Imperfect choice or imperfect attention? Understanding strategic thinking in private information games *

Isabelle Brocas
University of Southern California and CEPR

Juan D. Carrillo
University of Southern California and CEPR

Stephanie W. Wang
University of Pittsburgh

Colin F. Camerer
California Institute of Technology

Published in The Review of Economic Studies; Jul 2014, Vol. 81 Issue 3, p 944

Abstract

To understand the thinking process in private information games, we use "Mousetracking" to record which payoffs subjects attend to. The games have three information states and vary in strategic complexity. Subjects consistently deviate from Nash equilibrium choices and often fail to look at payoffs which they need to in order to compute an equilibrium response. Choices and lookup are similar when stakes are higher. When cluster analysis is used to group subjects according to lookup patterns and choices, three clusters appear to correspond approximately to level-3, level-2 and level-1 thinking in level-k models, and a fourth cluster is consistent with inferential mistakes (as, for example, in QRE or Cursed Equilibrium theories). Deviations from Nash play are associated with failure to look at the necessary payoffs. 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 (e.g., Camerer, 2003). Private information games are especially interesting because strategically naïve agents do not fully understand the link between the information and choices of other agents. Many types of lab and field evidence 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 naïvete, then distortions due to hidden information can be larger than predicted by equilibrium analysis, and ideal policy responses may be different (e.g., Crawford, 2003).

In this paper, we consider 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 choose whether to bet or not bet. Unless both players bet, they earn a known sure payoff. If both bet, they earn the payoff corresponding to the realized state. These games capture the essence of two-sided adverse selection. We also get information about decision processes by hiding the payoffs in opaque boxes. These are only revealed by a "lookup", when the computer mouse is moved into the box and a mouse button is held down. As in earlier experiments, subjects get feedback about the state after each trial, so they can learn.

Lab and field evidence suggest there are strategic thinking limits in many different games with private information.[1] This evidence can be explained by two types of theories. (1) *Imperfect choice*; or (2) *Imperfect attention*.

By imperfect choice, we mean stochastic response to payoffs (rather than optimization) or a simplification of some structural feature of likely behavior. Quantal response equilibrium (QRE) assumes players' beliefs are statistically accurate but players respond noisily to expected payoffs (McKelvey and Palfrey, 1995). Two other theories assume optimization, but also assume an imperfection in recognizing aspects of behavior. In cursed equilibrium (CE), players correctly forecast the distribution of actions chosen by other players, but underestimate the link between the private information and the strate-

[1] Examples include the winner's curse in common value auctions (Kagel and Levin, 2002), overbidding in private value auctions (Crawford and Iriberri, 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 over-communication in sender-receiver games (Cai and Wang, 2006; Wang et al., 2010). Field examples include overbidding 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., 2012), and excessive trading by individual stock investors (Odean, 1999).

gies of other players (Eyster and Rabin, 2005). In analogy-based expectation equilibrium (ABEE), players draw analogies between similar player types and use coarse statistics within analogy classes to form beliefs (Jehiel, 2005; Jehiel and Koessler, 2008; Huck et al., 2011). These different theories can all potentially explain non-equilibrium choices. However, they reflect fundamentally different cognitive processes.

By imperfect attention, we mean that some players do not attend to all elements of the game structure. CH and level-k models were first defined as specifications of how players beliefs about other players' choices can be incorrect. However, early in the development of these models it was clear, for many researchers, that particular incorrect beliefs were associated with imperfect attention (e.g., Camerer et al., 2004; Crawford, 2008). Then such models can be tested by a *combination* of choice and attention data, such as mouse-based or camera-recorded lookups.[2] A strong form of this joint choice-attention specification of level-k posits, for example, that if level 0 players randomize equally, then level 1 players will make choices that maximize expected payoff given those beliefs *and* will not need to look at the other players' payoffs (although they must still be aware of their opponents' set of possible actions).

The imperfect choice theories mentioned above could conceivably be associated with some type of imperfect attention as well, but no such association has been proposed. Therefore, we do not attempt to test QRE, CE and ABEE using the attentional data.[3]

Note that whether deviations from equilibrium are due to imperfect choice or imperfect attention (or both) is important for making predictions and giving advice. An imperfect attention account implies that variables like time pressure, information displays, and explicit instruction guiding attention could change behavior (as shown by Johnson et al., 2002).

Here is a summary of the basic results. As in previous related experiments, subjects play Nash equilibrium about half the time in simple situations (as defined later) and a quarter of the time in more complex situations. The combination of choices and lookup analysis suggest two clearly different types of non-equilibrium choices: some subjects look at the payoffs necessary to play the equilibrium choice but often fail to do so; in other cases, subjects do not look at all the necessary payoffs and also do not make equilibrium choices. At the same time, lookup is an imperfect but reasonably good predictor of Nash

[2]For example, Camerer, Ho and Chong wrote (2004, p. 891): "Since the Poisson-CH model makes a prediction about the kinds of algorithms that players use in thinking about games, cognitive data other than choices like... information lookups... can be used to test the model." The earliest Mouselab studies in economics begun more than ten years earlier (Camerer et al., 1993; and especially Johnson et al 2002) clearly treated levels as attention patterns, as did the more formal analysis in Costa-Gomes, Crawford and Broseta (2001).

[3]Other theories of potential interest include asymmetric logit equilibrium (Weiszacker 2003) and heterogeneous QRE (Rogers et al., 2009)

choice: the likelihood of equilibrium behavior is two to five times greater among subjects who look at the necessary payoffs.

The analysis at the individual level suggests heterogeneity in both lookups and choices. Subjects can be (endogenously) classified into four clusters using both measures. In cluster 1, subjects usually look at the necessary payoffs and play Nash but this cluster is small. Clusters 2 and 3 look at the necessary payoffs and usually play Nash in the simpler situations, but look and choose less strategically in the more complex situations. In cluster 4– the most common– subjects spend less total time making choices, look at necessary payoffs less often, and rarely play Nash. Differences in lookup patterns can be used to predict choice with some reliability.

Heterogeneous clustering suggests an approximate fit with the level-k model. Indeed, payoffs in our game are such that a level 1 player should never look at the necessary payoffs and never play the Nash equilibrium strategy. A level 3 should always look at the necessary payoffs and always play Nash. Finally, a level 2 should behave like a level 3 in the simple situations and like a level 1 in the complex ones. The choice and lookup data suggests that clusters 1 and 4 correspond approximately to levels 3 and 1 respectively and cluster 3 is a reasonable candidate for level 2. More generally, apart from a few rational players (12% in cluster 1) there appear to be a majority of players showing *imperfect attention* (67% in clusters 3 and 4) and a few others who exhibit *imperfect choice* (21% in cluster 2).

Actual earnings are interesting. Playing Nash is an empirical best response in the simple situations but not in the complex ones. The first cluster does not earn the most money because they act as if they overestimate the rationality of their opponents in complex situations (which misses the opportunity to earn more by exploiting mistakes of others). The middle clusters earn the most. The fourth cluster, who "underlook" compared to theory and rarely play Nash, earns the least. The highest earnings of the middle clusters could be due to either 'optimal inattention' (it does not pay to be sophisticated when rivals are not) or 'luck'. Our analysis tentatively favors the second explanation, although more research in this direction is needed.

Overall, the paper sheds light on the question: Why do people seem to underestimate adverse selection? Since theories of imperfect choice and imperfect attention could both explain the choice data, measuring information use directly is an efficient way to contribute to resolving the empirical puzzle. Also, equilibrium play in these games sometimes requires players to look up the numerical value of payoffs which they know they will never receive. Paying this kind of attention-to-impossible-payoffs is quite counterintuitive and therefore highly diagnostic of sophisticated strategic thinking.

Before proceeding to the formal analysis, we review the related literature. Three previous experimental studies have reported behavior in private information betting games: Sonsino et al. (2002), Sovik (2009) and Rogers et al. (2010). All these papers find substantial rates of non-equilibrium betting, and surprisingly little learning over many trials. Similar results are also obtained in simple experimental assets markets in which no trade is predicted to occur (Carrillo and Palfrey, 2011; Angrisani et al., 2011). Compared to the previous betting game experiments, our games have two novel properties (besides 'mouse-tracking'): Nash theory predicts a mixture of betting and no-betting across the games (so there is a rich statistical variety of behavior) and we consider 3-state games that are simpler than the 4-state or 8-state games in the existing literature.

Several studies have used attention measures to understand cognitive processes in economic decisions.[4] 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 choices in alternating-offer bargaining (Johnson et al., 2002). Information acquisition measures have then been used to study forward induction (Johnson and Camerer, 2004), matrix games (Costa-Gomes et al., 2001; Devetag, De Guida, Polonio 2013), two-person "beauty contest" games (Costa-Gomes and Crawford, 2006), learning in normal-form games (Knoepfle et al., 2009) and cheap-talk games (Wang et al., 2010).[5] Crawford (2008) summarizes these findings and argues for the value of lookups.[6]

2 Theory and Design

2.1 Theory: Equilibrium

We consider the following game. Nature draws a state, A, B or C, with equal probability. Player 1 learns whether the state is C or not C (that is, A or B) and player 2 learns whether the state is A or not A (that is, B or C). Players choose to bet on the state (action Y for 'yes') or to secure the sure payoff S (action N for 'not bet'). If either of them chooses to not bet, they each earn the number in the box under the sure payoff S. If both choose to bet, the payoffs for players 1 and 2 depend on the state A, B or C and are shown in the top and bottom rows of the matrix, as described in Figures 1

[4]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), search under time pressure (Reutskaja et al., 2011), and look for within-attribute vs. within-choice comparisons (Arieli et al., 2011).

[5]Wang 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.

[6]Other variables could be measured along with lookups, such as response times (Rubinstein, 2006) 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 cognitive processes. Attentional data is complementary to these other techniques.

or 2 for example. To imagine practical applications, think of imperfect knowledge of the state that results from coarse categorization or imperfect perception. So, if the states are Bad, Medium and Good, some players 1 may have expertise in detecting Good states but cannot distinguish Bad from Medium, whereas some players 2 have expertise in detecting Bad states but cannot distinguish between Good and Medium. Importantly, while player 1 cannot distinguish Bad from Medium, he knows that player 2 can (and vice versa). Formally, there are two information sets for each player: one is a singleton and the other contains two states.

