Information Gatekeepers: Theory and Experimental Evidence *

Isabelle Brocas
University of Southern California and CEPR

Juan D. Carrillo
University of Southern California and CEPR

Thomas R. Palfrey
California Institute of Technology

June 2009

Abstract

We consider a model where two agents with opposite interests can spend resources in acquiring public evidence about the unknown state of the world in order to influence the choice of a decision-maker. We characterize the sampling strategies of agents in the Nash equilibrium of the game. We show that, as the cost of news acquisition for an individual increases, that agent collects less evidence whereas his opponent collects more evidence. We then test the results in a controlled laboratory setting. The behavior of subjects is close to the theoretical predictions. Mistakes are relatively infrequent (15%). They occur in both directions, with more over-sampling (39%) than under-sampling (8%). The main difference with the theory is the smooth decline in sampling around the theoretical equilibrium. Comparative statics are also consistent with the theory, with subjects sampling more when their cost is low and when their rival's cost is high. Finally, there is little evidence of learning over the 40 matches of the experiment.

*We thank the audiences at the CIRANO Political Economy workshop in Montreal, the Hong Kong University of Science and Technology and the National University of Singapore for comments, Dustin Beckett for research assistance, and Chris Crabbe for developing the software. We also thank the LUSK Center, the Office of the Provost at the University of Southern California and the Microsoft Corporation for financial support. Addresses for correspondence: Isabelle Brocas and Juan D. Carrillo, Department of Economics, University of Southern California, Los Angeles, CA 90089, <brocas@usc.edu> and <juandc@usc.edu>; Thomas R. Palfrey, Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, <trp@hss.caltech.edu>.

1 Motivation

The economics of information literature has devoted considerable effort in understanding the strategic use of private information by agents in the economy. However, much less is known about the *strategic collection of information*. And yet, economic examples of this situation abound. For example, lobbies and special interest groups spend substantial resources in collecting and disseminating evidence that supports their views. The US legal system is based on a similar advocacy principle: prosecutor and defense attorney have opposite objectives, and they each search and provide evidence on a given case in an attempt to tilt the outcome towards their preferred alternative. In this paper, we build a theoretical model to understand the incentives of individuals to collect information in strategic contexts. The results are then tested in a controlled laboratory setting.

We consider a simple theoretical framework where two individuals with opposite interests can acquire costly evidence. Formally, there are two possible states of the world. Nature draws one state from a common prior distribution. Individuals can then acquire signals that are imperfectly correlated with the true state. Information affects the belief about the true state and is public, in the sense that all the news collected by one party are automatically shared with the other agents. This means in particular that parties have common beliefs at every point in time. When both individuals choose to stop the acquisition of information, a third agent (the decision-maker) makes a binary choice. The mapping between information and decision maker's choice is deterministic and known: it favors the interests of one agent if the belief about the relative likelihood of states is below a certain threshold and it favors the interests of the other if the belief is above that threshold.

Opposite incentives implies that agents never acquire costly information simultaneously. Indeed, when the current evidence implies that the decision-maker will favor the interests of one agent, that individual will not have incentives to provide information, as extra evidence can only hurt him. We then determine the optimal sampling strategies. The individual who would be harmed by the final decision if no further information were accumulated has to trade-off the cost of information and the likelihood that such information will change beliefs in his favor. As the belief becomes more and more adverse, the likelihood of reversing it decreases and the expected sampling cost necessary to achieve a belief reversal increases, so the net gain of accumulating evidence goes down. Overall, when the belief is mildly against the interests of one individual, that agent acquires information. He keeps sampling up to a point where either the belief is reversed, in which case the other agent starts the sampling process over, or else it has become so unfavorable that it is preferable to give up. Solving this problem analytically is challenging, since the value function of each agent depends on the sampling strategy of both agents. Indeed, the value of sampling for information in order to 'send the belief to the other camp' depends on how intensely the opponent will sample for information himself, and therefore how likely he is to 'bring the belief back'. In Proposition 1, we determine the best response strategies of each individual as a function of the common belief about the state and the cost of sampling for each subject. We provide an analytical characterization of the Nash equilibrium, and perform some comparative statics: when the cost of news acquisition for an individual increases, that agent has fewer incentives to collecting evidence, whereas his rival has more incentives to collect evidence.

We then report an experiment that analyzes behavior in this two-person information acquisition game. We study variations of the game where each agent may have a low or a high unit cost of sampling. The structure of the game is therefore identical in all treatments, but the equilibrium levels of sampling is not. Our first and main result, is that the empirical behavior in all treatments is close to the predictions of the theory both in the action space (Result 1) and in the payoff space (Result 2). This conclusion is remarkable given that the optimal stopping rule is fairly sophisticated, it involves strategic considerations about the opponent's choice, and it prescribes corner choices (either never sample or always sample). To be more precise, when the agent is unfavored by the existing evidence, his optimal action depends on whether the belief is mildly adverse (in which case he should sample) or strongly adverse (in which case he should stop), where the cutoff between 'mildly' and 'strongly' depends on the cost of sampling. We show that the individual takes the decision predicted by Nash theory with probability 0.85 (0.92 when he should sample for information and 0.61 when he should not). Furthermore, the best response to the empirical strategy of the opponent is to play the Nash equilibrium, which reinforces the idea that deviations from equilibrium play are small. Similar results are obtained when we analyze choices in the payoff space: given their empirical behavior, subjects lose less than 5% of the payoff they would obtain if they best responded to the strategy of their rival.

