple: In the level-k approach, for a fixed distribution of types, there is generally a level k^* at which players start to actually play suboptimally because their model of what other players will do starts to deviate from reality. Level k^* players think they face level k^*-1 opponents but they do not. This means— depending on the game structure— that their strategies are no longer empirical best responses. Thus, cluster 1 players represent an interesting case of people who analyze the game the most ‘rationally’

normalization uses random earnings for each subject (from choosing bet and no bet equally often) and earnings from empirically best response to the location-specific sample (which represents an upper bound). Each subject’s expected earnings in each type of game is then normalized by subtracting random earnings and dividing by the difference between the empirical best response earnings and random earnings.

but do not earn the most.³⁴ We can only speculate as to why the cluster 1 players did not shift to the non-Nash choices that would have lead to higher earnings. Perhaps they were hoping to teach other players to play the equilibrium strategies or perhaps they did not revise their erroneous beliefs about the other players' types.

Averaging across the three game types, subjects in cluster 2 earn the most. There are two possible interpretations of why the cluster 2's are doing so well in earnings which cannot be decisively separated. One interpretation is that cluster 2 players are "worldly" (Stahl and Wilson, 1995) in the sense that they can compute Nash equilibrium, but they also manage to figure out when it pays to play it and when it does not. The support for this view is that the cluster 2's look at MIN sets almost always (96% and 92% in D1 and D2 games). However, they have a much lower rate of making Nash choices in D2 games, given MIN lookups (only 25%) than cluster 1 subjects do (70%).

A different interpretation is that subjects in cluster 2 are not thinking that shrewdly, but are simply lucky, in a sense, because enough other players deviate from Nash choices that not playing Nash in D2 games is an empirical best response. Several pieces of evidence point in that direction: (i) they look for less time and make fewer transitions in D2 than the cluster 1 (even though they look longer than cluster 1 in D1 games³⁵); (ii) they do not learn with experience (if anything, they choose MIN-Nash more often by the end of the session, see Appendix section 7.2.2); (iii) they look at the impossible $[1C]$ payoff less than at their own possible payoffs; (iv) they oversearch much more than other clusters in the simple F games (see Table 7); and (v) they still get lower payoffs in D2 than the clueless cluster 4 subjects.

Another hypothesis that the data can say something about is the idea that players are "optimally inattentive"—i.e., it does not pay to look at a lot of payoffs, or to look at the MIN payoff sets, to choose the best strategies since other players are not necessarily playing Nash. This hypothesis is appealing but there is a lot of evidence against it.

First, looking is not actually very costly (in any conventional sense of consuming scarce resources with an opportunity cost). These student subjects are highly practiced using a mouse to retrieve information. Also, in F games where the players know they need to compare only two payoffs, they often look at some of the other cells (40% of their search is on other cells). This baseline tendency gives a rough way to calibrate the revealed cost of looking, and suggests it is low.

³⁴Note that this property does not hold in cognitive hierarchy approaches when it is assumed that level- k players have correct (truncated) beliefs. Higher thinking types have less truncation. Therefore, having increasingly more accurate beliefs, they are expected to actually earn the most.

³⁵See Tables 5 and 6 for the differences in looking.

Second, as noted earlier, learning appears to guide subjects to look at *more* of the MIN cells over time, not fewer (see Appendix section 7.2 for details on learning effects). If subjects were being optimally inattentive, or learning to be inattentive in the face of non-Nash play, there would be a decline in MIN looking occurrences (particularly in D2 games). The increase in MIN looking occurs for both subject pools and in both games (except a tiny drop in UCLA-D2). Since Nash strategies are played less often in D2 games, the fact that MIN looking *increases* or decreases very slightly over time in those games suggests they are not learning to look less in the face of non-Nash play. (And keep in mind that the matching protocol of subjects is random each period, so subjects cannot build up a subject-specific model which can be finely tuned.)

Third, the expected payoffs of subjects in each of the four clusters give a clear picture of the correlation between looking and earning. Combining the looking times in Tables 5 and 6 with the earnings statistics in Table 11 shows there is generally a positive relation between looking and earning. The largest cluster 4 (36% of the subjects) spends the least time and earns the least money. The other three clusters look a lot more, increase looking over time and earn more. Among clusters 1, 2 and 3, the secret to success seems to be looking at MIN payoffs almost always (as cluster 2 subjects do), but not necessarily playing the Nash choice in D2 games.

4.5 Summary of results

The main results of section 4 are as follows. Lookup patterns over relevant pieces of information and conversion of this information into Nash choice are heterogeneous. We performed a model-based clustering analysis on a combination of choice and lookup measures in D1 and D2 games, where the number of clusters and clustering criterion are optimized. The model groups the 58 subjects into four clusters. Some patterns of lookups help discriminate sharply between groups. For instance, the clusters differ by how much they look at the MIN boxes [2A] in D1 and [1C] and [2B] in D2. Also, subjects in clusters 1 and 2 increase the time spent before making a decision in D2 and also look frequently at [1C] in D2 games. Subjects in cluster 3 increase lookup from D1 to D2 less and look much less at [1C] in D2 games. Subjects in cluster 4 look the least and mostly concentrate on their own payoffs. The clusters map very roughly onto level-k thinking types. Cluster 4 corresponds to level-1 and cluster 1 corresponds to level-3. In theory, level-2 players should play MIN-Nash in D1 games and notMIN-notNash in D2 games. Cluster 2 (and to a lesser extent cluster 3) shows tendencies in this direction but not very sharply. Their combination of lookups and choices suggests a possible role for an element of stochastic choice and imperfect response as in QRE. There is some learning by subjects in cluster 1

(for the complex game) and in cluster 3 (for the simple game).

Regression analysis also show lookup duration in some MIN cells to be somewhat predictive of Nash choice. We also find differences in normalized expected payoffs for D1 and D2 games across clusters. Besides showing that longer lookup durations lead to greater payoffs, these differences offer greater insight into the processes linking lookup patterns and eventual choice.