Each information set of a game with the structure described above falls into one of three *situations*, characterized by the type of knowledge required for equilibrium choice. We call them F, D1 and D2. Figure 1 illustrates these three situations using the payoffs of one of the games in the experiment.

	A	B	C	S
1	25	5	20	$\underline{\mathbf{10}}$
2	$\underline{\mathbf{0}}$	30	5	

F: 2 in $\{A\}$

	A	B	C	S
1	$\underline{\mathbf{25}}$	$\underline{\mathbf{5}}$	20	$\underline{\mathbf{10}}$
2	$\underline{0}$	30	5	

D1: 1 in $\{A, B\}$

	A	B	C	S
1	$\underline{25}$	$\underline{5}$	20	$\underline{\mathbf{10}}$
2	$\underline{0}$	$\underline{\mathbf{30}}$	$\underline{\mathbf{5}}$	

D2: 2 in $\{B, C\}$

Figure 1 Different situation complexity for a given game. **Possible payoffs** are in bold. Payoffs in the MIN set are underlined.

The left table illustrates the Full information, F situation. In F, Player 2 has a singleton information set $\{A\}$. To find the equilibrium action, she only needs to compare the two sure payoffs, $[2A]$ and $[S]$. In this example, she knows that if she bets she will obtain 0 (if the other player bets), which is lower than the sure payoff $S = 10$. We call the set of information lookups that is necessary to choose the equilibrium action the "Minimum Information Necessary" (MIN). In this F situation, the MIN set is $[2A, S]$.

The table in the center illustrates the 1-step Dominance, D1 situation. Player 1 is in the information set $\{A, B\}$. She must look at $[2A]$ to determine if Player 2 will bet in $\{A\}$ and also look at her payoffs in $[1A]$ and $[1B]$ (as well as $[S]$, of course). This process is strategic because it requires looking at another player's payoff, realizing that the opponent has different information sets (which is explained in the instructions), and making an inference assuming the simplest level of rationality (dominance) of the opponent. In this example, once Player 1 realizes that Player 2 will not bet in $\{A\}$, she has no incentive to bet in $\{A, B\}$ as the payoff in $[1B]$ is lower than $[S]$.

The right table illustrates the 2-step Dominance, D2 situation. Player 2 is in the 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 $[S]$).[7] Notice also that looking at $[2A]$ means looking at a payoff she knows with certainty will not occur (an impossible counterfactual). This is quite counterintuitive and is therefore a clear hallmark of strategic reasoning.

These three situations can be unambiguously categorized from simplest to most complex, based on the degree of strategic thinking required. F is a trivial situation: the information required to play correctly is minimal and no strategic thinking is involved. D1 is a simple situation: it involves a more subtle reasoning and paying attention to the rival's payoff. D2 is a complex situation, as it require a further step in reasoning: the subject must anticipate how a rival will behave in D1 and best respond to it. It is then possible to perform comparative statics on behavior as a function of the complexity of the situation.

Finally, it is crucial to realize that MIN sets are defined from the perspective of an outside observer who is aware of the payoffs in all boxes. Our subjects, who need to lookup at the boxes to reveal the payoffs, have no way of knowing *a priori* whether they face a D1 or a D2 situation. It is therefore unlikely that they will open *only* the boxes in the MIN set. For that reason, in our analysis we will only distinguish between the subjects who open all the MIN boxes (and possibly some other that are not in MIN) from those who open a strict subset of the MIN boxes (and, again, possibly some other that are not in MIN). The former are coded 'MIN' and the latter are coded 'notMIN'. Subjects using different information processing algorithms will look at payoffs in different orders. However, they will have to open all the MIN boxes if they are to infer the equilibrium strategy.

2.2 Experimental games

Figure 2 shows the five payoff variants used in the experiment. (Game 1 is the example that was used above for illustration.) Among the 10 two-state information sets, 6 are in D1 (3 predict betting and 3 predict no betting) and 4 are in D2 (2 predict betting and 2 predict no betting). Games are designed to minimize the number of two-state information sets in which naïvely comparing the average of the two payoffs to $[S]$ gives the Nash choice (the only ones are $\{B, C\}$ in game 3 and $\{A, B\}$ in game 5). Also, because Nash theory predicts betting in some games and not in others, mistakes can occur in both directions

[7] In fact, it is not necessary to look at $[1A]$ in D1 or D2. Player 1 anticipates that Player 2 will not 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 $[2B]$. 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 a player who looks at one payoff in an information set almost invariably looks also at the other payoff in that set.

(overbetting as well as underbetting).[8]

Game 1

	A	B	C	S
1	25	5	20	10
2	0	30	5	

Game 2

	A	B	C	S
1	0	20	5	15
2	0	20	5	

Game 3

	A	B	C	S
1	5	20	20	15
2	10	25	10	

Game 4

	A	B	C	S
1	15	0	5	10
2	15	5	25	

Game 5

	A	B	C	S
1	10	25	25	20
2	25	40	15	

Figure 2 Payoff-variations of the betting game.

Theoretical predictions from Nash equilibrium are shown in Figure 3 for the two-state information sets, with 'Y' denoting the choice to 'bet' (Yes) and 'N' the choice to 'not bet' (N). We include the steps-of-dominance situation in parentheses (D1 or D2). The asterisk * appears 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 (as a non-strategic player would do for example). Predictions in the singleton 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. Our subjects, however, participate in all the games so within-subject choice patterns suggest that risk attitudes are unlikely to be the primary factor behind equilibrium deviations. Furthermore, Rogers et al. (2010) controlled for risk attitudes by using gambles as outside options and still found substantial non-equilibrium betting.[9]

Game 1

	A	B	C
1	N(D1)		
2		N(D2)	

Game 2

	A	B	C
1	Y(D1)		
2		Y(D1)	

Game 3

	A	B	C
1	Y(D1)		
2		Y(D2)*	

Game 4

	A	B	C
1	Y(D2)		
2		N(D1)	

Game 5

	A	B	C
1	N(D1)*		
2		N(D2)	

[8] Earlier studies have been criticized for predicting "inaction" in all trials (never bet, never trade). With boundary predictions of this sort, any deviation from equilibrium will look like a systematic bias. It is also possible that subjects would presume that experimenters would not create a game in which they are always supposed to take the same (in)action. In our setting, they could just as easily underbet as overbet.

[9] We focus on the unique Bayesian Nash Equilibrium (BNE) that survives iterated elimination of weakly dominated strategies. There are other BNE (e.g., both players choosing always N) which we neglect.

Figure 3 Nash equilibrium and situation complexity in each game. Note: Y and N denote equilibrium bet (Yes) and no bet (N); * denotes situations in which naïve averaging and Nash coincide.

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.[10] Mousetracking is a simple 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.[11]

There were six Mousetracking sessions run 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.[12] In each session, subjects played betting games with the five sets of payoffs described in Figure 2. Furthermore, each game was 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 randomly re-matched with another subject and randomly assigned to be either Player 1 or Player 2 in each trial. After each trial they learned the true state, the other player's action and their payoff. 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.

Five of the six sessions constitute what we will call hereafter the "Baseline" treatment, in which a total of 58 subjects participated. Subjects earned $20 on average plus a $5 show-up fee. The sixth session was a "High stakes" treatment in which we multiplied the

[10]In earlier versions of Mouselab, subjects do not have to hold down a button to keep information visible. 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.

[11]Cheap scaling-up is an advantage for studying multi-person games and markets compared to single-subject eye-tracking using Tobii or Eyelink camera-based systems, which cost about $35,000 each.

[12]Documentation and instructions for downloading the software can be found at `http://multistage.ssel.caltech.edu`.

earnings by 5 through a change in the conversion rate. Thus the payoffs in the boxes did not differ from the baseline sessions and any difference we observe in lookup and choice patterns should be attributed to the incentive effect. 20 subjects participated in that session and earned \$100 on average plus the same \$5 show-up fee.

We ran 2 additional sessions, the “Open boxes” treatment, in CASSEL at UCLA. The experimental design was the same as in the baseline experiment except that the Mousetracking technique was not implemented, and the subjects were facing a standard screen with visible payoffs. A total of 40 subjects participated in those sessions and earned \$20 on average plus a \$5 show-up fee. These sessions allow us to assess whether Mousetracking acts as a constraint on information processing.

No subject participated in more than one session. Table 1 summarizes the 8 sessions.

Session #	Treatment	Location	Subjects
1	Baseline	Caltech	12
2	Baseline	Caltech	12
3	Baseline	UCLA	10
4	Baseline	UCLA	10
5	Baseline	UCLA	14
6	High stakes	UCLA	20
7	Open boxes	UCLA	20
8	Open boxes	UCLA	20

Table 1: Summary of experimental sessions.

3 Aggregate analysis

3.1 Equilibrium play in the Baseline and Open boxes treatments

We first report the frequency of Nash play in the two baseline populations, Caltech (sessions 1 and 2) and UCLA (sessions 3, 4 and 5), and in the open boxes treatments (sessions 7 and 8 both run at UCLA).

Equilibrium play in F is 99% at Caltech and 96% at UCLA in baseline treatments, and 95% at UCLA in the open boxes treatment. Proportions of equilibrium play in the two-state information sets (D1 and D2) are reported in Figure 4. Because equilibrium play in the singleton information sets (F) is always close to 1, we omit it to aid visual clarity.

Caltech

	Game 1	Game 2	Game 3	Game 4	Game 5
1	.51	**.65**	**.76**	**.29**	.72
2	.35	**.73**	**.79**	.51	.13

UCLA

	Game 1	Game 2	Game 3	Game 4	Game 5
1	.58	**.39**	**.60**	**.32**	.59
2	.33	**.46**	**.88**	.39	.03

Open

	Game 1	Game 2	Game 3	Game 4	Game 5
1	.59	**.58**	**.72**	**.32**	.72
2	.35	**.70**	**.74**	.51	.06

Figure 4 Empirical frequency of equilibrium play in D1 and D2 by information set and game. **Bold frequencies** are cases where Nash predicts betting.

The proportion of equilibrium behavior is generally quite low (random behavior would predict .50). Equilibrium behavior differs across complexity of the situation (D1 v. D2) and, to a lesser extent, across populations (with Caltech students playing closer to Nash than UCLA students), but there are no systematic differences between games in which subjects are or are not supposed to bet.[13] The overall betting probabilities are also quite similar between our study and Rogers et al. (2010) and Sonsino et al. (2002).