Second, we study in more detail the deviations observed in the data. The main difference with the theoretical predictions is the smooth rather than sharp decline in sampling around the theoretical equilibrium. We also show that mistakes occur in both directions. In general, there is more over-sampling than under-sampling. Also, under-sampling occurs relatively more often when the agent's cost is low and over-sampling occurs relatively more often when the agent's cost is high (Result 3). Because the decline in sampling is smoother than it should, it is also instructive to perform some comparative statics. The predictions of the theory are also supported by the data in that dimension. As the cost of sampling for an individual increases, that person is less likely to acquire information whereas his opponent is (weakly) more likely to do so. That conclusion holds in the empirical analysis at the aggregate level using Probit regressions on the probability of sampling, and at the state-by-state level using mean comparisons of sampling between cost pairs (Result 4). Finally, there is little evidence of learning by the agent unfavored by the existing evidence, possibly because the problem is difficult, the feedback is limited and, most importantly, the choices are close to equilibrium right from the outset. The agent favored by the existing evidence makes few mistakes at the beginning and learns to avoid them almost completely by the end of the session (Result 5).

The paper is related to two strands of the literature, one theoretical and one experimental. On the theory side, Brocas and Carrillo (2007) and Gul and Pesendorfer (2008) provide models of acquisition of public information. The former is technically similar to ours but considers only one agent and free information. It thus ignores the strategic component of optimal sampling. It shows how an individual can affect the choices of others by optimally refraining from acquiring public knowledge. The latter considers a closer setting to ours where two agents with conflicting preferences collect costly information. It shares with the present paper the strategic substitutability result of the sampling cutoffs of players. The model, however, is specified in quite different terms: decisions are made in continuous time and agents choose the instantaneous intensity of sampling. This formulation is elegant but unsuitable for an experiment. Finally, there is an older literature on games of persuasion (Matthews and Postlewaite (1985) and Milgrom and Roberts (1986) among others) that studies the ex-ante incentives of individuals to acquire verifiable information given the ex-post willingness to reveal it depending on its content.

On the experimental side, there is to our knowledge no laboratory test of *search for* *information* nor of *search in strategic contexts*. There is, however, an extensive literature on search for payoffs by a single individual (see e.g. Schotter and Braunstein (1981) and the surveys by Camerer (1995) and Cox and Oaxaca (2008)). The common finding is that subjects stop the search process either optimally or excessively soon. Insufficient experimentation can be rationalized by assuming agent's risk-aversion. As suggested by Camerer (1995, p. 673), a search for information could lead to different conclusions, although it is not clear which behavioral reason could account for it. More importantly, we are interested in the *strategic effect* that the sampling behavior of one individual has on the choice of his rival. Despite the game being substantially more complex than the individual decision-making counterpart, we still observe a behavior that is close to the theory. Also and contrary to the above mentioned literature, excessive experimentation is more frequent than insufficient experimentation. Finally, games of persuasion have been studied in the laboratory by Forsythe, Isaac and Palfrey (1989).

The paper is organized as follows. In section 2, we present the model and the main theoretical proposition. In section 3, we describe the experimental procedures. In section 4, we analyze the results, including aggregate behavior in the action space and the payoff space, deviations from equilibrium as a function of the costs of both agents, comparative statics (aggregate and state-by-state), and learning. In section 5, we provide some concluding remarks. The proof of the main proposition is relegated to the appendix.

2 The model

2.1 The game

Consider a game with 3 agents. Agent 0 is a *judge*, who must undertake an action that affects the payoff of every individual. Agents 1 and 2 are *advocates*, who can collect costly evidence about an event that has realized in order to affect the belief (hence, the action) of the judge. We assume that all the information collected by advocates becomes publicly available. Thus, at any point in time, all agents share the same belief about which event has realized. However, because agents have different preferences over actions, they will also have different incentives to stop or continue gathering evidence as a function of the current belief. Whether public information is a realistic assumption or not depends very much on the issue under consideration. We introduce it mainly because we want to *isolate* the incentives for information gathering. In that respect, adding private information would only pollute the analysis by mixing the incentives to collect information with the incentives to transmit information.

To formalize the information collection process, we consider a simple model. There are two possible events, $S \in \{B, R\}$ (for "blue" and "red"). One event is drawn by nature but not communicated to any agent. The judge must choose between two actions, $a \in \{b, r\}$. Both the payoff of the judge (agent-0) and the payoff of advocate i ($\in \{1, 2\}$) depend on the action and the event realized, and they are denoted by $v(a \,|\, S)$ and $u_i(a \,|\, S)$, respectively. The expected payoffs are:

$$v(a) \equiv \sum_S \Pr(S) v(a \,|\, S) \quad \text{and} \quad u_i(a) \equiv \sum_S \Pr(S) u_i(a \,|\, S)$$

To preserve symmetry, we assume that the common prior belief that the event is B is $\Pr(B) = 1/2$. At each stage, each agent i simultaneously decides whether to pay a cost c_i (> 0) in order to acquire a signal $s \in \{\beta, \rho\}$, which is imperfectly correlated with the event realized. Formally:

$$\Pr[\beta \,|\, B] = \Pr[\rho \,|\, R] = \theta \quad \text{and} \quad \Pr[\beta \,|\, R] = \Pr[\rho \,|\, B] = 1 - \theta$$

where $\theta \in (1/2, 1)$. Because the prior is common and all the information is public, all players have common posterior beliefs about the likelihood of each event. Also, in this simple framework, bayesian updating implies that the posterior belief depends exclusively on the difference between n_β, the number of β-signals, and n_ρ, the number of ρ-signals accumulated by agents. Formally:

$$\Pr(B \,|\, n_\beta, n_\rho) \;\equiv\; \mu(n) = \frac{1}{1 + \left(\frac{1-\theta}{\theta}\right)^n}$$

where $n \equiv n_\beta - n_\rho \in \mathbb{Z}$. Thus, for the purpose of the posterior held, two opposite signals cancel each other out. It is immediate to check that $\mu(n+1) > \mu(n)$ for all n, $\lim_{n \to -\infty} \mu(n) = 0$ and $\lim_{n \to +\infty} \mu(n) = 1$. We assume that from the judge's viewpoint there is one "correct" action for each event: action b if the event is B, and action r if the event is R. Formally, $v(b|B) > v(r|B)$ and $v(b|R) < v(r|R)$. As a result, there will always exist a belief $\mu^* \in (0, 1)$ such that $v(b) \gtreqless v(r)$ if and only if $\mu(n) \gtreqless \mu^*$.