5 Clustering based only on lookups

The analysis in the previous section clustered subjects according to *both* information search and play. A different analysis uses only the pattern of lookups of a subject to cluster them, then asks how well choices can be predicted from these lookup-only clusters. The immediate point of the analysis is to supplement the regressions reported in the last section, which show how well choices at the individual and trial level can be predicted from lookup statistics, with a cluster analysis. The long-run potential of this analysis is using clues about information processing as “leading indicators”, early in a choice process, to predict what choices might be made in the future, which could be quite distant.³⁶

The clustering and regression analysis in the last section suggest that predicting choices from lookups is possible to some extent. We therefore use four lookup statistics to cluster our population: (1) the difference in total lookup duration in D2 compared to D1; (2) the total lookup time in D2 spent at [1C]; (3) the total lookup time in D2; and (4) the percentage lookup time in D2 spent at [2B]. These statistics were chosen in the hope that (1) and (2) might discriminate between subjects in clusters 1 and 2 and subjects in clusters 3 and 4. Statistics (3) and (4) could then help discriminate between subjects in cluster 3 and subjects in cluster 4.³⁷

A two-step sequential clustering method leads to three main groups (and a single-subject fourth group which we ignore). In the first step only the first two measures (1)-(2) are used. That led to two clusters, which are group 1 below and a second cluster. Then measures (3)-(4) were used to separate the second cluster into groups 2 and 3 (a single step with all four measures at the same time did not generate reasonable clusters).³⁸

³⁶In these simple lab settings, information search and choices are close together in time so there is no great advantage in using information search to predict choices a few seconds later. However, in many naturally occurring settings, information search and choice can occur far apart in time. Potential applications include some of the most important choices people make, such as mates, education, career, buying houses, etc.

³⁷We tried different clustering models and we report only one in this section.

³⁸The fourth group subject is cross-classified as cluster 3 with a type profile closely resembling group 2. This subject is placed in a separate group because his or her total duration was much larger than those in

Table 12 presents for each group the proportion of observations that can be categorized as MIN-Nash, MIN-notNash and notMIN according to the patterns of lookup and choice. This table can be compared to Table 4 that clustered the subjects using both lookups and choices. If lookups were a perfect guide to choices, the frequencies reported in these two tables would match up closely. The results show that group 1 roughly corresponds to a combination of our previous clusters 1 and 2. These subjects look at the relevant information most of the time and play Nash in D1 but significantly less Nash in D2. Group 2 closely corresponds to cluster 3. Subjects play Nash in D1 but miss the relevant information in D2. Group 3 corresponds to cluster 4. These subjects do not pay attention to the relevant information and often do not play Nash in D1 or D2.

D1				
Group	MIN-Nash	MIN-notNash	notMIN	# obs.
Group 1	.71 (.053)	.21 (.038)	.08 (.044)	326
Group 2	.48 (.089)	.40 (.083)	.12 (.038)	211
Group 3	.24 (.051)	.42 (.047)	.33 (.064)	238

D2				
Group	MIN-Nash	MIN-notNash	notMIN	# obs.
Group 1	.31 (.055)	.50 (.045)	.20 (.044)	179
Group 2	.15 (.035)	.33 (.055)	.51 (.076)	123
Group 3	.05 (.022)	.29 (.083)	.66 (.086)	147

Table 12: Proportion of types by group in D1 and D2.

The main reason that the groups derived only from lookup statistics do not match the MIN-choice derived clusters more closely is because there are two distinct sets of subjects who look at the right information most of the time: cluster 1 in the earlier analysis who consistently play Nash in both games, and cluster 2 in the earlier analysis who play Nash much less often, especially in D2 games. These two clusters are lumped together into group 1 when only lookups are used. This is a reminder that lookups alone are correlated with choice, but do not perfectly predict choice. It also suggests an area for future research: Find other measures besides simple lookup statistics which can pin down which subjects

group 2. We tried other specifications, both sequential and simultaneous. Simultaneous clustering tends to generate a relatively large number of clusters with two or three main clusters and a series of small clusters (containing one or two subjects each). We conjecture this is due to the small sample. Also, some subjects have sometimes very similar patterns of look-ups but a subset of them have longer lookups in all boxes. The clustering method may classify them differently even though they are behaviorally similar.

will make particular choices, given the information they have attended to.

Table 13 reports for the three lookup-only groups which cluster these subjects belong to and what is their overall percentage of Nash play in D1 and D2 games. The groups preserve the ordering in the likelihood of Nash play although, naturally, the differences in choice frequencies are not as sharp as those resulting from the earlier cluster analysis which also included choices. The fact that there is any differentiation in choices among the three lookup-only groups means lookups have some predictive power in this type of cluster analysis (as shown in section 4 for individuals and trials). However, the fact that the choice frequency differences are much clearer in the lookup-plus-choice clusters means lookups have limited predictive power.

	Cluster				% Nash		
	1	2	3	4	D1	D2	All
Group 1 ($N = 23$)	6	8	7	2	.72	.47	.74
Group 2 ($N = 16$)	1	2	6	7	.53	.39	.65
Group 3 ($N = 18$)	0	2	4	12	.37	.31	.55
	$N = 7$	$N = 12$	$N = 17$	$N = 21$			

Table 13: Classification of subjects by cluster and group.

6 Conclusion

Having accurate models of human strategic thinking is important for explaining patterns in field data, for predicting new comparative statics, for designing institutional rules that lead to efficient outcomes when agents are boundedly rational, and perhaps for supplying lifelike simulated agents in various fields (e.g., agent-based social science). This paper contributes to development of such models by using a Mousetracking system to record information search in two-person games with private information about payoff-relevant states. Previous research has shown that subjects' behavior in these games departs substantially from Nash predictions and does not improve strongly with experience. These results are also replicated in our experiment, using a new set of games.

Besides finding that there are substantial deviations from Nash choice, the main conclusion is that there is heterogeneity among agents in their choice and lookup behavior. A small percentage (cluster 1) looks at all the payoffs required to make an optimal choice and chooses the equilibrium strategy quite often. A larger percentage (cluster 4) rarely looks at the required information and rarely choose the Nash strategy. A third cluster (cluster

3) looks at the required information and chooses the right strategy in simple games, but neither looks at MIN nor plays Nash in complex games. A last cluster (cluster 2) typically looks at the required information in both simple and complex games but chooses the equilibrium strategy only in the simple ones. The behavior of this last cluster clearly shows a limit in processing information, rather than in acquiring information per se. The fact that subjects in cluster 3 look and choose optimally in the simpler game but not in the game that requires more steps of inference also shows that game complexity is playing a role.

Two classes of theories can explain non-equilibrium choices. One theory emphasizes imperfect choice: subjects analyze the game fully but make inferential mistakes or believe others make mistakes (as in QRE or cursed equilibrium). Another theory emphasizes imperfect attention: heuristics or limits on cognition cause some subjects to ignore relevant information (as in level-k and cognitive hierarchy).