The frequency of equilibrium choice with open boxes is similar to the Caltech baseline treatment for some information sets and to the UCLA baseline treatment for the others. We ran a series of Mann-Whitney tests (both pooling all baseline sessions and treating Caltech and UCLA baseline sessions separately) to compare the distribution of play in both treatments and we found that the difference in behavior was marginally statistically significant only for the UCLA baseline population in D1 (at the 5% level). This implies that subjects are not failing to play Nash due to imperfect memory or search costs. Our results, which may surprise some readers, are in fact in accordance with previous studies showing that behavior is very similar in sessions with open and closed boxes (Costa-Gomes et al., 2001; Johnson et al., 2002). If there is a friction in the subject's ability to process,

[13] This counters the suggestion that excessive betting in previous betting game experiments have been driven by an experimenter's demand effect or by the excitement of an uncertain payoff.

interpret or evaluate information, such problem is not significantly exacerbated in the closed boxes variant.[14] It also implies that the inferences about the information used by other players in the game should not be different with open and closed boxes. Having established there is little difference between choices in these two treatments, we will now concentrate on the baseline population.

3.2 Lookups and play in the F situation

Recall that F corresponds to a trivial individual decision-making problem in a singleton information set. As expected, subjects look at MIN and perform well. More precisely, subjects look at MIN 97% (Caltech) and 95% (UCLA) of the time and then play the equilibrium action 99% (Caltech) and 96% (UCLA) of the time. This suggests that both populations understand the fundamentals of the game: they compare the payoff of betting with the sure payoff and choose 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 on the 2 MIN boxes and the rest in some of the 5 other boxes. They also look at all 7 payoff boxes of the game only 18% (Caltech) and 22% (UCLA) of the time. This suggests that subjects look at the payoffs *strategically and succinctly.*

Behavior in this situation sheds some light on social preference theories. 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 information sets these subjects would deviate from Nash. Empirically, however, such sacrificial non-Nash behavior is extremely rare in the simplest F situation, when some social preferences might predict Nash deviations most clearly.[15,16]

There is little else to learn from the F situation. Unless otherwise stated, in the rest

[14] As we will see later (section 5), behavior is also similar with high stakes. Thus, contrary to the predictions of a cost-benefit search theory with option value of unveiling information, lowering the cost of search (open boxes) or increasing the benefit (high stakes) has a negligible marginal impact on behavior.

[15] Consider Player 2 in game 5 state $\{A\}$. While the Nash choice is to bet, an inequity-averse player might sacrifice and not bet. Similarly, Player 1 in game 4 state $\{C\}$ could bet contrary to the Nash solution due to efficiency concerns. However, Non-Nash behavior in these two scenarios is extremely rare (about 3% of the time). It could be that a player behaves pro-socially in D1 and D2 but not in F. However, we have no reason to suspect so. If anything, one would think that social preferences would be more prevalent in the simplest situations where the results of doing so are most palpable.

[16] The use of Mousetracking is also helpful in ruling out the influence of social preferences in these situations. Socially-regarding players would need to look at payoffs of other players to decide whether it is worth deviating from Nash in order to express their social preference. But in F situations players generally look at the possible payoff of the other player only rarely and quickly.

of the analysis we will focus on the two-states information sets, where subjects face the more interesting D1 and D2 situations.

3.3 Order of lookups

In this section, we address the sequence of lookups in D1 and D2. Note that subjects look at [S] 98% of the time in both D1 and D2. We therefore concentrate on the algorithm they use to look at the payoffs in the 6-box matrix. In a given trial, a subject reveals one box at a time and may need to open the same box several times. The actual sequence of lookups is difficult to analyze, in part due to the large number of lookup combinations and transitions over the 6 boxes. However, we are interested primarily in the first-order properties of sequencing. We consider a simple measure that allows us to extract that information by measuring lookups in information sets as opposed to lookups in individual boxes. That is, we are not interested in whether a player 1 in information set $\{A, B\}$ reveals first box [1A] or box [1B], but rather if and when he reveals any of these boxes. Moreover, we restrict to the first time there is a lookup in one of the boxes of an information set, and ignore the subsequent clicks on boxes in that information set.

Given there are 4 information sets, each of them can be opened in the first, second, third, fourth position or not at all. For expositional purposes, we take the perspective of Player 1 in $\{A, B\}$ and analyze the sequence chosen to reveal [1A,1B], her own payoffs in the states that can realize, [1C], her own payoff in the state that cannot realize, [2A], the payoff of her rival if she is fully informed, and [2B,2C], the payoffs of her rival in her two-state information set. Given the sequence is similar across populations, we pool Caltech and UCLA subjects together (a series of two-sample K-S tests showed that the distributions of sequences were not statistically different). Table 2 reports this information.

The order in which boxes are open is almost identical between D1 and D2 (we ran two-sample K-S tests to compare the distribution of sequences in D1 and D2 and found no statistical difference). The subject's own possible payoffs, [1A,1B] which are most salient, are revealed first and the sure payoff their rival can obtain [2A] is revealed second. This is quite reasonable: since players do not know a priori whether they face a D1 or a D2 situation, they look first at the 'most natural' boxes. The payoffs of the rival in the two-state information set [2B,2C] is revealed third, as it contains a payoff that can realize. The proportions of these first three lookups are very similar between D1 and D2. Finally, the payoff of the player in the state that cannot realize [1C] is either revealed in the last position or not revealed at all. The former is more common in D1 whereas the latter occurs more often in D2. Interestingly, subjects in D1 do not need to reveal [2B,2C] whereas subjects in D2 do need to reveal [1C] in order to find the equilibrium. Hence,

D1

Lookup #	$[1A, 1B]$	$[1C]$	$[2A]$	$[2B, 2C]$	None
1	<u>0.86</u>	0.02	0.09	0.03	0
2	0.08	0.15	<u>0.51</u>	0.17	0.09
3	0.03	0.08	0.17	<u>0.62</u>	0.09
4	0.02	<u>0.34</u>	0.14	0.10	<u>0.40</u>

D2

Lookup #	$[1A, 1B]$	$[1C]$	$[2A]$	$[2B, 2C]$	None
1	<u>0.86</u>	0.02	0.07	0.04	0
2	0.09	0.15	<u>0.53</u>	0.15	0.09
3	0.04	0.06	0.19	<u>0.63</u>	0.08
4	0.02	<u>0.45</u>	0.13	0.12	<u>0.30</u>

Table 2: Frequencies of the n^{th} lookup over the four information sets in D1 and D2 situations (highest frequencies underlined). Subject is Player 1 in $\{A, B\}$

there is an average tendency to overlook in D1 and underlook in D2.

Notice that there is no explicit cost for opening boxes. Still, some subjects stop quickly and do not check some boxes that turns out to be relevant. One possible interpretation is the existence of implicit costs of opening boxes. This interpretation, compatible with an optimal stopping rule, could rationalize limited search. However, the results obtained in section 3.1 suggest that subjects play similarly with open and closed boxes. Therefore, implicit costs should not be attributed to the Mousetracking design. If such costs exist, they are present in both treatments. The objective of the next sections is to correlate measures of lookups with equilibrium play.

3.4 Occurrence of lookups and equilibrium play

Having established the basic patterns of lookup search, we now study jointly attention and behavior. *Occurrence of lookups* is the simplest measure of attention, and the most conservative one (the only assumption is that a payoff which is never seen cannot be used in decision making). It is a binary variable that takes value 1 if the box under scrutiny has been opened at least once during the game for any amount of time.[17]

In order to distinguish between luck and strategic reasoning, we separate the * cases where Nash equilibrium coincides with the naïve averaging strategy from the cases where equilibrium and averaging differ. Table 3 provides basic statistics of the frequencies of behavior, Nash / notNash, and lookups in the necessary payoff boxes, MIN / notMIN (recall that looking exactly at MIN and looking at MIN plus some other boxes are both recorded as "MIN").

When averaging and equilibrium choices coincide (D1* and D2*), subjects play Nash significantly more often than when they do not but it is poorly or even negatively related to whether they look at MIN (rows 5 and 6). It suggests that these situations mix sophisticated (or cognitive) and naïve (or lucky) subjects.

When averaging and equilibrium choices do not coincide (D1 and D2), 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 natural, since the MIN set is smaller in the former than in the latter case, and does not involve looking at a payoff the subject knows for sure will not realize. Interestingly and in sharp contrast to D1* and D2*, the frequency of equilibrium choice is 2 to 5 times higher when subjects look at MIN than when they do not. These differences suggest that lookups can be predictive of equilibrium choice. 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).

We then compare the patterns of lookup and play between Caltech and UCLA. 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 situations. 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.

[17]Given our software configuration, the minimal duration of a click is around 50 milliseconds. 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).

% 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 3: Occurrence of lookups and equilibrium play (st. errors clustered at subject level).

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). We find a consistent but modest increase in the frequency of MIN-Nash and a modest decrease in the total lookup duration (see Appendix section 7.2.1 for details). In general and consistent with the previous literature on this game, some learning occurs but it is rather limited.

3.5 Duration (attention) of lookups

This section associates choices with *duration of lookups*. Contrary to section 3.3, we are now interested in collecting precise information for each single box. Figure 5 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.[18] To study the relationship between lookup and choice, we build on section 3.4 and classify observations in three types: MIN-Nash, MIN-notNash, and notMIN. Note that there is no information contained in Nash behavior if the subject does not look at MIN, so all notMIN observations will henceforth be pooled together into one category.