2.2 Optimal stopping rule with competing advocacy

Suppose that two advocates with conflicting preferences can collect public evidence. For simplicity, suppose that one advocate wants the judge to take action b independently of the event realized, and the other advocate wants the judge to take action r also independently of the event realized.[1] From now on, we call them the blue advocate and the red advocate, respectively. Formally, we assume that the payoffs of the blue and red advocates are, respectively, $(u_B(b) = \pi_B, u_B(r) = 0)$ and $(u_R(b) = 0, u_R(r) = \pi_R)$ with $\pi_B > 0$ and $\pi_R > 0$.

Advocates can acquire as many signals $s \in \{\beta, \rho\}$ as they wish. The cost of each signal is c_B for the blue advocate and c_R for the red advocate. The timing is as follows. At each stage, advocates simultaneously decide whether to pay the cost of acquiring one signal or not. Any signal acquired is observed by all agents (judge, blue advocate, and red advocate). Agents update their beliefs and move to a new stage where advocates can again acquire public signals. When both advocates decide that they do not wish to collect any more information, the judge takes an action and the payoffs of all players are realized.

In this setting, advocates have opposite incentives and compete to provide information. Given the judge's utility described in section 2.1, there will exist a value k for the difference between the number of β-signals and the number of ρ-signals such that $\mu(k-1) < \mu^* < \mu(k)$. We normalize the payoffs in such a way that $k = 0$.[2] This means that if advocates stop collecting information when $n \geq 0$, then the judge takes action b. Conversely, if advocates stop when $n \leq -1$, then the judge takes action r.[3] It is then immediate that the blue advocate will never collect evidence if $n \geq 0$, as the current belief implies the optimal action from his viewpoint. For identical reasons, the red advocate will never collect evidence if $n \leq -1$ (from now on, we will say that the blue advocate is "ahead" if $n \geq 0$ and "behind" if $n \leq -1$). Define $\lambda \equiv (1-\theta)/\theta$ (< 1), $F_B \equiv c_B(1+\lambda)/(1-\lambda)$ and $F_R \equiv c_R(1+\lambda)/(1-\lambda)$. Although technically challenging, it is possible to characterize analytically the optimal sampling strategies under competing advocacy. This is done in the following proposition.

Proposition 1 *Under competing advocacy, the red advocate samples if and only if* $n \in \{0, \dots, h^* - 1\}$ *and the blue advocate samples if and only if* $n \in \{-l^* + 1, \dots, -1\}$. *The*

[1] This assumption is without loss of generality. What we need for the theory is a vector of preferences such that the judge has conflicting interests with one advocate for beliefs in one compact set and conflicting interests with the other advocate for beliefs in another compact set.

[2] One of the inequalities could be weak. We assume that both are strict to avoid tie-breaking rules.

[3] In other words, the blue advocate has an initial advantage. The advantage is lost after one ρ-signal.

equilibrium cutoffs are $h^*(l) = \arg\max_h \Pi^r_n(l,h)$ *and* $l^*(h) = \arg\max_l \Pi^b_n(l,h)$*, where:*

$$\Pi^r_n(l,h) = \frac{1}{1+\lambda^n}\left[\left(\pi_R(1+\lambda^l) - F_R(h+1)(1-\lambda^l)\right)\left[\frac{\lambda^n - \lambda^h}{1-\lambda^{h+l}}\right] - F_R(h-n)(1-\lambda^n)\right],$$

$$\Pi^b_n(l,h) = \frac{1}{1+\lambda^n}\left[\left(\pi_B(1+\lambda^h) - F_B(l-1)(1-\lambda^h)\right)\left[\frac{1 - \lambda^{n+l}}{1-\lambda^{h+l}}\right] + F_B(n+l)(1-\lambda^n)\right].$$

Advocates sample more if their cost is lower and their benefit is higher. Also, the stopping thresholds of advocates are strategic substitutes, so advocates sample more if the cost of the rival is higher and the benefit of the rival is smaller.[4]

Proof: see Appendix A1.

The idea is simple. Two advocates with conflicting goals will never accumulate evidence simultaneously. Indeed, for any given belief, one of the advocates will be ahead and therefore will not have incentives to collect information as it can only hurt his interests. Suppose now that $n \geq 0$. The red advocate (who is currently behind) can choose to collect evidence until he is ahead (that is, until he reaches $n = -1$), in which case either the other advocates samples or action r is undertaken yielding a payoff π_R. Alternatively, he can cut his losses, stop the sampling process, and accept action b that yields a payoff of 0. As the difference between the number of blue and red draws increases, the likelihood of reaching $n = -1$ decreases and the expected number of draws in order to get to -1 increases, making the sampling option less interesting. This results in an upper cutoff h^* where sampling by the red advocate is stopped. A symmetric reasoning when $n \leq -1$ implies a lower cutoff $-l^*$ where sampling by the blue advocate is stopped. Overall, when the event is very likely to be B the red advocate gives up sampling, and when the event is very likely to be R the blue advocate gives up sampling. For beliefs in between, the advocate currently behind acquires evidence while the other does not. The strategies are graphically illustrated in Figure 1.

[4] These comparative statics are determined by taking derivatives in the profit functions $\Pi^r_n(h,l)$ and $\Pi^b_n(h,l)$ (see Appendix A1). Obviously, there is a strong mathematical abuse in doing so, since h and l have to be integers. To avoid this technical issue in the experiment, we simply determine for each cost pair treatment the equilibrium cutoffs by creating a grid: for each integer l we find the integer h that maximizes $\Pi^r_n(l,h)$ and for each integer h we find the integer l that maximizes $\Pi^b_n(l,h)$ and use these values to find the Nash equilibrium. Naturally, the same comparative statics hold.