In the introduction, we noted that Rogers et al. (2010) found that CH and QRE models fit about equally well to their experimental data from 4-state betting games. There is some support for the imperfect attention approach in our analysis, since clusters 1 and 4 seem to correspond closely to level 3 and level 1 strategic thinkers, and cluster 3 is a relatively good candidate for level 2.

However, cluster 2 subjects are often looking at all the information required to do iterated thinking, but they are not drawing the right conclusion, as in the imperfect choice approach. This combination of looking and choice is what one might expect to see if, for example, a subject were trained to execute a QRE choice (i.e., they would look at all the payoffs, which is necessary to compute choice probabilities, but would not always optimize). Since QRE has never been specified as a theory of joint information search and choice, we cannot conclude that these clusters represent those looking patterns, but such an approach is promising.

At the same time, recall that in F games (where players only need to compare one certain betting payoff to the sure option $[S]$) subjects optimize 98% of the time. It is probably difficult for a QRE model to explain the failure to optimize in the D1 games *and* the near-perfect optimization in the F games with a single common parameter for response sensitivity. The best model would require a high parameter to fit the transparent easy choice (in F games), and a substantially lower parameter to fit the many non-Nash choices in more complex games requiring more analysis.

An interesting possible generalization is that cognitive difficulty increases implicit payoff imprecision (à la Van Damme and Weibull, 2002). Adding such a feature to QRE seems to be necessary to explain behavior in these games and could be a major advance for QRE

and related theories. (Of course, doing so could use some guidance from data on both choices and lookups, like those we are reporting here.)

Similarly, the simplest forms of cursed equilibrium do not fit the combination of lookups and choices especially well. The most congenial interpretation of cursed equilibrium is that there are two types of players. One group of players ignore information-action links, so they need not look at other players's payoffs at all, corresponding roughly to the empirical cluster 4. The other group of players expect other players to choose Nash (so they look at MIN payoffs). In our data, those two types of agents correspond roughly to clusters 4 and 1, but that specification leaves out clusters 2 and 3, which are half the subject pool. As with QRE, one can imagine modifications of cursed equilibrium in which degrees of cursedness are manifested by intermediate looking and choice patterns, but such a possibility is not part of the standard specification.

The heterogeneity of our lookup and choice data as revealed through the cluster analysis indicates a good fit with the levels of reasoning model. However, QRE and CE can be elaborated upon to explain both the lookup and choice data in future work. Furthermore, QRE, CE, and ABEE may be more appropriate for experimental contexts where convergence to equilibrium is more clearly observed as these are all equilibrium models. In our setting, there was limited learning and it is unlikely that any convergence can occur without more learning.

More generally, economists often talk casually about “contemplation costs” (Ergin and Sarver, in press), “control costs” (van Damme and Weibull, 2002), “thinking aversion” (Ortoleva, 2009) and cognitive difficulty. These costs are usually inferred from higher-order choices. The combination of choice and Mousetracking makes an empirically grounded approach to these topics. These implicit costs should, in principle, be linked to the “spending” of actual cognitive resources such as attention, information acquisition, time spent looking at information, transitions between pieces of information consistent with comparison and other mathematical operations, and so forth. Our view is that this exciting area of research cannot move forward merely by pure speculation about the nature of these processes without some direct measurement of attention. Mousetracking is one of the simplest such techniques.

7 Appendix: Additional analyses

The Appendix reports details of several analyses which are referred to in the text but are likely to be of secondary interest to most readers.

7.1 Transitions

Patterns of *lookup transitions* from one payoff cell to another can be diagnostic of computational algorithms. Figure 8 presents a graphical summary of the results, where the arrows in the matrix indicate transitions between betting payoffs and the outside arrows capture transitions to and from the sure payoff. We divide the analysis into D1 and D2 classes, pool both populations together, and separate the data into the usual three combinations of attention and choice.³⁹ We report percentage of transitions of each type, as well as the average number of transitions per match. The analysis can a priori be cumbersome since there are 28 combinations of transitions (including consecutive clicks in the same box). Fortunately, diagonal and non-adjacent transitions are rare, and so are two or more consecutive clicks in the same box. We report below only transitions that represent at least 5% of the total. With 5 to 9 of these transitions we can account for 88% to 93% of the observations.

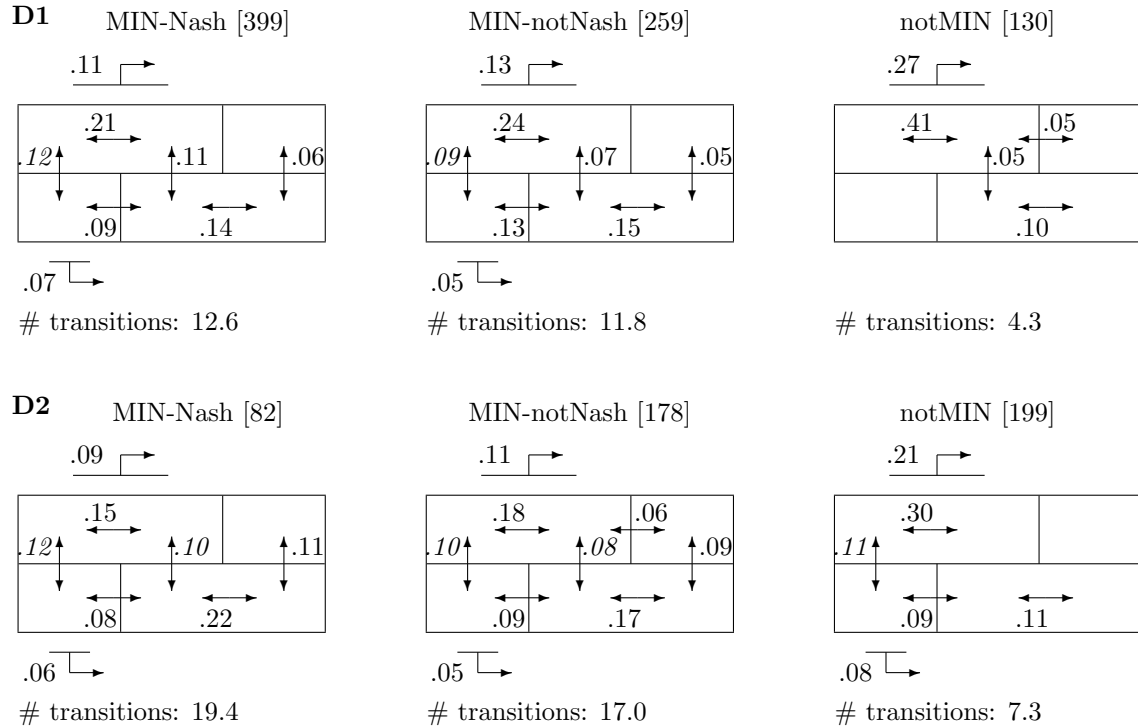


Figure 8 transitions of lookups in D1 and D2 [# observations in brackets]. *Italics* indicates between-player transitions needed for MIN. Subject is Player 1 in $\{A, B\}$.