Because duration is quite similar across populations, we pool Caltech and UCLA subjects together. As before, we express results as if the subject is always Player 1 in information set $\{A, B\}$.[19]

D1 MIN-Nash [399]

	A	B	C	S
1	**$\underline{.17}$**	**$\underline{.19}$**	.06	.14
2	$\underline{.20}$	.14	.09	

Total duration: 6.4s
\# clicks: 15.7

MIN-notNash [259]

	A	B	C	S
1	**$\underline{.19}$**	**$\underline{.18}$**	.05	.14
2	$\underline{.16}$	.18	.10	

Total duration: 5.1s
\# clicks: 14.4

notMIN [130]

	A	B	C	S
1	**$\underline{.29}$**	**$\underline{.27}$**	.09	.21
2	$\underline{.03}$	.06	.06	

Total duration: 2.2s
\# clicks: 5.7

D2 MIN-Nash [82]

	A	B	C	S
1	**$\underline{.13}$**	**$\underline{.12}$**	$\underline{.13}$	.09
2	.16	$\underline{.20}$	$\underline{.17}$	

Total duration: 11.4s
\# clicks: 23.6

MIN-notNash [178]

	A	B	C	S
1	**$\underline{.16}$**	**$\underline{.15}$**	$\underline{.12}$	.12
2	.15	$\underline{.16}$	$\underline{.14}$	

Total duration: 7.9s
\# clicks: 20.7

notMIN [199]

	A	B	C	S
1	**$\underline{.22}$**	**$\underline{.20}$**	$\underline{.01}$	.21
2	.17	$\underline{.12}$	$\underline{.09}$	

Total duration: 3.7s
\# clicks: 9.3

Figure 5 % of lookup duration in D1 and D2 [# of observations in brackets]. Underlined boxes are in the MIN set. Subject is Player 1 in $\{A, B\}$

Let us focus first on the observations where subjects look at MIN. The relative time spent in the different boxes is remarkably similar when we compare MIN-Nash and MIN-notNash within each situation. The only difference is that subjects who play correctly

[18] The 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.

[19] So, 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 D1 and looking at the box of her rival in the full information set.

spend marginally more time thinking throughout the game. Lookups, however, are vastly different between situations. Differences are expected since MIN is defined differently in D1 and D2. 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 and sequentially at the information they *think* will be most relevant for decision-making. In particular, and as noted before, it makes sense that subjects spend a substantial amount of time at $[2A]$ in D2 (a box not in MIN) since they cannot know a priori whether they face a D1 or a D2 situation. As for the differences, 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. Total duration increases significantly from D1 to D2, reflecting the fact that subjects realize the increased difficulty of the situation. Also, subjects spend similar amount of time looking at their own and at their opponent's relevant payoffs. Taken together, the results suggest 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 a cognitive capacity to solve the equilibrium or an expectation about others' behavior than with an inability to understand that the game has strategic elements.[20]

Subjects in notMIN spend very little time thinking. In D1 they barely look at the payoffs of the other player or their payoff in $\{C\}$. From the lookup pattern, it seems that they 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 similar to that of MIN types in D1. This suggests that some of these subjects realize the strategic component of the decision-making process. However, they fail to follow all the necessary dominance steps to find the equilibrium.

Transitions between boxes can also be informative although it turns out that they are very closely linked to occurrences (see Appendix section 7.1 for details). The main feature is that there are substantial transitions from player 1's own payoff boxes to the crucial $[2A]$ in D1 for observations exhibiting Nash play. This is consistent with the remarkable fact that MIN-Nash players in D1 look *more often* at $[2A]$ than at their own $[1A, 1B]$ boxes.

3.6 Regression analysis: Predicting choice from specific lookup data

In this section we use lookup data to predict choices. Table 4 presents OLS regressions at the subject level using average lookup durations (in seconds) for those payoffs in MIN which are most likely to be predictive of Nash behavior. Looking at $[2A]$ is associated with an increase in the frequency of Nash play in D1. Looking at $[2B]$ and $[1C]$ is associated

[20]Players do look at more than just MIN in a number of trials but this type of "overlooking" has been observed in other studies on attention (Knoepfle et al., 2009; Wang et al., 2010). Since looking is costless, such overlooking could also be due to curiosity or limited memory.

with an increase in the frequency of Nash play in D2. 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 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	-.0018	.0022	-.0054***	.0017
Duration in [2A]	.0243**	.0096	—	—
Duration in [1C]	—	—	.0201***	.0056
Duration in [2B]	—	—	.0284***	.0082
# of transitions	1.183	.7559	-.5353	.5160
Constant	28.79***	8.50	23.36***	6.12
# Observations	58		58	
R^2	0.21		0.34	

Table 4: OLS regression on Percentage of Nash play
(* = significant at 10% level, ** = 5% level, *** = 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 5. The probability of playing Nash is affected by the time spent in [2A] for D1 and by the time spent in [2B] for D2. In D1 there is a significant positive effect of experience (match number). These trial-specific results are consistent with the OLS regression, although the R^2 values are significantly smaller. One difference is that looking at [1C] is not significant. There are no effects of individual difference measures.

3.7 Summary of aggregate analysis

The results obtained so far can be summarized as follows. (i) Behavior is similar with open and closed boxes, so implicit costs should not be attributed to the Mousetracking design. (ii) Deviations from Nash equilibrium abound and reflect severe limits on strategic thinking. (iii) Occurrence of lookup is an imperfect but reasonably good predictor of Nash choice. (iv) Order of lookups is similar between D1 and D2 but not the total number of open boxes, which suggests that subjects try to solve the game as they go and stop when they (rightly or wrongly) believe they have enough information to make a decision. (v) MIN-Nash and MIN-notNash players have similar lookup patterns, which indicates that sorting out the correct information is important but not sufficient for equilibrium play.

	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 [$2A$]	2.7×10^{-4}*	1.4×10^{-4}	—	—
Duration in [$1C$]	—	—	1.1×10^{-4}	0.7×10^{-4}
Duration in [$2B$]	—	—	2.8×10^{-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 5: Probit regression on probability of Nash play (st. errors clustered at subject level) (* = significant at 10% level, ** = 5% level, *** = 1% level).

(vi) Regression analysis also shows lookup duration in some MIN boxes to be predictive of Nash choice to some extent.

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 (see e.g. Costa-Gomes et al. (2001)). 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. One advantage is that it creates statistical evidence on how well the extreme single-cluster and subject-specific approaches capture what is going on. Another advantage is that it is model free. It does not impose any model of heterogeneity, but rather describes the heterogeneity found in the data as it is. As such, clustering 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.

Note that there are many combinations of lookups and behavior as a function of the situation. Nash theory predicts subjects should use one specific combination: look at MIN and play Nash in both D1 and D2. There are at least three alternative theories to explain the data: Quantal response equilibrium (QRE), Cursed equilibrium (CE), and Level-k theories.

QRE is an equilibrium model with noisy best responses which makes the same prediction as Nash theory in terms of lookups: it is necessary to look at the MIN set to play at equilibrium. In a χ-cursed equilibrium, all subjects believe that with probability χ the opponent's decision is independent of their private information and with probability $(1-\chi)$ other subjects are also χ-cursed. The model predicts only two combinations of lookups. One group of players ($\chi < 1$) should look as predicted by Nash theory. The other (fully cursed subjects or $\chi = 1$) should completely ignore information-action links. They only need to look at their own payoffs, and act like level-1 players. (So strong support for level-1 is also evidence for fully cursed behavior.)

Both QRE and CE seem to place many restrictions on the combinations of lookups and behavior in our game which do not match the heterogeneity in our data. The level-k model offers clearer predictions about both lookups and choices, so it might be the most natural theory to test with our data.[21] We describe the level-k model in detail in the next subsection.

4.1 Level-k

This section describes one version of how 'level-k subjects' look and choose in this game. We use the level-k specification rather than the cognitive hierarchy one because it allows us to make crisp parameter-free predictions about lookup and choices for each level.[22]

The first step is to define the behavior of a level-0 subject. We assume that a level-0 randomizes between betting and not betting in all information sets, ignoring their information.[23]

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

Level-2 players best respond to level-1. Subjects must therefore open the boxes a level-1 would open to deduce the level-1 action, 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 boxes and also open $[1A, 1B]$ to figure out her best response. Given this type of reasoning,

[21] One could design other games where Nash theory, QRE and CE also have rich lookup predictions. Such games would then be more suitable than the current one to compare the different theories.

[22] Under cognitive hierarchy, in contrast, the predictions depend on the parameter τ of the distribution of types. For example, if τ is sufficiently low, then all levels above 0 could exhibit the same behavior.

[23] The most natural alternative is a dominance-satisfying level-0 who makes Nash choices in F. Since level-1 players always satisfy (strict) dominance, level-2 players in our specification act like level-1 players in this alternative specification. Therefore, if our specification is wrong relative to this alternative, it simply "misnumbers" levels (inferring level 2 when it should be level 1).

a level-2 subject will inevitably miss $[1C]$. This means that in D1 situations she will look at MIN, but in D2 situations she will not. In our games, payoffs are designed in such a way that level-2 agents always play Nash in D1 and 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.

Figure 6 summarizes the lookups and choices of each level (X denotes a box the subject opens). Subjects who look at more boxes than a level-k but do not look at some of the boxes predicted by a level-k+1 are also categorized as level-k lookup.

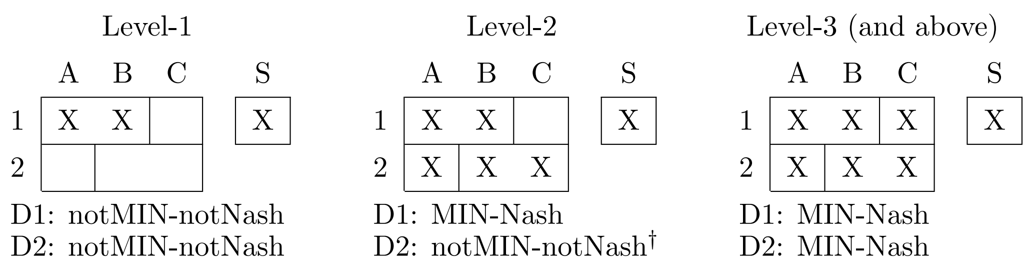

Figure 6 Lookup behavior and choice by level-k, k = $1, 2, 3$. Subject is Player 1 in $\{A, B\}$ (†with the exception of game 5).

4.2 Clusters based on lookup and choices

To find the clusters, we use the six aggregate statistics described in section 3 at the subject level: three combinations of lookup and play ("MIN-Nash", "MIN-notNash", and "notMIN") in D1 and the same three combinations in D2. We then compute the percentage of trials in which each subject's combination of lookup and choice is of either type in each situations. Each subject is thus measured by six percentages (or four variables since the three percentages for each situation must add up to one).

We group the 58 subjects of our baseline experiment (see Table 1) in clusters based on these six percentages. 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.