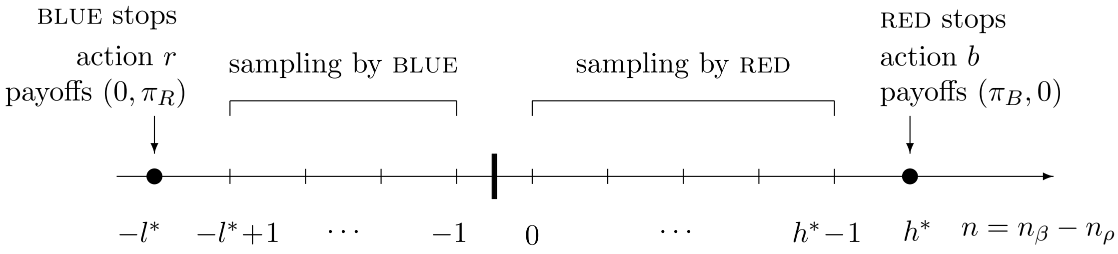

Figure 1. Sampling strategies

The comparative statics results are rather intuitive: an advocate has higher incentives to sample the lower his cost and the higher the benefit when his favorite action is taken by the judge. More interestingly, the stopping thresholds of advocates, h^* and l^*, are strategic substitutes. If the red advocate decides to sample more (h^* increases), the value for the blue advocate of reaching $n = 0$ is decreased, since his rival is more likely to find evidence that brings the belief back to $n = -1$. As a result, the blue advocate has less incentives to sample (l^* decreases). Combined with the previous result, it also means that if the cost of one advocate decreases or the benefit under his favorite action increases, then the other advocate will engage in less sampling.

3 Experimental design and procedures

We conducted 8 sessions of the two-advocate game with a total of 78 subjects. Subjects were registered students at the California Institute of Technology who were recruited by email solicitation. All sessions were conducted at The Social Science Experimental Laboratory (SSEL). All interaction between subjects was computerized, using an extension of the open source software package, Multistage Games.[5] No subject participated in more than one session. In each session, subjects made decisions over 40 paid matches. For each match, each subject was randomly paired with one other subject, with random rematching after each match.

The experimental game closely followed the setting described in section 2. More specifically, at the beginning of each match, each subject in a pair was randomly assigned a

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

role as either a "red" player or a "blue" player. The computer chose an urn for that pair, red or blue, with equal probability (the realized event). The urn was presented to the subjects. A red urn contained two red balls and one blue ball. A blue urn contained one red ball and two blue balls. Subjects knew the number of red and blue balls in each urn but did not observe the colors of the balls in the urn selected by the computer. That is, the event realized remained unknown to subjects.

Each subject had to decide simultaneously whether to draw one ball from the urn or not (the sampling strategy). Because there were twice as many red balls than blue balls in the red urn and twice as many blue balls than red balls in the blue urn, the correlation between signal and event (ball color and urn color) was $\theta = 2/3$. The cost of drawing a ball for the red and blue player, c_R and c_B respectively, was know to all subjects but varied on a match-by-match basis as detailed below. If one or both subjects drew a ball, then *both subjects* observed the color(s) of the ball(s) drawn. The ball was then replaced in the urn.[6] If at least one subject drew a ball, they both moved to another round of ball drawing. The process continued round after round until neither of them chose to draw a ball in a given round. At that point, the match ended, and the computer allocated a number of points to each subject which depended exclusively on the color of the balls drawn by both players.[7] More precisely, if the difference between the number of blue and the number of red balls drawn was 0 or greater, then the blue player earned 150 points and the red player earned 50 points. Form now on, we will say the blue payer "won" the match and the red player "lost" the match. If the difference was -1 or smaller, then the blue player lost the match and earned 50 points, whereas the red player won the match and earned 150 points. From these earnings, subjects had their ball drawing costs (number of balls they drew times cost per draw) subtracted. Subjects then moved to another match where they were randomly rematched, randomly reassigned a role and a new urn was randomly drawn.

There are a few comments on the experimental procedures. First, we wanted to mini-

[6]Even though the decision of drawing a ball within a round was taken simultaneously, the balls were drawn with replacement. That is, players always had 3 balls to draw from (this point was clearly spelled out in the instructions).

[7]As shown in Proposition 1, if the advocate unfavored by the evidence accumulated so far prefers not to draw a ball, then he has no incentives to start the sampling process afterwards. Thus, ending the match if no player draws a ball in a given round shortens the duration of the experiment without affecting the outcome.

mize the likelihood that a subject earned negative points in a given match once the costs were subtracted because this could result in loss aversion effects. For that reason, we allocated 50 points for losing a match.[8] Second, as in the theory section, roles were not symmetric. We gave an initial advantage to the blue player in order to implement a simple, deterministic and objective rule for the case $n = 0$. Finally, we computerized the judge's role to make sure that sampling did not depend on (possibly incorrect) beliefs about the judge's choice.

At the beginning of each session, instructions were read by the experimenter standing on a stage in the front of the experiment room, which fully explained the rules, information structure, and computer interface. After the instructions were finished, two practice matches were conducted, for which subjects received no payment. After the practice matches, there was an interactive computerized comprehension quiz that all subjects had to answer correctly before proceeding to the paid matches. Subjects then participated in 40 paid matches, with opponents and roles (red or blue) randomly reassigned and urns randomly drawn at the beginning of each match. The design included four blocks of ten matches, where the costs pair (c_R, c_B) was identical within blocks and different across blocks. The four cost pairs were the same in all sessions. However, to control for order effects, the sequences were different. Subjects were paid the sum of their earnings over all 40 paid matches, in cash, in private, immediately following the session. Sessions averaged one hour in length, and subject earnings averaged $25. Table 1 displays the pertinent details of the eight sessions.