³⁹ A comparison of the distributions of transitions at Caltech and UCLA reveals no significant differences.

Among subjects who look at MIN, the patterns are similar within a class whether they play Nash or not, and rather different between D1 and D2. NotMIN subjects in D1 spend most of the time in transitions between their possible payoffs ($[1A, 1B, S]$). In D2, they also transition to $[2A]$. Notice that transitions between the subject's own payoffs ($[1A] \leftrightarrow [1B]$) are relatively more frequent than between other boxes, even though Figure 7 suggests that duration is rather similar in all the relevant boxes. This is the only new insight relative to the duration analysis. Overall, in our games, there seems to be little information in the transition data that could not be deduced from the occurrence and duration data.

7.2 Experience

7.2.1 Learning at the aggregate level

This section contains results on learning effects at the aggregate level. In Table 14 we look at differences in lookups and choices between the first 20 and last 20 trials of a session. Recall that subjects play the same set of 10 games four times, so each subsample consists of all 10 games played twice (standard errors clustered at the individual level are in parentheses).

Caltech	% observations			
	D1		D2	
	1-20	21-40	1-20	21-40
MIN & Nash	.53 (.073)	.69 (.073)	.14 (.041)	.27 (.063)
MIN & not Nash	.30 (.053)	.19 (.041)	.45 (.066)	.40 (.066)
Not MIN	.17 (.047)	.12 (.044)	.41 (.070)	.33 (.080)
# observations	174	156	85	96

Caltech	mean duration (in seconds)			
	D1		D2	
	1-20	21-40	1-20	21-40
MIN & Nash	7.2 (1.34)	5.9 (1.07)	9.4 (1.09)	9.8 (1.60)
MIN & not Nash	6.4 (0.85)	4.9 (1.07)	8.4 (1.34)	8.0 (1.43)
Not MIN	2.4 (0.60)	2.0 (0.59)	4.7 (0.80)	2.8 (0.56)
# observations	174	156	85	96

UCLA	% observations			
	D1		D2	
	1-20	21-40	1-20	21-40
MIN & Nash	.38 (.060)	.48 (.065)	.14 (.038)	.18 (.049)
MIN & not Nash	.40 (.051)	.37 (.057)	.41 (.058)	.33 (.058)
Not MIN	.21 (.052)	.15 (.051)	.46 (.064)	.49 (.082)
# observations	221	237	138	140

UCLA	mean duration (in seconds)			
	D1		D2	
	1-20	21-40	1-20	21-40
MIN & Nash	7.6 (0.89)	5.5 (0.56)	13.7 (2.76)	12.3 (1.97)
MIN & not Nash	5.2 (0.62)	4.3 (0.67)	7.6 (1.57)	8.0 (1.52)
Not MIN	2.7 (0.42)	1.6 (0.30)	4.4 (0.92)	3.1 (0.44)
# observations	221	237	138	140

Table 14: Effect of experience on lookups and behavior.

In every subject pool and game except one, the frequency of Min-Nash increases from early to late trials and the frequencies of the other two classifications falls (the exception is UCLA-D2 notMIN). The likelihood of MIN-Nash increases proportionally between 26.3% and 92.9% for all treatments. Using a two-sample Wilcoxon signed-rank test to test for equality of distributions of types, behavior in D1 are statistically different both in Caltech and UCLA (at $p < .01$) and in D2 at Caltech only (at $p < .02$).⁴⁰

We also observe a modest decrease in total lookup time in D1. A two sample Wilcoxon signed-rank test of difference in total duration between the first 20 and the last 20 trials shows a significant decrease in D1 ($p < .001$, pooling both groups) but not in D2.⁴¹ Overall, the combination of an increase in MIN lookups and a decrease in total time suggests that subjects develop some skills in sorting out the relevant boxes, and this translates into small improvements in behavior. However, low proportions of equilibrium play and significant heterogeneity remain at the end of each session, especially in the more complex D2 games.

7.2.2 Learning at the cluster level

This section contains results on learning effects by cluster. Table 15 reports learning statistics in clusters 1-4. As before, we divide the data between early trials (1 to 20) and late trials (21 to 40). In D1, only subjects in cluster 3 benefit from experience and increase both their MIN lookup and Nash play rate substantially. Clusters 1 and 2 know how to play that game from the beginning and their performance is not varying over time. Subjects in cluster 4 do not learn at all. They do not increase their MIN lookup rate over time and do not learn how to convert MIN lookups into Nash play. In D2, the frequency of Nash play rises over time in cluster 1. All other clusters do not seem to play better with experience.

Overall, subjects in clusters 2 and 4 have no significant learning trends. Subjects in cluster 1 learn in complex games while subjects in cluster 3 only learn in simple games. These observations are confirmed by two-sample Wilcoxon signed rank tests. When comparing the distributions of types in the 20 early trial and the 20 last trials, we find that the differences are significant for cluster 3 in D1 and for cluster 1 in D2.⁴² One should notice however that the number of observations is relatively small.

⁴⁰More precisely, the p-values in D1 are 0.0007 and 0.0065 for Caltech and UCLA respectively. The p-values in D2 are 0.0164 and 0.4532 for Caltech and UCLA respectively.

⁴¹More precisely, in D1 the p-value are 0.1643 (Caltech) and < 0.001 (UCLA). It is < 0.001 if we combine the populations. In D2, the p-values are 0.9255 (Caltech) and 0.2914 (UCLA).

⁴²The respective p-values are < 0.001 and 0.0578.