We implement model-based clustering analysis with the Mclust package in R (Fraley and Raftery, 2006). A maximum of nine 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 models 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 endogenously yields *four clusters* maximizes the BIC at -1117. Table 6 shows the frequencies of Caltech and UCLA subjects in each cluster, listed from high to low according to frequency of Nash choice and MIN lookup. 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 6: 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 6: 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 7 displays the empirical frequency of "MIN-Nash", "MIN-notNash", and "not-MIN" in D1 and D2 for subjects in each cluster. It also gives the predicted proportions for the various level-k types as described in section 4.1 and summarized in Figure 6.

Table 7 suggests that the lookup and choice behavior of subjects in clusters 1, 3 and 4 roughly correspond to the predictions of levels 3, 2 and 1 in the level-k model. Subjects in cluster 2 do not correspond to any specific level.

To be more precise, cluster 1 looks at MIN and plays Nash reasonably often in both D1 and D2 situations (like a level-3 player) whereas cluster 4 plays Nash rarely and often

D1	MIN-Nash	MIN-notNash	notMIN	# obs.
Cluster 1	.95 (.028)	.05 (.028)	.00 (n/a)	94
Cluster 2	.73 (.069)	.23 (.069)	.04 (.014)	159
Cluster 3	.57 (.048)	.38 (.043)	.05 (.014)	248
Cluster 4	.18 (.039)	.43 (.059)	.39 (.061)	287
Level 3+	1	0	0	
Level 2	1	0	0	
Level 1	0	0	1	

D2	MIN-Nash	MIN-notNash	notMIN	# obs.
Cluster 1	.58 (.071)	.27 (.031)	.15 (.070)	60
Cluster 2	.21 (.051)	.71 (.056)	.08 (.026)	96
Cluster 3	.15 (.034)	.44 (.032)	.41 (.058)	139
Cluster 4	.04 (.012)	.20 (.058)	.76 (.058)	164
Level 3+	1	0	0	
Level 2	0	0	1	
Level 1	0	0	1	

Table 7: Proportion of lookups and choices in D1 and D2: empirical by cluster and theoretical by level-k type (st. errors clustered at subject level).

does not look at MIN in D1 and D2 (like a level-1 player).[24] Clusters 2 and 3 switch from playing Nash in D1 to not playing Nash in D2 but they exhibit different lookup patterns. Cluster 2 subjects look at MIN payoffs in both D1 and D2 whereas cluster 3 subjects look at MIN in D1 but not in D2, where they consistently miss [$1C$]. The lookups and choices of cluster 3 then correspond roughly to level-2 (MIN-Nash in D1 and NotMIN in D2).[25] Cluster 2 on the other hand, exhibit the lookups of a level-3 and the choices of a level-2. Their behavior suggests that these individuals (21% of the population) always pay attention to the relevant information but make mistakes in complex situations (D2).

[24]Notice that level-1 predicts 'notMIN' but also 'notNash' 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.

[25]In Appendix section 7.5 we analyze strategic choice in situations where Nash play coincides with naïve averaging of payoffs in the information set. 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.

It reinforces the idea that correct lookups are necessary but not sufficient for equilibrium choice.

There are interesting differences across clusters regarding the effect of experience on lookup and choices. (See Appendix section 7.2.2 for details). We find no evidence of learning by subjects in clusters 2 and 4. Subjects in cluster 1 show significant learning in D2 situations (MIN-Nash rates increase from .40 to .71) whereas subjects in cluster 3 learn in D1 situations (MIN-Nash rates increase from .44 to .69), suggesting that these subjects conform more to level 3 and level 2 over time. There are also socio-demographic differences across clusters but the results are often weak in significance. (See Appendix section 7.3 for details). Male subjects with experience in game theory, 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. These suggestive results are encouraging but are not strongly established in our data.

4.3 Lookup statistics within clusters

Lookup frequencies and durations within clusters offer some clues about choice patterns. Some diagnostic summary statistics are presented in Table 8. It reports average total duration and average numbers of transitions. It also reports ratios of the time spent looking at the crucial MIN boxes containing other player's payoffs and their impossible own payoff [$1C$], compared to the average time spent looking at one's own possible payoffs.

Subjects in cluster 1 increase their looking the most– almost doubling it– from D1 to D2. The fact that they look *longer* at the MIN payoffs of the other player than they look at their own possible payoffs is a clear informational marker of strategic thinking. Subjects in cluster 2 spend a lot of total time looking, and spend about a third more time on D2 than on D1. They also make a lot of transitions. Subjects in cluster 3 show a similar pattern of looking to cluster 2 except they look at their own impossible payoff only half as often as they look at their own possible payoffs. Finally, looking patterns of cluster 4 subjects are entirely different. Duration is short and similar in D1 than in D2. Also, the ratios of looking time in the MIN payoffs of the other player and the impossible payoff box in D2 situations compared to their own possible payoffs are low.

The order of lookups is also interesting. Recall that the aggregate analysis revealed a typical order: [1A,1B], then [2A], then [2B,2C], and finally [1C] (or nothing). Table 9 reports the percentage of times the n^{th} lookup lies in the typical n^{th} information set. High numbers reveal that subjects are consistent with this typical order. We can see that most of the numbers up to the third lookup are above 50% suggesting that, on average, subjects in all four clusters open the boxes in the typical order. However, clusters differ in the proportion of times subjects adhere to this typical pattern and, more importantly,

D1

Cluster	$\text{avg}[1A, 1B]$	$[2A]$	Ratio $\frac{[2A]}{\text{avg}[1A,1B]}$	Total Duration	Number of Transitions
1	1061	1357	1.28	6166	11
2	1155	1113	.96	6603	15
3	1149	1218	1.06	6065	13
4	770	435	.57	3574	7

D2

Cluster	$\text{avg}[1A, 1B]$	[2B]	Ratio $\frac{[2B]}{\text{avg}[1A,1B]}$	[1C]	Ratio $\frac{[1C]}{\text{avg}[1A,1B]}$	Total Duration	Number of Transitions
1	1420	2063	1.45	1692	1.19	11501	18
2	1227	1364	1.11	1083	.88	8390	18
3	1239	1282	1.03	524	.42	7616	15
4	843	469	.56	158	.19	3952	8

Table 8: Average lookup times in MIN boxes (in milliseconds) and between-boxes ratios by cluster in D1 and D2.

in the proportion of times they complete the full search.

Subjects in cluster 1 complete the search in D2, but they stop earlier in D1. They act as if they make a running choice as they open boxes and reveal information on a need-to-know basis. Subjects in cluster 2 are more heterogeneous and disorganized. They tend to complete the full search more often, but the order of the search is more erratic. Subjects in cluster 3 exhibit similar sequencing patterns in D1 and D2 and tend to not complete the full search. Lastly, and consistent with the previous results, subjects in cluster 4 stop very early, completing only a small part of the typical sequence.

In Appendix section 7.4 we perform a comparative study of occurrence and duration of lookups in each box by cluster. The analysis reinforces the main results of this section. Also, the behavior of subjects in cluster 1 is particularly striking. 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 payoffs $[1A, 1B]$. We also perform Probit regressions separately for all clusters and find similar results as in the aggregate analysis.

D1 Cluster	1st lookup in $[1A, 1B]$	2nd lookup in $[2A]$	3rd lookup in $[2B, 2C]$	4th lookup in $[1C]$	Stop after 3rd lookup
1	0.90	0.68	0.61	0.37	0.36
2	0.69	0.52	0.65	0.53	0.18
3	0.90	0.53	0.67	0.31	0.38
4	0.92	0.44	0.55	0.19	0.62

D2 Cluster	1st lookup in $[1A, 1B]$	2nd lookup in $[2A]$	3rd lookup in $[2B, 2C]$	4th lookup in $[1C]$	Stop after 3rd lookup
1	0.87	0.68	0.73	0.65	0.12
2	0.67	0.46	0.55	0.61	0.06
3	0.93	0.59	0.73	0.42	0.33
4	0.91	0.46	0.53	0.12	0.57

Table 9: Frequency of n^{th} lookup in the typical n^{th} information set revealed.

4.4 Are players optimally inattentive? Earning and learning

One important question we have postponed until now is whether, given the behavior of other subjects, it is actually optimal to play Nash strategies. This is an empirical question, which the data can answer. If it is not always optimal, 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 an earnings mistake.

The first empirical observation is that playing Nash is optimal in D1 (averaging across the entire sample) but is *not* optimal in D2. The reason is that the average payoffs in the two-state information sets are often well above the sure payoff when not betting is the Nash choice (and viceversa). So if other players choose non-equilibrium betting often enough, then a player who bets obtains an average payoff above $[S]$. Because Nash play is optimal in D1 situations but not in D2, it is possible that cluster 1 players, who play Nash most often, may not earn the most money. Perhaps surprisingly, this is true.

Table 10 shows average expected normalized earnings for each cluster.[26] It also shows

[26] The normalization scales earnings such that a value 0 matches random choice and 1 matches empirical best response. To calculate expected earnings, we take the expectations over ex-ante probabilities of each state of the information set, to smooth out luck (resp. bad luck) from ending up in a high (resp. low) payoff 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 role with equal probability. A similar approach was used by Lucking-Reiley

for comparison the earnings that a level-k player would obtain. Notice that 'level 0' corresponds to the normalized earnings of a random player and 'level 2' to the normalized earnings of a player who best responds to the empirical choice of others (Nash in F and D1 and notNash in D2). Also, 'level 3+' corresponds to the normalized earnings of a subject who always plays the Nash equilibrium strategy.

	Cluster				Level			
	1	2	3	4	0	1	2	3+
F	.95	.95	.94	.95	.00	1.0	1.0	1.0
D1	.79	.46	.17	-.34	.00	-1.0	1.0	1.0
D2	-.025	.46	.60	.59	.00	1.0	1.0	-1.0

Table 10: Expected normalized earnings by cluster and level-k type (Random = level 0, Best Response = level 2, Nash = level 3+)

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

In D1, clusters are clearly ranked from 1 to 4 in earnings. Since the empirical best response is Nash, the clusters that choose Nash most often also earn the most. Note that cluster 4 does much worse than random, since they play Nash less than half the time.

In D2, clusters 2, 3 and 4 earn similar amounts. From earnings alone, they seem to have an understanding based on the game structure, and perhaps history, that it does not pay to play Nash in these complex situations. Cluster 1 subjects– who frequently play Nash– earn about the same as random players, and less than any other cluster including the clueless cluster 4 (although not as little as a Nash player)! They represent an interesting case of people who analyze the game the most 'rationally' but do not translate that into earnings. We can only speculate as to why cluster 1 players did not shift to notNash. Perhaps they were hoping to teach other players to play the equilibrium actions or perhaps they did not have enough observations to revise their erroneous beliefs about the other players' types.