4 Results

4.1 Aggregate sampling frequencies

Using Proposition 1, we can compute the theoretical levels of sampling as a function of the costs of both players. This can serve as a benchmark for comparison with the empirical behavior. Recall that h^* and $-l^*$ correspond to the states where the red and blue player stop sampling, respectively (see Figure 1). The result is reported in Table 2.

The first cut at the data consists of determining the empirical probabilities of sampling by the blue and red players as a function of n, the difference between the number of blue

[8] The payoff conversion rate was 200 points = $1.00.

Session (date)	# subjects	costs (c_R, c_B) in matches 1-10	11-20	21-30	31-40
1 (06/03/2008)	8	(3,3)	(3,13)	(13,3)	(13,13)
2 (06/04/2008)	10	(3,3)	(13,13)	(3,13)	(13,3)
3 (06/09/2008)	10	(3,13)	(3,3)	(13,3)	(13,13)
4 (06/09/2008)	10	(3,13)	(13,3)	(13,13)	(3,3)
5 (06/11/2008)	10	(13,3)	(3,13)	(13,13)	(3,3)
6 (06/12/2008)	10	(13,3)	(13,13)	(3,3)	(3,13)
7 (06/16/2008)	10	(13,13)	(3,3)	(3,13)	(13,3)
8 (06/16/2008)	10	(13,13)	(13,3)	(3,3)	(3,13)

Table 1: Session details.

(c_R, c_B)	$-l^*$	h^*
$(3,3)$	-3	3
$(3,13)$	-2	3
$(13,3)$	-4	1
$(13,13)$	-2	1

Table 2: Nash equilibrium.

draws and the number of red draws, which from now on will be referred to as the "state". Table 11 shows the empirical sampling frequencies and the Nash equilibrium predictions (reported in Table 2) for each cost pair and pooling all eight sessions together. A graphical representation of the same data is provided in Figure 2. Although, the state space is $n \in \{-6, \dots, 7\}$ in the data, we restrict the analysis to $n \in \{-4, \dots, 4\}$, only because there are few observations for choices in states outside this range (between 0 and 15 depending on the cost treatment).

[Insert Table 11 and Figure 2 here]

Despite the data being rather coarse, it allows us to draw two main conclusions. First, subjects understand the fundamentals of the game. Indeed, the theory predicts that both subjects should never simultaneously draw balls. It is a dominated strategy for blue to draw when $n \geq 0$ and for red to draw when $n \leq -1$. Among the 7879 observations where both players had to simultaneously choose whether to sample, only in 4.8% of the cases the player ahead in the game did draw a ball. Furthermore, 3.2% correspond to a blue subject drawing when $n = 0$, which may well be due to a misunderstanding of the tie-breaking rule. For the rest of the analysis and except otherwise noted, we will focus on the sampling strategy of the player behind in the game (red when $n \geq 0$ and blue when $n \leq -1$).

Second, sampling behavior is reasonably close to Nash equilibrium predictions. Using Table 11, we can determine the number of instances where the subject behind in the game played according to the predictions of the theory. We separate the analysis in two groups. First, the aggregate data. These include all the observations of the subject behind in the game, separated into the cases where the theory predicts that the subject should draw and the cases where the theory predicts that the subject should not draw (the data is then pooled across roles). Second, the 'marginal states.' These include the observations in the last state where the theory predicts that the subject should draw and the observations in the first state where the theory predicts that the subject should not draw. The data is compiled in Table 3 (standard errors clustered at the individual level reported in parentheses and number of observations reported in brackets).

	all states behind	marginal states
Theory is DRAW	.920 (.007) [$N = 6130$]	.870 (.013) [$N = 2404$]
Theory is NO DRAW	.609 (.029) [$N = 1819$]	.549 (.031) [$N = 1156$]
All	.849 (.008) [$N = 7949$]	.766 (.010) [$N = 3560$]

Table 3: Proportion of Nash Equilibrium behavior when player is behind.

The proportion of total observations consistent with the theoretical predictions is high, 85%, especially given that we only consider the choices of the individual behind in the game. Also, there is a substantially lower proportion of under-sampling (8%) than over-sampling (39%). There are at least two reasons for such asymmetry. First and by construction, the number of observations where the player behind has to make his initial sampling decision (red in $n = 0$ or blue in $n = -1$) is relatively large. In these states, subjects almost invariably follow the theoretical prediction and sample (between 92% and 100% of the time). Perhaps more subtly, in equilibrium, a subject can only under-sample once in each match, unless his rival chooses to sample despite being ahead (a rare event). By contrast, he can keep over-sampling indefinitely. This, together with the lower number of observations where the theory predicts no sampling, can partly account for the asymmetry in percentages.

To investigate under- and over-sampling in more detail, we also study the decision of subjects in the marginal states, where they are supposed either to draw for the last time or not draw for the first time. By definition, the cost-benefit analysis is most difficult to perform in these states, so we can expect the greatest number of mistakes. As Table 3 shows, there are 8% fewer observations consistent with the theory than when all states are considered. If we divide the analysis into under- and over-sampling, then the increase in mistakes is around 5% in the former and 6% in the latter.[9] Therefore, although the fraction of mistakes is non-negligible (13% and 45%), we can still argue that behavior is reasonably consistent with the theory, especially for the 'no draw' case. Notice that the same caveat applies as previously. For example, a red player who decides to stop in state $x + 1$ when theory predicts that he should stop in x will, on average, produce several observations of over-sampling in a given match.

Finally, we can also determine the optimal strategy of an individual who knows the empirical sampling frequencies of the population. The problem turns out to be challenging because, contrary to the theoretical model, both subjects sometimes sample simultaneously, and therefore move the state from x to $x \pm 2$. Using numerical methods, we computed the best response to the empirical strategies for each player (blue and red) and in each cost treatment (c_R, c_B). In all eight cases, the best response coincides with the Nash equilibrium play described in Table 2. This result provides further support to the idea that subjects' choices are close to the theoretical predictions. Indeed, if the strategies of players were to depart systematically and substantially from Nash equilibrium, the best responses of the rivals would also depart from the Nash equilibrium behavior. The results of this section are summarized as follows.