	Cluster 1		Cluster 2		Cluster 3		Cluster 4	
	1-20	21-40	1-20	21-40	1-20	21-40	1-20	21-40
D1								
MIN-Nash	.94 (.03)	.95 (.03)	.68 (.07)	.78 (.08)	.44 (.05)	.69 (.07)	.15 (.03)	.22 (.06)
MIN-notNash	.06 (.03)	.05 (.03)	.27 (.07)	.20 (.08)	.47 (.04)	.30 (.06)	.43 (.07)	.42 (.07)
notMIN	.00 (—)	.00 (—)	.05 (.02)	.02 (.02)	.09 (.03)	.01 (.01)	.42 (.07)	.36 (.08)
# obs.	53	41	78	81	119	129	145	142
D2								
MIN-Nash	.40 (.11)	.71 (.10)	.14 (.05)	.28 (.06)	.13 (.05)	.18 (.05)	.06 (.03)	.01 (.01)
MIN-notNash	.36 (.05)	.20 (.05)	.73 (.08)	.68 (.05)	.41 (.06)	.47 (.06)	.26 (.07)	.15 (.07)
notMIN	.24 (.10)	.09 (.08)	.12 (.05)	.04 (.03)	.46 (.08)	.35 (.07)	.68 (.07)	.84 (.07)
# obs.	25	35	49	47	71	68	78	86

Table 15: Effect of experience on behavior by cluster.

7.3 Individual differences

In this section we look at the demographic information we collected through the Questionnaire for each subject and ask whether these factors account for the individual differences in choices. Table 16 reports, for each population (Caltech and UCLA) the following statistics: gender, exposure to game theory (0 class = ‘no exposure’ and 1 class or more = ‘exposure’), experience with Poker or Bridge (never or rarely = ‘no experience’ and occasionally or often = ‘experience’), and the number of correct responses in each of the CRT questions. Numbers in parentheses report percentages of the sample. Note that Caltech subjects differ from UCLA subjects only in the CRT test.

	Caltech ($N = 24$)	UCLA ($N = 34$)
Females	9 (37.5%)	17 (50.0%)
Exposure to Game Theory	9 (37.5%)	13 (38.2%)
Experience in Poker or Bridge	8 (33.3%)	11 (32.3%)
Correct CRT 1	21 (87.5%)	14 (41.2%)
Correct CRT 2	21 (87.5%)	9 (26.4%)
Correct CRT 3	21 (87.5%)	11 (32.3%)

Table 16: Summary of Questionnaire answers.

Table 17 groups CRT answers for each subject into all correct (‘Correct’) and at least one wrong (‘Wrong’). It then reports the fraction of times that looking at MIN was translated into Nash play (standard errors clustered at the individual level are reported in parentheses). Male subjects, subjects who have been exposed to Game Theory, subjects with experience in Poker or Bridge, and subjects who answer correctly all three CRT questions seem to perform slightly better than their counterparts.⁴³ A variety of two-sample Wilcoxon rank-sum tests show that

⁴³The fraction of times subjects do not look at MIN is stable across genders, exposure to game theory

the distributions of the probabilities of playing Nash conditional on looking at MIN are indeed statistically different in some cases. In particular in D1, there are gender, game theory, and CRT effects.⁴⁴ However, differences are not significant in D2.

	Pr[Nash MIN]			
	Gender		Game Theory	
	Female	Male	No Exposure	Exposure
D1	.52 (.066)	.68 (.051)	.54 (.053)	.71 (.062)
# observations	297	361	403	255
D2	.21 (.054)	.39 (.050)	.26 (.051)	.39 (.058)
# observations	109	151	152	108
	Poker/Bridge		CRT	
	No Experience	Experience	Wrong	Correct
D1	.58 (.047)	.65 (.083)	.49 (.050)	.75 (.056)
# observations	443	215	368	290
D2	.29 (.048)	.37 (.069)	.26 (.057)	.37 (.053)
# observations	172	86	136	124

Table 17: Effect of Questionnaire answers.

There are also a few interesting features when we look at the Questionnaire by cluster. Clusters 2, 3 and 4 have high numbers of female subjects, subjects who have not been exposed to Game Theory, do not play Poker or Bridge or did poorly in the CRT test. This is reported in Table 18. Cluster 1 by contrast is a cluster made of male subjects who took Game Theory classes and answered correctly all CRT questions. This is suggestive but it is again difficult to make strong conclusions given the limited number of observations.

Last, we also analyzed the relationship of durations with the answers to the Questionnaire and we found no impact of gender, game theory background or poker expertise. Interestingly, two-sample Wilcoxon rank-sum tests for the differences in distribution of duration lookups reveals that subjects who do answer correctly to CRT questions have different look-up patterns in D2.⁴⁵

and experience in poker/bridge. However, subjects who answer correctly CRT questions tend to look at MIN slightly more often in both D1 and D2.

⁴⁴When testing for the differences in distributions, the p-values are 0.0302, 0.0467 and 0.0012 for gender, game theory and CRT answers respectively.

⁴⁵When testing for the difference between the distributions of total durations in D2, the p-value is 0.0020. When testing for the difference between the distributions of durations in box [1C] in D2, the p-value is 0.0795. Interestingly, effects are different across questions. We ran the same tests question by question and we obtained consistently significant results for total durations in D1 and D2, as well as [1C] in D2 for the first question. Results were often not significant for the other questions.

Cluster	Gender		Game Theory	
	Female	Male	No Exposure	Exposure
1	1 (4%)	6 (19%)	2 (6%)	5 (23%)
2	7 (27%)	5 (16%)	8 (22%)	4 (18%)
3	7 (27%)	11 (34%)	10 (28%)	8 (36%)
4	11 (42%)	10 (31%)	16 (44%)	5 (23%)
	$N = 26$	$N = 32$	$N = 36$	$N = 22$

Cluster	Poker/Bridge		CRT	
	No Experience	Experience	Wrong	Correct
1	4 (10%)	3 (16%)	1 (3%)	6 (24%)
2	8 (21%)	4 (20%)	5 (15%)	7 (28%)
3	12 (31%)	6 (32%)	12 (36%)	6 (24%)
4	15 (38%)	6 (32%)	15 (46%)	6 (24%)
	$N = 39$	$N = 19$	$N = 33$	$N = 25$

Table 18: Frequencies based on Questionnaire answers.

7.4 Lookup occurrence and duration within clusters

This section contains analysis similar to that in section 3.2 except that we consider only subjects in the same cluster.⁴⁶

For illustrative purposes, we first present in Figure 9 a scatterplot that depicts the clusters. Each data point corresponds to one subject. While the clustering is done using six statistics (the three types of lookup and play in each of the two game classes), the plot uses only the two statistics which illustrate the clustering most sharply in two dimensions. The x-axis is the percentage of trials on which the subject looked at MIN but did not play Nash in D2 games. The y-axis is the percentage of trials on which the subject looked at MIN and played Nash in D1 games.