It is also instructive to compare earnings by cluster with earnings by level-k types. Clusters 1 and 4 are similar to levels 3+ and 1 in the earnings space, although cluster 1 does not do as bad in D2 as level 3+ and cluster 4 not as bad in D1 as level 1. This corroborates our findings in section 4.2. Cluster 2 is the closest in earnings to level 2. However, as we noted earlier, they look at payoffs more like a level 3+ type.

and Mullin (2006). This reduces variability introduced by the matching mechanism. The 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 game is then normalized by subtracting random earnings and dividing by the difference between the empirical best response earnings and random earnings.

Averaging across the three situations, subjects in cluster 2 earn the most. There are two possible reasons. One 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 cluster 2 subjects look at MIN almost always but have a much lower rate of Nash choices in D2 than cluster 1 subjects. A different interpretation is that subjects in cluster 2 are not thinking that shrewdly, but are simply lucky because enough other players deviate from Nash choices making notNash an empirical best response in D2.[27] Several pieces of evidence point in that direction: (i) they earn significantly less than cluster 1 subjects in D1; (ii) they do not learn with experience; (iii) they look at the impossible $[1C]$ payoff less than at their own possible payoffs; and (iv) they still get lower payoffs in D2 than the clueless cluster 4 subjects. Although we favor this second alternative, there is not enough evidence in our data for a firm conclusion. Future research should design tasks to better disentangle between these two possibilities.

The data can also say something about the possibility that players are "optimally inattentive"– i.e., it does not pay to look at MIN to choose the best strategies since other players are not playing Nash.[28] This hypothesis is appealing but there is a lot of evidence against it.

First, these subjects are highly practiced using a mouse to retrieve information. The fact that they often look at notMIN boxes in F (40% of their search is on other boxes) reveals that the cost is low. The fact that choices are very similar with open boxes suggests that the cost may even be comparable to that of an eye fixation.

Second, learning appears to guide subjects to look at *more* of the MIN boxes over time, not less (see Appendix section 7.2 for details on learning effects). If subjects were learning to be inattentive in the face of non-Nash play, there would be a decline in MIN looking occurrences, particularly in D2 situations where Nash strategies are played less often. The fact that MIN looking increases or decreases very slightly over time in those situations suggests subjects are not learning to look less in the face of non-Nash play.

Third, the combination of looking times in Table 8 with earnings statistics in Table 10 shows there is generally a positive relation between looking and earning.

[27] In this respect, our two-action game is a handicap: when Nash and best response to empirical strategy differ, a player who deviates from Nash (whatever the reason) is categorized as playing best response.

[28] A different type of "rational inattention", namely the idea that individuals are exogenously constrained in their capacity to process information, has been the object of recent research in macroeconomics (see e.g. Sims (2003) and Woodford (2008)).

4.5 Summary of results

The cluster analysis reveals that lookup patterns over relevant information and conversion of this information into Nash choice are heterogeneous. Three clusters map roughly onto level-k thinking types. Cluster 4 corresponds to level-1 and cluster 1 corresponds to level-3. Cluster 3 shows a tendency in the direction of level-2 but not very sharply. Finally, the behavior of subjects in cluster 2, who look like Nash players but deviate from equilibrium choice in complex situations, suggests a possible role for stochastic choice and imperfect responses, as in QRE. There is some learning by subjects in cluster 1 (for complex situations) and in cluster 3 (for simple situations). We also find that longer lookup durations are correlated with greater payoffs both in D1 and D2.

5 High stakes

We now turn to the high stakes (×5) treatment. In F, subjects played Nash 95% of the time and looked at MIN 90% of the time. These frequencies are similar to those in the baseline treatment. The proportions of Nash equilibrium play in D1 and D2 situations are reported in Figure 7.

Game 1

	A	B	C
1	.59		
2		.26	

Game 2

	A	B	C
1	**.55**		
2		**.60**	

Game 3

	A	B	C
1	**.62**		
2		**.70**	

Game 4

	A	B	C
1	**.34**		
2		.51	

Game 5

	A	B	C
1	.54		
2		.08	

Figure 7 Empirical frequency of equilibrium play in high stakes games (×5).

The results are comparable to the frequencies obtained in the UCLA baseline treatment (see Figure 4). Subjects did not play Nash more often when stakes were high. The occurrence of look-up and behavior is reported in Table 11. The results are also strikingly similar to the baseline (see Table 3). Subjects looked at MIN in the same proportions as in the baseline treatment and transformed it into Nash choice at comparable rates. If anything, MIN lookup with high stakes is not as good predictor of Nash choice as MIN lookup in the baseline treatment (rows 5 and 6). Figure 8 reports durations in each box for each game. Total durations and durations per box are again similar to those observed in the baseline (see Figure 5). The only noticeable difference is an increase in the number of clicks: with high stakes, subjects open the same boxes, spend less time on each of them and reopen each box more times.

% of observations	High Stakes			
	D1	D2	D1*	D2*
MIN-Nash	.51 (.030)	.12 (.026)	.52 (.067)	.36 (.069)
MIN-notNash	.34 (.029)	.35 (.037)	.32 (.063)	.18 (.055)
notMIN-Nash	.07 (.016)	.10 (.024)	.02 (.018)	.34 (.068)
notMIN-notNash	.09 (.017)	.43 (.039)	.14 (.047)	.12 (.046)
Pr[Nash | MIN]	.60 (.033)	.26 (.051)	.62 (.074)	.67 (.092)
Pr[Nash | not MIN]	.44 (.077)	.20 (.043)	.11 (.11)	.74 (.094)
# observations	271	162	56	50

Table 11: Occurrence of lookups and equilibrium play with high stakes.

D1 MIN-Nash [137]

	A	B	C	S
1	**.18**	**.19**	.05	.13
2	.17	.17	.10	

Total duration: 6.7s
clicks: 21.5

MIN-notNash [91]

	A	B	C	S
1	**.22**	**.21**	.05	.13
2	.17	.14	.08	

Total duration: 6.0s
clicks: 17.1

notMIN [43]

	A	B	C	S
1	**.23**	**.23**	.08	.14
2	.09	.12	.11	

Total duration: 2.9s
clicks: 13.6

D2 MIN-Nash [20]

	A	B	C	S
1	**.17**	**.16**	.12	.10
2	.16	.16	.14	

Total duration: 10.6s
clicks: 31.5

MIN-notNash [56]

	A	B	C	S
1	**.17**	**.19**	.10	.11
2	.14	.17	.12	

Total duration: 7.0s
clicks: 24.7

notMIN [86]

	A	B	C	S
1	**.23**	**.24**	.03	.15
2	.14	.13	.08	

Total duration: 4.4s
clicks: 13.7

Figure 8 Duration of lookups with high stakes [# of observations in brackets].

To further assess the impact of stakes on behavior, we report in Table 12 probit regressions similar to the ones presented in Table 5. We pooled the observations from the baseline and high stakes treatments and added a dummy variable taking value 1 for observations in the high stakes treatment and 0 otherwise. The results we obtained were comparable to the baseline treatment and the high stakes dummy variable was not found to be significant.

Taking these results together suggests that increasing the stakes does not affect significantly the lookup patterns of subjects. They do not spend more time analyzing the

	D1		D2	
	coefficient	st. error	coefficient	st. error
Total duration	1.6×10^{-5}	1.5×10^{-5}	2.6×10^{-6}	1.9×10^{-5}
Duration in $[2A]$	2.3×10^{-4}**	1.0×10^{-4}	—	—
Duration in $[1C]$	—	—	7.9×10^{-5}	6.5×10^{-5}
Duration in $[2B]$	—	—	2.3×10^{-4}***	7.6×10^{-5}
Experience (trial #)	.016***	.0041	.0047	.0056
Highstakes	.054	.17	.025	.17
Constant	-.49***	.15	-1.17***	.18
# Observations	1059		621	
R^2	0.05		0.07	

Table 12: Probit regression on probability of Nash play (st. errors clustered at subject level) (* = significant at 10% level, ** = at 5% level, *** = at 1% level)

game and they do not shift attention to other boxes. Stakes do not affect their choices either. There are two possible interpretations. It could be that the threshold multiplier that triggers a different attention and behavior is way above 5. Or, it could be that the ability of subjects to solve the game is limited, making the increased incentives ineffective. We believe the second interpretation is a more plausible explanation of the data.

6 Conclusion

The objective of this paper was to improve our understanding of human strategic thinking. To do this, we used a Mousetracking system to record information search in two-person games with private information about payoff-relevant states. We found significant heterogeneity among agents in their choice *and* lookup behavior. The results suggest that some subjects are limited in their attention information, while others are limited in their ability to process information. Game complexity also plays a key, intuitive role.

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

There is some support for the imperfect attention approach in our analysis. Subjects are endogenously clustered in four groups according to their choice and lookup patterns.[29] Clusters 1 and 4 seem to correspond closely to level 3 and level 1 strategic thinkers, and cluster 3 is a reasonable 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 Nash 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 making an imperfect QRE choice. Since QRE has never been specified as a theory of joint information search and choice, we cannot conclude that this cluster represents those looking patterns, but such an approach is promising.

However, it is also true that subjects almost always optimize in the simplest F games. It is probably difficult for a single-parameter 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 easy choice (F games), and a substantially lower parameter to fit the many non-Nash choices in more complex situations. 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 imperfect choice.

Multiple-type variants of cursed equilibrium could fit the combination of lookups and choices better than QRE. One interpretation of CE which is consistent with cluster-level choice data is that there are different types of players, with cursed parameters $\chi = 0$ (Nash players as in cluster 1) and $\chi = 1$ (fully naïve players as in cluster 4). This CE specification leaves out clusters 2 and 3, which are half the subject pool. Interestingly, cluster 2 might fit a generalized version of CE in which $\chi = 0$ types think there are $\chi > 0$ types (Eyster and Rabin, 2005, Appendix A). In that case, players could look at all payoffs (since some perceived types have $\chi = 0$ which requires full analysis) but then decide to play non-Nash if the perceived χ is high enough.[30]

As with QRE, one can imagine modifications of cursed equilibrium in which degrees of cursedness are manifested by intermediate looking and choice patterns. While such a

[29]This result contrasts with Carrillo and Palfrey (2009) who found heterogeneity in behavior but no clustering of subjects around a few strategies in their "compromise game." A combination of cursedness and smooth imperfections (cursed-QRE) fitted their data best. Although a comparison is difficult because the compromise game is not conducive to mousetracking, the differences in results pose a challenge for future experimental and theoretical research in private information games.