Result 1 *The empirical behavior is close to the theoretical prediction in the action space. Best response to the empirical strategies coincide with Nash equilibrium behavior. Deviations are infrequent and occur in both directions (under- and over-sampling).*

[9] Although the differences are small in magnitude, they are all statistically significant given the large number of observations.

4.2 Aggregate payoffs

The next step consists in determining the expected payoffs of players in the states where each of them should start sampling (blue at $n = -1$ and red at $n = 0$) under different scenarios. More precisely, we compute three cases: (1) the expected payoffs given the empirical behavior of both players; (2) the expected payoffs if both subjects played according to the Nash equilibrium; and (3) the expected payoff of a player who best responds to the empirical strategy of the rival which, given our previous result, coincides with the Nash equilibrium play. To facilitate comparisons, we normalize the payoffs of losing and winning the match to 0 and 100 respectively. The results are summarized in Table 4.[10]

state	$n = -1$				$n = 0$			
(c_R, c_B)	$(3,3)$	$(3,13)$	$(13,3)$	$(13,13)$	$(3,3)$	$(3,13)$	$(13,3)$	$(13,13)$
BLUE payoff								
(1) Empirical	20.0	-1.2	28.6	5.3	47.4	33.0	62.4	46.7
(2) Nash	18.6	0.5	32.1	11.9	43.0	30.3	66.0	55.9
(3) Best response	22.6	3.3	31.0	9.4	49.7	36.6	64.2	50.3
RED payoff								
(1) Empirical	64.2	73.2	46.8	60.1	35.3	40.8	7.3	14.7
(2) Nash	63.4	75.2	43.1	64.0	36.0	44.2	8.5	19.0
(3) Best response	66.6	75.8	50.2	62.2	38.5	44.7	12.3	18.4

Table 4: Expected payoffs of BLUE and RED players at $n = -1$ and $n = 0$.

Comparing (1) and (3), we notice that by deviating from the best response strategy, subjects lose at most 3.9 points if their drawing cost is low, and at most 5.0 points if their drawing cost is high. This is relatively small given that the difference between winning and losing is 100 points and that the cost per draw is 3 or 13 points. As discussed in section 4.1, it suggests that subjects are not far from best responding to the strategy of their rivals. Comparing (1) and (2), we notice that the empirical choices of subjects translate into net gains relative to the Nash equilibrium in 5 cases and net losses in the other 11, with the magnitudes being always rather small. This provides further evidence that sampling errors occur in both directions. Indeed, recall that the sum of benefits

[10] For more extreme states, the analysis is less informative: payoffs are mostly driven by costs so the differences between the three cases is small (data not reported but available upon request). Below we perform what we think is a more interesting comparison of payoffs in the marginal states.

is constant across matches. Joint under-sampling is likely to result in lower costs and therefore higher average payoffs for both players whereas joint over-sampling is likely to result in higher costs and therefore lower average payoffs (of course, all this assumes a roughly symmetric under- and over-sampling behavior).

The previous comparisons are suggestive but incomplete. Indeed, there can be two very different reasons for observing small payoff differences between predicted and empirical choices: behaviors are close to each other or payoff functions are flat so even large departures imply small utility losses. In order to evaluate the cost of deviating from equilibrium behavior, we conduct the following numerical analysis. We fix the cost treatment, assume that the first player follows the empirical strategy and that the second player best responds to it (which, remember, also corresponds to the Nash equilibrium) at all states but n. We then determine the expected payoff in state n of the second player if he also plays the equilibrium strategy at n and if he plays the alternative strategy.[11] This exercise captures how much is lost by deviating from best response in one and only one state. The results are summarized in Table 12. We highlight in bold the payoffs given equilibrium play at all states. So, for example, since $h^* = 3$ for the red player in the (3,3) treatment, the bold value is for "draw" in states $n \in \{0, 1, 2\}$ and for "no draw" otherwise. As before, the payoffs of winning and losing are normalized to 100 and 0 respectively.

[Insert Table 12 here]

The most interesting information in this table consists in determining the utility loss of under-sampling and over-sampling in the marginal states, for each pair of costs and each role. We can notice a wide spread in the cost of one-unit deviations, which ranges from 0.6 to 17.1 points across treatments. Also, there are no systematic patterns on the relative losses of under- and over-sampling within a treatment. Under-sampling is more costly than over-sampling in 5 cases and less costly in the other 3. Erring on either side sometimes results in similar costs (3.4 vs. 2.6 points) and some other times in substantially different ones (17.9 vs. 1.9 points). This may be simply due to the integer nature of the sampling strategies considered. Indeed, when the optimal stopping point is somewhere between $x - 1$ and x, the individual obtains a similar payoff when he stops at either of these thresholds. In that respect, using a discrete information accumulation process

[11] Notice that he may reach state n several times. The assumption is that he either always or never plays the equilibrium strategy.

makes the model more intuitive and easier to explain to subjects but, at the same time, introduces integer effects that can have an impact on the results. Finally, we performed the same computations as in Table 12 with one difference. Instead of best responding, we assumed that the second player followed the empirical strategy at all states but n and then determined the expected payoff given drawing at n and given not drawing at n. The results were very similar so for the sake of brevity we decided not to report them (they are available upon request). We summarize the findings of this section as follows.