Cluster 1 is represented by a dark blue filled square, cluster 2 by a light blue hollow square, cluster 3 by a red filled circle and cluster 4 by a green hollow circle. Clusters 1 and 3 are the tightest, reflecting low within-cluster variation.⁴⁷

Figure 10 shows average number of transitions, total time and percentage of time in each cell for subjects in cluster 1. As before, while subjects switch roles between player 1 and 2, we express everything from the point of view of player 1 in information set $\{A, B\}$ (see section 3.2 for details). For this cluster, the looking patterns are a portrait of rationality: they look a lot at the crucial $[2A]$ cell in D1 games (longer than they look at their own payoff cells). They also look a lot at $[2B]$ and $[1C]$ in D2 games compared to their looking rates in D1 games (and, again, longer than

⁴⁶Statistics are obtained from the average data collected at the individual level. However, some subjects are more frequently exposed to a given situation (F, D1, D2). To account for this, we computed the same statistics from the trial-by-trial dataset. Results were both qualitatively and quantitatively similar.

⁴⁷Note that the algorithm produces some apparent misclassifications— for example, there are red circles (cluster 3 subjects) inside the ellipse for cluster 4. This is because the algorithm is clustering using four other statistics that are not plotted.




Figure 9 Four clusters of 58 subjects.

they look at their potential payoffs $[1A, 1B, S]$.⁴⁸ When they do look correctly, they convert their information into Nash choices at high rates (although lower in D2).

D1 [94]	A	B	C	S	D2 [60]	A	B	C	S
1	.17	.17	.10	.14	1	.13	.12	.15	.10
2	.22	.12	.08		2	.16	.18	.17	
	Total duration: 6.2s					Total duration: 11.5s			
	# transitions: 11					# transitions: 18			

Figure 10 Lookup behavior by subjects in cluster 1 [# observations in brackets].

A key characteristic of the lookups in this cluster is the increase in duration from D1 to D2. When the situation is more complex, subjects pay more attention in general. Subjects also shift attention to the state they know cannot realize ($\{C\}$) which is key in D2 but irrelevant in D1 (they spend, on average, 1.7 seconds in D2 and 0.6 seconds in D1).

Figure 11 shows lookup statistics for cluster 2. Recall that their MIN percentages are very high in both games (above 90%), but they only play Nash 76% and 24% of the time. Put it simply,

⁴⁸Because of the extremely high compliance with Nash and MIN, this cluster is useful in a special way. One method for classifying types is to train subjects to execute a particular algorithm, then use their looking patterns as a filter to determine whether other subjects are “looking rationally.” Johnson et al. (2002) were the first to do this in a casual way (simply suggesting backward induction) and Costa-Gomes et al. (2001) did so more precisely. Comparing regular subjects’ looking to looking by trained subjects controls for many types of noise and idiosyncrasy (e.g., forgetting which requires repeated lookups, particular transitions, etc.). But these cluster 1 subjects already provide a strong and possibly more natural template in D1 (where choice and MIN are almost perfect) and also a strong template for MIN in D2.

they look at payoffs like cluster 1 subjects, but choose Nash at rates more like cluster 3 subjects (see Table 4 and the discussion below).

D1 [159]	A	B	C	S		D2 [96]	A	B	C	S
1	.17	.18	.07	.14		1	.15	.14	.13	.13
2	.17	.15	.12			2	.15	.16	.14	
	Total duration: 6.6s					Total duration: 8.4s				
	# transitions: 15					# transitions: 18				

Figure 11 Lookup behavior by subjects in cluster 2 [# observations in brackets].

The main and perhaps only significant difference between Figures 10 and 11 is that subjects in cluster 2 do not increase lookup duration between D1 and D2 as much as subjects in cluster 1 (27% increase for cluster 2 for 86% increase for cluster 1). Cluster 2 also shows what can be learned from using choices and lookups together. Looking only at the Nash compliance rate, one might have guessed that these subjects were not attending to the correct cells of the payoff matrix; but that guess is wrong. Looking only at the MIN compliance rate, one might have guessed that these subjects would play Nash most of the time; but that guess is wrong too.

Figure 12 shows lookup statistics for cluster 3. Recall that their MIN rates are very high in D1 games, which is consistent with a long lookup duration in the crucial cell [2A]. However, while they look at the required information frequently and in similar proportions, they convert it to Nash choice at somewhat lower rates than cluster 2 (58% compared to 72%). In D2 games, subjects do not realize it is necessary to look at [1C]. This, together with a lower total lookup duration, is the main attentional difference between these individuals and subjects in clusters 1 and 2.

D1 [248]	A	B	C	S		D2 [139]	A	B	C	S
1	.17	.21	.05	.12		1	.18	.15	.07	.12
2	.20	.17	.09			2	.17	.17	.15	
	Total duration: 6.1s					Total duration: 7.6s				
	# transitions: 13					# transitions: 15				

Figure 12 Lookup behavior by subjects in cluster 3 [# observations in brackets].

Finally, Figure 13 shows lookup statistics for cluster 4. Recall that their MIN percentages are modest and low (61% and 24%) and Nash choice percentages are very low too (25% and 10%). The duration and number of transitions are about half of those in the other clusters in D1, and they do not increase in D2 games. This is indeed a paradigmatic case of non-strategic behavior: lookups are virtually identical in D1 and D2, and the patterns suggest that subjects concentrate their attention almost exclusively on their own payoffs (for example, they spend on average 0.16 seconds in [1C], the crucial box in D2 games). It also illustrates the importance of clustering. In section 3.2, we argued that lookup patterns in D2 among subjects who do not look at MIN resembled lookups in D1 among subject who look at MIN. It turns out that this is true for cluster 3 but not for cluster 4, where duration is much shorter and lookups at the other player's boxes (including [2A]) are rare.

D1 [287]		A	B	C		S		D2 [164]		A	B	C		S	
	1	.22	.21	.04		.19			1	.22	.21	.04		.21	
	2	.12	.14	.07					2	.13	.12	.08			
		Total duration: 3.6s								Total duration: 4.0s					
		# transitions: 7								# transitions: 8					

Figure 13 Lookup behavior by subjects in cluster 4 [# observations in brackets].