[30]By contrast, in this generalized version, $\chi = 1$ types cannot think there are $\chi < 1$ types, otherwise they would also have to look at the other players' payoffs.

possibility is not part of the standard specification, it is a challenging direction for future research.

More generally, economists often talk casually about "contemplation costs" (Ergin and Sarver, 2010), "control costs" (van Damme and Weibull, 2002), "thinking aversion" (Ortoleva, in press) 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.

References

1. Angrisani, M, Guarino, A, Huck, S., and N. Larson (2011). No-trade in the laboratory. B.E. Journal of Theoretical Economics, Advances, 11(1): Advances.

2. Arieli, A., Ben-Ami, Y., and A. Rubinstein (2011). Tracking decision-makers under uncertainty. American Economic Journal: Microeconomics, 3(4), 68-76.

3. 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.

4. Brown, A.L., Camerer, C. and D. Lovallo (2012). To review or not to review? Limited strategic thinking at the box office. American Economic Journal: Microeconomics, 4(2), 1-28.

5. Cai, H., Wang, J. (2006) Overcommunication in strategic information transmission games. Games and Economic Behavior, 56, 7-36.

6. Camerer, C. (2003). Behavioral game theory: Experiments on strategic interaction. Princeton University Press, Princeton.

7. Camerer, C. (2008). The neural basis of strategic thinking. In Glimcher, Camerer, Fehr, Poldrack (Eds.), Neuroeconomics, Elsevier.

8. Camerer, C. and T. Ho (1999). Experience-weighted attraction learning in normal form games. Econometrica, 67, 827-874.

9. Camerer, C., Ho, T. and J. Chong (2004). A cognitive hierarchy model of games. Quarterly Journal of Economics, 119, 861-898.

10. 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.

11. Capen, E. C., R. V. Clap, and W. M. Campbell (1971). Competitive Bidding in High-Risk Situations. Journal of Petroleum Technology, 23, 641-653.

12. Carrillo, J.D. and T. Palfrey (2009). The Compromise Game: Two-Sided Adverse Selection in the Laboratory. American Economic Journal: Microeconomics, 1(1), 151-181.

13. Carrillo, J.D. and T. Palfrey (2011). No Trade. Games and Economic Behavior, 71(1), 66-87.

14. Charness, G. and D. Levin (2009). The Origin of the Winner's Curse: A Laboratory Study. American Economic Journal: Microeconomics, 1, 207-236.

15. Charness, G. and M. Rabin (2002). Understanding social preferences with simple tests. Quarterly Journal of Economics, 117, 817-869.

16. 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.

17. Costa-Gomes, M. and V. Crawford (2006). Cognition and behavior in two-person guessing games: an experimental study. American Economic Review, 96, 1737-1768.

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

19. Costa-Gomes, M. and G. Weizsäcker (2008). Stated beliefs and play in normal form games. Review of Economics Studies, 75, 729-762.

20. Crawford, V. (2003). Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions. American Economic Review, 93, 133-149.

21. 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.

22. Crawford, V. and N. Iriberri (2007). Level-k auctions: can a non-equilibrium model of strategic thinking explain the winner's curse and overbidding in private-value acutions? Econometrica, 75, 1721-1770.

23. van Damme, E. and J. Weibull (2002). Evolution in games with endogenous mistake probabilities. Journal of Economic Theory, 106, 296-315.

24. Devetag, G., S. De Guida, and L. Polonio (2013). An eye-tracking study of feature-based choice in one-shot games. ECORE working paper 2013/6.

25. Ergin, H. and T. Sarver (2010). A unique costly contemplation representation. Econometrica, 78 (4), 12851339.

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 (2011). Feedback Spillover and Analogy-based Expectations: A Multi-Game Experiment. Games and Economic Behavior, 71(2), 351-365.

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. 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. Lucking-Reiley, D. and C. H. Mullin. (2006). Recombinant Estimation for Normal-Form Games, with Applications to Auctions and Bargaining. Games and Economic Behavior, 54 (1), 159-182.

42. 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.

43. McKelvey, R. and T. Palfrey (1995). Quantal response equilibria in normal form games. Games and Economic Behavior, 10, 6-38.

44. Odean, T. (1999). Do Investors Trade Too Much? American Economic Review, 89, 1279-1298.

45. Ortoleva, P. (in press). The price of flexibility: towards a theory of thinking aversion. Journal of Economic Theory.

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 (2011). Economic decision making under conditions of extreme time pressure: An eyetracking study. American Economic Review, 101, 900-906.

48. Rogers, B. , Palfrey, T. and C. Camerer (2010). Heterogeneous quantal response equilibrium and cognitive hierarchies. Journal of Economic Theory, 144, 1440-1467.

49. Rubinstein, A. (2006). Instinctive and cognitive reasoning: A study of response times. Working Paper.

50. Sims, C.A. (2003). Implications of Rational Inattention. Journal of Monetary Economics, 50, 665-96.

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. (2009). Strength of dominance and depths of reasoning: an experimental study. Journal of Economic Behavior and Organization, 70(1), 196-205.

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). Belief Updating in Sequential Games of Two-Sided Incomplete Information: An Experimental Study of a Crisis Bargaining Model. Quarterly Journal of Political Science, 5(3), 243-255.

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, 984-1007.

56. Weizsäcker, G. (2003). Ignoring the rationality of others: Evidence from experimental normal-form games. Games and Economic Behavior, 44, 145-171.

57. Woodford, M. (2008). Inattentiveness as a source of randomized discrete adjustment", working paper.

7 Online 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 box to another can be diagnostic of computational algorithms. Figure 9 presents a graphical summary of the results, where the arrows in the matrix indicate transitions between betting payoff boxes and the outside arrows capture transitions to and from the sure payoff. We divide the analysis into D1 and D2 situations, pool both populations together, and separate the data into the usual three combinations of attention and choice (a comparison of the distributions of transitions at Caltech and UCLA reveals no significant differences). 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.

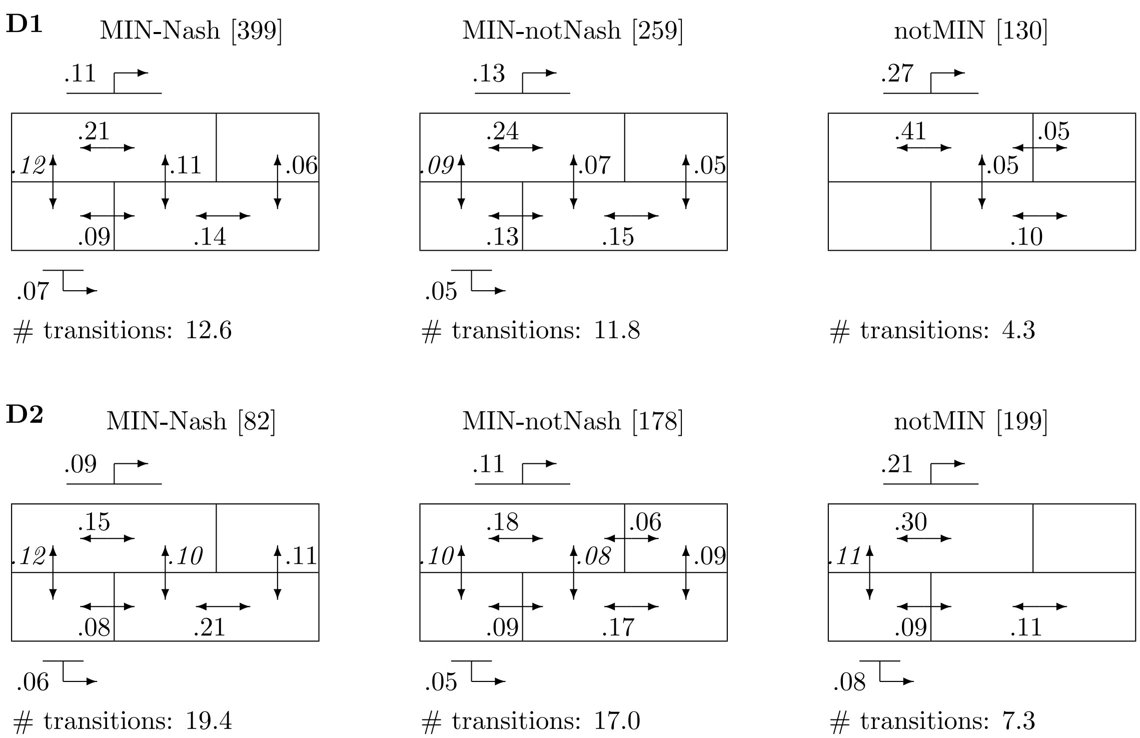

Figure 9 transitions of lookups in D1 and D2 [# observations in brackets]. *Italics* indicates between-player transitions needed for MIN. Subject is Player 1 in $\{A, B\}$.

Among subjects who look at MIN, the patterns are similar within a situation 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 5 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 13 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 13: Effect of experience on lookups and behavior.

In every subject pool and situation 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 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 and in D2 at Caltech only.[31]

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 but not in D2.[32] 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 situations.

7.2.2 Learning at the cluster level

Table 14 reports learning statistics by cluster. As before, we divide the data between early trials (1 to 20) and late trials (21 to 40). In simple D1 situations, only subjects in cluster 3 benefit from

[31]More 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.

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

experience and increase both their MIN lookup and Nash play rate substantially. Clusters 1 and 2 know how to play from the beginning and their performance does not vary over time. Subjects in cluster 4 are clueless and do not learn at all. In complex D2 situations, the frequency of Nash play rises over time in cluster 1. All other clusters do not seem to play better with experience. These observations are confirmed by two-sample Wilcoxon signed rank tests. When comparing the distributions of types in early and late trials, we find that the differences are significant for cluster 3 in D1 (p-value < 0.001) and for cluster 1 in D2 (p-value 0.057). One should notice however that the number of observations is rather small.

	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 14: 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 15 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.