Result 2 *The empirical behavior is close to the theoretical prediction in the payoff space.*

4.3 Deviations

We now explore in more detail the deviations from Nash equilibrium observed in the data. We start with an analysis of the subjects' actions. From inspection of Table 11 and Figure 2, it is apparent that the main difference with the theoretical prediction is the absence of a sharp decline in the likelihood of sampling around the equilibrium level. In Table 5, we separate the marginal states into the last state where subjects are supposed to draw (that is, $h^* - 1$ for red and $l^* + 1$ for blue, which we call 'last draw') and the first state where subjects are supposed not to draw (that is, h^* for red and l^* for blue, which we call 'first no draw'). We then report the proportion of sampling in each of these two cases as well as the percentage decline between the two.

subject	RED				BLUE			
(c_R, c_B)	$(3,3)$	$(3,13)$	$(13,3)$	$(13,13)$	$(3,3)$	$(3,13)$	$(13,3)$	$(13,13)$
Last DRAW	.62	.67	.93	.97	.77	.92	.52	.97
First NO DRAW	.30	.48	.52	.55	.45	.31	.37	.33
Decline	.52	.29	.44	.43	.42	.66	.29	.66

Table 5: Percentage of draw in marginal states.

Instead of a 100% decline, we observe in the data a decline of 29% to 66%. There are at least two reasons for this smooth pattern. One is a significant heterogeneity in individual behavior. Although it is worth noting this possibility, we will not conduct a detailed individual analysis. Indeed, since the observed behavior is close to the theoretical prediction, we feel that the added value of an exhaustive exploration at the individual level would be rather small. The second reason is related to the integer nature of the sampling strategy, and the idea that when the optimal stopping point is between two cutoffs then similar payoffs may be obtained by stopping at either of them (see the discussion in section 4.2). Notice that subjects draw with a substantially higher probability in $h^* - 1$ and $l^* + 1$ when their cost is high than when it is low. Also, in 3 out of 4 cases, their percentage decrease is also greater. This suggests that under-sampling is more frequent and the decline in drawing around the equilibrium is less steep when subjects have low costs than when they have high cost.

To further explore how costs affect deviations from equilibrium, we perform the same analysis as in Table 3, except that we separate the proportion of Nash play according to the subject's own cost. The results are displayed in Table 6 (standard errors clustered at the individual level are reported in parentheses and the number of observations are reported in brackets).

	all states behind		marginal states	
	LOW cost	HIGH cost	LOW cost	HIGH cost
Theory is DRAW	.909 (.008) [$N = 4204$]	.943 (.010) [$N = 1926$]	.727 (.029) [$N = 779$]	.939 (.011) [$N = 1625$]
Theory is NO DRAW	.585 (.042) [$N = 535$]	.619 (.029) [$N = 1284$]	.592 (.043) [$N = 387$]	.528 (.034) [$N = 769$]
All	.873 (.007) [$N = 4739$]	.814 (.015) [$N = 3210$]	.682 (.016) [$N = 1166$]	.807 (.013) [$N = 2394$]

Table 6: Proportion of Nash behavior by subjects' own cost.

When we pool together all states where the subject is behind, the results are similar for low and high costs, simply because in non-marginal states subjects generally play close to the Nash equilibrium predictions. More interestingly, in the marginal states, under-sampling is overall infrequent and more pronounced with low than with high costs (27% vs. 6%). Over-sampling is more frequent and slightly more pronounced with high than with low costs (47% vs. 41%).

Next, we study how deviations affect payoffs in the different cost treatments. Comparing (1) and (2) in Table 4, we notice that for the (13,13) treatment, the Nash payoffs exceed the empirical payoffs of players in all four cases. By contrast, for the (3,3) treatment the empirical payoffs exceed the Nash payoffs of players in three out of four cases. This is consistent with the sampling biases discussed previously: joint under-sampling in the (3,3) treatment results in lower costs for both players and similar benefits whereas joint over-sampling in the (13,13) treatment results in higher costs for both players and similar benefits.[12] The result is confirmed if we compare Nash equilibrium and best response to empirical behavior. When the red player's cost is low, the blue player gets a higher payoff in (3) than in (2) whereas when the red player's cost is high, the blue player gets a higher payoff in (2) than in (3). Since in both cases the blue player is choosing the same (optimal) strategy, this reinforces the idea that the red player has a tendency to under-sample when his cost is low and over-sample when his cost is high. The same result applies for the red player when the blue player has cost 3 but not when the blue player has cost 13 (in that case, payoffs are almost identical in all four cases). However and as previously noted, payoff differences are generally small.

Finally, it is also instructive to compare the utility loss incurred by deviating from best response for players with high and low cost of sampling. Using Table 12, we notice that in 3 out of 4 observations, the utility loss for the low cost player is bigger with under-sampling than with over-sampling. Conversely, in 3 out of 4 observations, the utility loss for the high cost player is bigger with over-sampling than with under-sampling. In either case, the average difference is relatively small. Also, either type of deviation implies generally a greater loss for a player with a high cost than for a player with a low cost: averaging across deviations and roles, the loss is 10.6 when $c = 13$ and 3.1 when $c = 3$. The reason for such difference is easily explained in the case of over-sampling by the direct cost of drawing, but it also occurs for under-sampling. Last, notice that the deviations we observe in the data are precisely the ones that imply higher utility losses: under-sampling for low cost and over-sampling for high cost. The result is summarized as follows.

Result 3 *The decline in sampling around the theoretical equilibrium is smoother than predicted by theory. There is under-sampling by subjects with low cost and over-sampling by subjects with high cost. In general, over-sampling is more pronounced than under-sampling.*

[12] The asymmetric cost cases are more difficult to interpret. Over-sampling by the high cost player implies a lower expected payoff for the low cost player independently of his choice, but also a lower marginal value of sampling.

4.4 Comparative statics

We now study whether the basic comparative statics predicted by the theory are observed in the data. To this purpose, we first run probit regressions to compute the probability of sampling by a player as a function of the state. We only include states where the subject is behind to ensure a monotonic theoretical relation.[13] For each role, we perform the regression on four subsamples, taking either the player's own cost or the rival's cost as fixed. In the former case, we introduce a dummy variable that codes whether the rival's cost was high (high rival c). In the later case, we introduce a dummy variable that codes whether the player's own cost was high (high own c). We also analyze sequencing effects by including a dummy variable that codes whether the particular cost treatment occurred in the first 20 or the last 20 matches of the experiment (seq. late). Furthermore, remember that subjects played 10 consecutive matches with the same cost pairs. We study a simple version of experience effects by introducing a dummy variable that separates the first 5 matches form the last 5 matches within a given cost pair (exp). We also include interactions terms. The results are summarized in Table 7. Standard errors clustered at the individual level are reported in parentheses (* denotes significant at the 5% level and ** denotes significant at the 1% level).