In section 4.3 we performed a regression analysis where we used different lookup measures to predict choice. The analysis was done at the subject level, at the trial level pooling all subjects, and at the trial level for each cluster separately. Table 19 reports some new probit results obtained at the trial level when subsets of clusters are pooled together (standard errors are clustered at the subject level). Note first that we include dummy variables for clusters. For instance, when we pool clusters 1 and 2 together, the cluster 2 dummy (‘Cluster 2’) takes value 1 for observations from subjects in cluster 2 and 0 otherwise. We find that, in all regressions, the dummy(ies) is(are) significant. Belonging to a cluster is a good indicator of Nash play. A few interesting insights on typical lookups should also be noted. When pooling clusters 1 and 2, an increased lookup at [2A] is significant in D1. This is driven by the fact that subjects in cluster 1 always look and play Nash while subjects in cluster 2 do not. Among those, paying marginally more attention to [2A] leads to more Nash play. When pooling clusters 1 and 3 or 1, 3 and 4, a long lookup in [2B] in D2 increases the likelihood to play Nash. Typically, subjects in cluster 1 do look at [2B] while a large proportion of subjects in cluster 3 and 4 do not. A closer look at D2 situations shows that increasing attention to [1C] is significant in D2 when we pool clusters 1, 3 and 4 or 1 and 4. This is the case because subjects in cluster 1 always look at [1C] and play Nash while subjects in clusters 3 and 4 do not look or play Nash. In all other combinations, a large fraction of subjects may look at [1C] and may not understand what to make of it. This explains why the aggregate analysis could not pick any effect of a long lookup in [1C]. This also suggests that lookup is an imperfect predictor of choice. However, a combination of specific lookups is a reasonably good predictor. Last, the answers to the Questionnaire have again little effect.⁴⁹

D1	Clusters		
	1+2	1+3	1+3+4
Total Duration	$-.44 \times 10^{-4}$	1.2×10^{-5}	1.4×10^{-5}
Duration [2A]	$4.0 \times 10^{-4**}$	6.7×10^{-5}	1.2×10^{-4}
Match #	.014*	.029***	.014**
Cluster 2	-.84**	-	-
Cluster 3	-	-1.43***	-1.4***
Cluster 4	-	-	-2.2***
Constant	1.18***	.95***	1.1***
# Observations	253	342	629
R^2	.11	.17	.22

⁴⁹The coefficients for the dummy variables gender, experience in poker or bridge and correct answers the CRT tests are significant when combining clusters 1 and 2 in D1.

D2	Clusters				
	1+2	1+3	1+3+4	1+3+4	1+4
Total Duration	-2.5×10^{-5}	-3.7×10^{-5}	-2.2×10^{-5}	-2.0×10^{-5}	4.3×10^{-6}
Duration [1C]	1.1×10^{-4}	6.1×10^{-5}	1.6×10^{-5}	$2.1 \times 10^{-4*}$	$4.2 \times 10^{-4**}$
Duration [2B]	1.8×10^{-4}	$2.8 \times 10^{-4**}$	$2.5 \times 10^{-4***}$	$2.6 \times 10^{-4**}$	1.4×10^{-4}
Match #	.013	.012	.0036	.0048	.011
Cluster 2	-.96***	-	-	-	-
Cluster 3	-	-1.0***	-1.1***	-	-
Cluster 4	-	-	-1.4***	-	-
Constant	-.17	-.13	.010	-1.1***	-1.4***
# Observations	156	199	363	363	224
R^2	.15	.17	.19	.09	.17

Table 19: Probit regression on Nash play by groups of clusters (* = significant at the 10% level, ** = significant at the 5% level, *** = significant at the 1% level).

7.5 Predicting choice in other games

Most of the analysis in the main text disregarded the games in which the Nash solution coincides with the naïve strategy of comparing [S] with the average of payoffs in the information set. In this section, we analyze the performance in those games for subjects in each cluster. Results are reported in Table 20 (standard errors clustered at the individual level are reported in parenthesis).

Cluster	D1				D2			
	MIN-Nash	MIN-notNash	notMIN	# obs.	MIN-Nash	MIN-notNash	notMIN	# obs.
1	.68 (.114)	.32 (.114)	.00 (n/a)	22	.75 (.115)	.05 (.050)	.20 (.086)	20
2	.54 (.116)	.46 (.116)	.00 (n/a)	26	.65 (.121)	.24 (.091)	.11 (.084)	37
3	.58 (.106)	.40 (.098)	.02 (.022)	45	.57 (.078)	.05 (.031)	.39 (.073)	44
4	.39 (.077)	.20 (.050)	.41 (.087)	56	.18 (.047)	.04 (.024)	.79 (.045)	57

Table 20: Proportion of types by cluster in D1 and D2 when Nash and naïve coincide.

Conditional on looking at MIN, subjects in cluster 4 play Nash significantly more often in both D1 and D2. We argued that these subjects are level-1. Therefore, they are expected to perform poorly in situations where the averaging strategy does not coincide with Nash play, and to perform well when it does. Our result simply confirms this hypothesis.⁵⁰

Subjects in cluster 3 perform in a similar fashion in D1 but they play Nash more often in D2. Indeed, it may be the case that those subjects realize that D2 situations are more complex and, in the absence of a good guiding principle, they sometimes end up using the averaging shortcut strategy. A similar pattern can be observed for cluster 2 in D2.

⁵⁰Cluster 4 subjects also play Nash significantly more often in these games when they do not look at MIN (data not reported for the sake of brevity).

Finally, the results for cluster 1 are surprising: subjects perform less well when they look at MIN in D1 (a shift from MIN-Nash to MIN-notNash) and they tend to look at MIN less often in D2 (a shift from MIN-notNash to notMIN). We also observe a similar shift in D1 for subjects in cluster 2. It is difficult to formulate a clear hypothesis for such behavior and the statistics above rely on a small number of observations. However, it is striking that the most careful types tend to play non-Nash more often when a simple heuristic leads to Nash play.

References

1. Angrisani, M, Guarino, A, Huck, S., and N. Larson (2008). No-trade in the laboratory. Working paper, ELSE.
2. Arieli, A., Ben-Ami, Y., and A. Rubinstein (2009). Fairness Motivations and Procedures of Choice between Lotteries as Revealed through Eye Movements. Working Paper.
3. Armel, K.C. and A. Rangel (2008). The impact of computational time and experience on decision values. *American Economic Review*, 98, 163-168.
4. Bazerman, M.H. and W.F. Samuelson (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, 27, 618-634.
5. Brown, A.L., Camerer, C. and D. Lovo (2009). To review or not to review? Limited strategic thinking at the box office. Working paper.
6. Cai, H., Wang, J. (2006) Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56, 7-36.
7. Camerer, C. (2003). *Behavioral game theory: Experiments on strategic interaction*. Princeton University Press, Princeton.
8. Camerer, C. (2008). The neural basis of strategic thinking. In Glimcher, Camerer, Fehr, Poldrack (Eds.), *Neuroeconomics*, Elsevier.
9. Camerer, C. and T. Ho (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67, 827-874.
10. Camerer, C., Ho, T. and J. Chong (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119, 861-898.
11. Camerer, C., E. Johnson, T. Rymon, S. Sen (1993). Cognition and framing in sequential bargaining for gains and losses. *Frontiers of Game Theory*. K. G. Binmore, A. P. Kirman and P. Tani. Cambridge, MIT Press: 27-47.
12. Capen, E. C., R. V. Clap, and W. M. Campbell (1971). Competitive Bidding in High-Risk Situations. *Journal of Petroleum Technology*, 23, 641-653.
13. Carrillo, J.D. and T. Palfrey (2008). No Trade. Working paper.
14. Carrillo, J.D. and T. Palfrey (2009). The Compromise Game: Two-Sided Adverse Selection in the Laboratory. *American Economic Journal: Microeconomics*, 1, 151-181.