Table 16 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

	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 15: Summary of Questionnaire answers.

their counterparts. A variety of two-sample Wilcoxon rank-sum tests show that 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.[33] However, differences are not significant in D2.

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 17. 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 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.[34]

7.4 Lookup patterns within clusters

7.4.1 Lookup durations within clusters

This section contains analysis similar to that in section 3.5 except that we consider only subjects in the same cluster.[35]

For illustrative purposes, we first present in Figure 10 a scatterplot that depicts the clusters. Each data point corresponds to one subject. While the clustering is done using four independent

[33] When 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.

[34] When 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. The 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.

[35] Statistics 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.

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 (.084)	.49 (.050)	.75 (.056)
# observations	443	215	368	290
D2	.29 (.048)	.37 (.069)	.26 (.057)	.37 (.053)
# observations	172	86	136	124

Table 16: Effect of Questionnaire answers.

statistics (the three types of lookup and play in each of the two situations, given that probabilities must add up to one), 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. The y-axis is the percentage of trials on which the subject looked at MIN and played Nash in D1.

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.[36]

Figure 11 shows average number of transitions, total time and percentage of time in each box 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 3.5 for details). For this cluster, the looking patterns are a portrait of rationality: they look a lot at the crucial $[2A]$ box in D1 (longer than they look at their own payoff boxes). They also look a lot at $[2B]$ and $[1C]$ in D2 compared to their looking rates in D1 (and, again, longer than they look at their potential payoffs $[1A, 1B, S]$).[37] When they do look correctly, they convert their information into

[36] Note 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 two other statistics that are not plotted.

[37] Because 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)

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 17: Frequencies based on Questionnaire answers.

Nash choices at high rates (although lower in D2).

D1 [94]	A	B	C	S
1	.18	.17	.09	.14
2	.22	.12	.08	

Total duration: 6.1s
transitions: 10

D2 [60]	A	B	C	S
1	.13	.11	.14	.11
2	.17	.18	.16	

Total duration: 11.3s
transitions: 18

Figure 11 Lookup behavior by subjects in cluster 1 [# observations in brackets], box entries are % of total lookup time in that box.

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).

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.

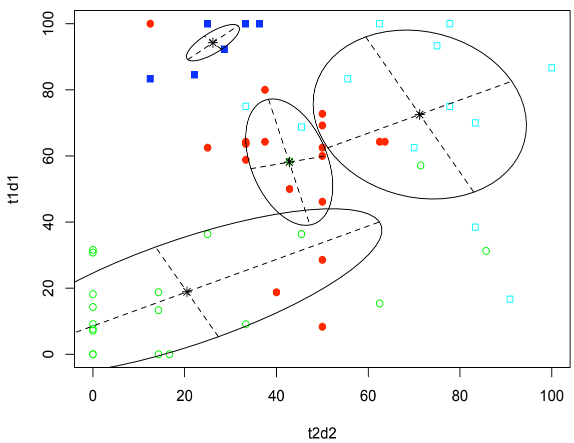

Figure 10 Four clusters of 58 subjects.

Figure 12 shows lookup statistics for cluster 2. Recall that their MIN percentages are very high in both situations (above 90%), but they only play Nash 76% and 24% of the time. Put it simply, they look at payoffs like cluster 1 subjects, but choose Nash at rates more like cluster 3 subjects (see Table 7 and the discussion below).

D1 [159]	A	B	C	S
1	.17	.18	.07	.14
2	.17	.15	.12	

Total duration: 6.8s
transitions: 15

D2 [96]	A	B	C	S
1	.15	.14	.13	.12
2	.16	.16	.14	

Total duration: 7.4s
transitions: 17

Figure 12 Lookup behavior by subjects in cluster 2 [# observations in brackets], box entries are % of total lookup time in that box.

The main and perhaps only significant difference between Figures 11 and 12 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 compared to 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 boxes 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 13 shows lookup statistics for cluster 3. Recall that their MIN rates are very high in D1, which is consistent with a long lookup duration in the crucial box [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, 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
1	.17	.21	.05	.12
2	.20	.16	.09	

Total duration: 6.0s
transitions: 13

D2 [139]	A	B	C	S
1	.17	.15	.07	.12
2	.17	.17	.15	

Total duration: 7.7s
transitions: 15

Figure 13 Lookup behavior by subjects in cluster 3 [# observations in brackets], box entries are % of total lookup time in that box.

Finally, Figure 14 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. 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). It also illustrates the importance of clustering. In section 3.5, 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
1	.22	.21	.04	.19
2	.12	.14	.07	

Total duration: 3.6s
transitions: 7

D2 [164]	A	B	C	S
1	.21	.21	.04	.19
2	.14	.12	.08	

Total duration: 3.9s
transitions: 8

Figure 14 Lookup behavior by subjects in cluster 4 [# observations in brackets], box entries are % of total lookup time in that box.

7.4.2 Information search by cluster in F situations

The data from F situations can be also used to shed light on the nature of information processing. Recall that in F the search problem is trivial (players need only compare the 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 is *not* used for clustering subjects. The analysis of F situations thus provides a nice template to study differences in search patterns across clusters.

Some lookup duration statistics are reported in Table 18. Subjects look at the MIN boxes $[1C]$ and $[S]$ about 350-700msec each, which is much higher than the 100-150msec they look at non-MIN ("Other") boxes (in fact, 14 out of the 15 median box durations are 0 in clusters 1, 3 and 4). The last column reports the fraction of total lookup that is fixated on the two MIN boxes. 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	589	555	104	1663	0.69
2	684	524	251	2461	0.49
3	649	435	156	1865	0.58
4	669	357	112	1522	0.65

Table 18: Average lookup times in MIN boxes by cluster in F situations

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

7.4.3 Regression analysis by cluster

Last, the probit regressions predicting choices on a trial-by-trial basis performed in section 3.6 can also be done separately for all the clusters. Table 19 reports these regressions. The results are generally consistent with the signs and magnitudes of the aggregate analysis but are much weaker in significance (probably due to modest sample sizes in each cluster). In D1, longer lookups at $[2A]$ predicts a higher likelihood of Nash play in three of the four clusters. In D2, longer lookups at $[2B]$ predicts a higher likelihood of Nash play only in cluster 3. Interestingly, the regressions reveal different learning pattern across clusters. Subjects in cluster 1 do learn in complex D2 situations. Subjects in cluster 3 learn only in simple D1 situations. All other subjects do not learn over the experiment.

7.4.4 Regression analysis with pooled clusters

Table 20 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 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 $[1C]$. It also

D1	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Total duration	-2.4 $\times 10^{-4}$***	-2.1 $\times 10^{-5}$	4.4 $\times 10^{-5}$*	5.3 $\times 10^{-6}$
Duration in [$2A$]	9.2 $\times 10^{-4}$***	3.0 $\times 10^{-4}$*	-1.9 $\times 10^{-5}$	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 $\times 10^{-6}$	-4.3 $\times 10^{-5}$	-4.3 $\times 10^{-5}$	9.6 $\times 10^{-5}$
Duration in [$1C$]	2.9 $\times 10^{-4}$	9.3 $\times 10^{-5}$	-9.5 $\times 10^{-5}$	-.0014**
Duration in [$2B$]	7.4 $\times 10^{-6}$	2.9 $\times 10^{-4}$	3.7 $\times 10^{-4}$***	2.3 $\times 10^{-4}$
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 19: Probit regression on Nash play by cluster (* = significant at 10% level, ** = significant at 5% level, *** = significant at 1% level).

suggests that lookup is a relevant but imperfect predictor of choice. By contrast, a combination of specific lookups is a reasonably good predictor. Last, the answers to the Questionnaire have again little effect.[38]

7.5 Predicting choice in other situations

Most of the analysis in the main text disregarded the situations in which Nash play coincides with the naïve strategy of comparing [S] with the average of payoffs in the information set (denoted D1* and D2* in section 3.4). In this section, we analyze the performance in those situations for subjects in each cluster. Results are reported in Table 21 (standard errors clustered at the individual level are reported in parenthesis).

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 confirms this hypothesis.[39]

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,

[38]The 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.

[39]Cluster 4 subjects also play Nash significantly more often in D1* when they do not look at MIN (data not reported for brevity).

D1	Clusters		
	1+2	1+3	1+3+4
Total Duration	$-.4.4 \times 10^{-5}$	1.2×10^{-5}	1.4×10^{-5}
Duration [$2A$]	4.0×10^{-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

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}	2.1×10^{-4}*	4.2×10^{-4}**
Duration [$2B$]	1.8×10^{-4}	2.8×10^{-4}**	2.5×10^{-4}***	2.6×10^{-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 20: Probit regression on Nash play by groups of clusters (* = significant at the 10% level, ** = significant at the 5% level, *** = significant at the 1% level).

	D1*				D2*			
Cluster	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 21: Proportion of types by cluster in D1* and D2* when Nash and naïve averaging coincide.

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*.

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. It is indeed puzzling that the most careful types play notNash more often when a simple heuristic leads to Nash play.

7.6 Clustering based only on lookups

The cluster analysis in the main text grouped subjects according to *both* information search and play. A different analysis uses only lookups to cluster subjects, 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 previously, 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.[40]

The clustering and regression analysis performed in section 4 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 clusters 1-2 and clusters 3-4. Statistics (3) and (4) could then help discriminate between cluster 3 and cluster 4 (we tried different clustering models and report only one in this section).

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).[41]

[40] In 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, home purchase, etc.

[41] The fourth group subject is cross-classified as cluster 3 with a profile closely resembling group 2. This

Table 22 presents for each group (classified according to lookups) the proportion of observations that can be categorized as MIN-Nash, MIN-notNash and notMIN. This table can be compared to Table 7 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 22: Proportion of types by group in D1 and D2.

The main reason that the groups derived only from lookup statistics do not match the lookup-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 situations, and cluster 2 in the earlier analysis who play Nash much less often, especially in D2. These two clusters are lumped together into group 1 when only lookups are used. This is a reminder that lookups are correlated with choice, but do not perfectly predict choice. It also suggests an area for future research: find other measures besides simple lookups which can pin down which subjects will make certain choices, given the information they have attended to.

Table 23 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. 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 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 analysis. However, the fact that the choice frequency differences are much clearer in the lookup-plus-choice clusters means lookups have limited predictive power.

subject is placed in a separate group because of a much larger total duration than those in group 2.

	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 23: Classification of subjects by cluster and group.