15. Charness, G. and D. Levin (2009). The Origin of the Winner's Curse: A Laboratory Study. *American Economic Journal: Microeconomics*, 1, 207-236.
16. Charness, G. and M. Rabin (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 117, 817-869.
17. Coricelli, G. and R. Nagel (2009). Neural correlates of depth of reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Science*, 106, 9163-9168.
18. Costa-Gomes, M. and V. Crawford (2006). Cognition and behavior in two-person guessing games: an experimental study. *American Economic Review*, 96, 1737-1768.
19. Costa-Gomes, M., V. Crawford, and B. Broseta (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69, 1193-1235.
20. Costa-Gomes, M. and G. Weiszäcker (2008). Stated beliefs and play in normal form games. *Review of Economics Studies*, 75, 729-762.
21. Crawford, V. (2003). Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions. *American Economic Review*, 93, 133-149.
22. Crawford, V. (2008). Look-ups as the windows of the strategic soul: Studying cognition via information search in game experiments. In A. Caplin and A. Schotter (Eds.), *Perspectives on the Future of Economics: Positive and Normative Foundations*, Oxford University Press.
23. Crawford, V. and N. Iriberry (2007). Level-k auctions: can a non-equilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica*, 75, 1721-1770.
24. van Damme, E. and J. Weibull (2002). Evolution in games with endogenous mistake probabilities. *Journal of Economic Theory*, 106, 296-315.
25. Ergin, H. and T. Sarver (in press). A unique costly contemplation representation. *Econometrica*.
26. Eyster, E. and M. Rabin (2005). Cursed Equilibrium. *Econometrica*, 73, 1623-1672.
27. Fehr, E. and K. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114, 817-868.
28. Fraley, C. and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97, 611-631.

29. Fraley, C. and A. E. Raftery (2006). MCLUST version 3 for R: normal mixture modeling and model-based clustering. Technical report no. 504, Department of Statistics, University of Washington.
30. Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19, 24-42.
31. Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96, 1043-1067.
32. Huck, S., Jehiel, P. and T. Rutter (2007). Learning Spillover and Analogy-based Expectations: A Multi-Game Experiment. Working paper.
33. Jehiel, P. (2005). Analogy-based Expectation Equilibrium, *Journal of Economic Theory*, 123, 81-104
34. Jehiel, P. and F. Koessler. (2008). Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior*, 62, 533-557.
35. Jin, G. Z. and P. Leslie (2003). The effect of information on product quality: Evidence from restaurant hygiene grade cards. *Quarterly Journal of Economics*, 118, 409-451.
36. Johnson, E. and C. Camerer (2004). Thinking backward and forward in games. In I. Brocas and J.D. Carrillo (Eds.), *The Psychology of Economic Decisions, Vol.II: Reasons and Choices*, Oxford University Press.
37. Johnson, E., C. Camerer, S. Sen and T. Rymon (2002). Detecting failures of backward induction: monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104, 16-47.
38. Kagel, J. and D. Levin (2002). *Common Value Auctions and The Winner's Curse*. Princeton University Press, Princeton.
39. Knoepfle, D., Wang, J., and C. Camerer (2009). Studying learning in games using eyetracking. *Journal of the European Economic Association*, 7, 388-398.
40. Lohse, G. (1997). Consumer eye movement patterns on Yellow Page advertising. *Journal of Advertising*, 26, 61-73.
41. Mathios, A. (2000). The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market. *Journal of Law and Economics*, 43, 651-678.

42. McKelvey, R. and T. Palfrey (1995). Quantal response equilibria in normal form games. *Games and Economic Behavior*, 10, 6-38.
43. Milgrom, P. and N. Stokey (1982). Information, trade and common knowledge. *Journal of Economic Theory*, 26, 17-27.
44. Odean, T. (1999). Do Investors Trade Too Much? *American Economic Review*, 89, 1279-1298.
45. Ortoleva, P. (2009). The price of flexibility: towards a theory of thinking aversion. Working Paper.
46. Payne, J., J. Bettman, and E. Johnson (1993). *The Adaptive Decision Maker*. Cambridge University Press, Cambridge, UK.
47. Reutskaja, E., J. Pulst-Korenberg, R. Nagel, C. Camerer and A. Rangel (in press). Economic decision making under conditions of extreme time pressure: An eyetracking study. *American Economic Review*.
48. Rogers, B. , Palfrey, T. and C. Camerer (2010). Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory*, 144, 1440-1467.
49. Ross, S. (1989). Discussion: Intertemporal Asset Pricing. In Bhattacharya, Constantinides, eds., *Theory of Valuation*, Rowman & Littlefield Publishers, Inc., 85-96.
50. Rubinstein, A. (2006). Instinctive and cognitive reasoning: A study of response times. Working Paper.
51. Sonsino, D., I. Erev, and S. Gilat (2002). On rationality, learning and zero-sum betting – an experimental study of the no-betting conjecture. Working Paper.
52. Sovik, Y. (2004). Strength of dominance and depths of reasoning. Working Paper.
53. Stahl, D. O. and P. W. Wilson (1995). On players' models of other players: theory and experimental evidence. *Games and Economic Behavior*, 10, 218-254.
54. Tingley, D. H. and S. W. Wang (2010). Beliefs and strategies in a crisis bargaining model: An experimental study. Working paper.
55. Wang, J., Spezio, M., and C. Camerer (2010). Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth-telling and deception in games. *American Economic Review*, 100, 9841007.
56. Weiszäcker, G. (2003). Ignoring the rationality of others: Evidence from experimental normal-form games. *Games and Economic Behavior*, 44, 145-171.