Not surprisingly, as the difference between unfavorable and favorable draws increases, subjects are less inclined to sample. The effect is strong and highly significant in all eight subsamples. Similarly, as a player's cost increases, his likelihood of sampling decreases. Again, the effect is strong and significant at the 1% level in all four subsamples. The strategic effect of the rival's cost on the behavior of a player is more involved. Proposition 1 states that thresholds are strategic substitutes, so a higher cost by one player translates into more sampling by the other. However, due to the integer constraints, the theory predicts that an increase in the cost of the red player should translate into a higher level of sampling by the blue player if his cost is low and to no change in sampling if his cost is high (see Table 2). This is precisely what we observe in the data with the coefficient 'high rival c' for the blue player being positive in both cases but significant only when $c_B = 3$. For the red player, the integer constraint implies no increase in sampling when the blue

[13] Also, we already know from the previous analysis that behavior is almost invariably in accordance with theory when the player is ahead.

	BLUE				RED			
	$c_B = 3$	$c_B = 13$	$c_R = 3$	$c_R = 13$	$c_R = 3$	$c_R = 13$	$c_B = 3$	$c_B = 13$
constant	2.34**	1.58**	2.30**	3.00**	3.11**	1.78**	3.50**	3.10**
	(.269)	(.389)	(.422)	(.347)	(.218)	(.340)	(.294)	(.344)
draws behind	-.710**	-.642*	-.641**	-.822**	-.990**	-.713**	-1.02**	-.779**
	(.144)	(.282)	(.216)	(.192)	(.090)	(.192)	(.126)	(.148)
seq. late	.241	.889	.199	.579	-.261	1.79**	-.163	.251
	(.216)	(.460)	(.371)	(.366)	(.335)	(.355)	(.345)	(.272)
draw × seq.	-.121	-.655*	-.140	-.381	.134	-1.19**	.097	-.179
	(.108)	(.316)	(.237)	(.219)	(.147)	(.209)	(.173)	(1.42)
exp.	.011	.541	.388	-.179	.045	.594*	.184	.141
	(.219)	(.410)	(.296)	(.354)	(.230)	(.300)	(.255)	(.243)
draw × exp.	.041	-.331	-.187	.125	.006	-.346	-.096	-.028
	(.114)	(.291)	(.186)	(.202)	(.098)	(.182)	(.131)	(.136)
high own c	—	—	-.836**	-1.17**	—	—	-1.26**	-1.08**
			(.165)	(.164)			(.118)	(.157)
high rival c	.244*	.022	—	—	.150	.248*	—	—
	(.110)	(.100)			(.096)	(.110)		
adj. R^2	0.28	0.27	0.25	0.32	0.35	0.36	0.35	0.33

Table 7: Probit regression on the probability of sampling.

player's cost increases both when $c_R = 3$ and when $c_R = 13$. In the data, the coefficient is significant when the red player's cost is high. Overall, all four coefficients for 'high rival c' are positive but two are significant even though only one should be. The analysis of experience and sequencing in this regression are deferred to the next subsection.

We next perform a different comparative statics on sampling as a function of costs. For each state n, we compare the average level of sampling across the different cost treatments. The results are summarized in Table 13. The table can be read as follows. For each state n, we consider only the player behind in the game (since we know that the other player virtually never samples). We then compute the empirical average difference in sampling between the column cost pair treatment and the row cost pair treatment. The number in parentheses beneath each average is the p-value for the statistical significance of the difference. Finally, we report in brackets the theoretical prediction: no change in sampling [o], a 100% decrease in sampling [-], or a 100% increase in sampling [+].

[Insert Table 13 here]

For each state, we then compare the empirical and theoretical change in sampling between cost pairs. Note that theory predicts either no change or a 1.0 change in the probability of sampling. We therefore code a (positive or negative) empirical change in probability as 'significant' when (i) the magnitude of the (positive or negative) change is at least 0.1, and (ii) the change is statistically significant at the 5% or 1% level.[14] Using this criterion and a 5% significant level, we obtain that 23 out of 24 mean comparisons for the red player follow the patterns predicted by theory: no difference in 15 cases, and a statistically significant decrease in 8 cases (with a 1% significant level, all 24 comparisons follow the predictions by theory). For the blue player, 21 out of 24 mean comparisons follow the patterns predicted by theory: no difference in 15 cases, a decrease in 4 cases, and an increase in 2 cases. The 3 misclassified observations are for $n = 3$. It is mainly due to an insufficient level of sampling in the $(13, 3)$ treatment and an excessive level of sampling in the $(3, 3)$ treatment, where the empirical draw rates are 0.52 and 0.45 whereas the predicted rates are 1.0 and 0.0. Notice that our method controls neither for joint correlation between tests (when one sampling departs significantly from the theory, several comparisons are affected) nor for multiplicity of tests (we make 48 comparisons at a 5% significance level). The fact that 44 out of 48 are correctly classified suggests that the comparative statics are in accordance with theory. The results of this section are summarized as follows.

Result 4 *The comparative statics follow the predictions of the theory both in aggregate and state-by-state: subjects sample more when their cost is low and when the cost of their rival is high.*

4.5 Learning

We now study whether subjects change their behavior over the course of the experiment. We know the proportion of mistakes by subjects ahead is low (4.8%). It is nevertheless instructive to determine how these mistakes evolve over time. The proportion of mistakes is 6.9% in the first 20 matches and 2.6% in the last 20 matches of the experiment; this suggests that subjects learn to avoid basic mistakes almost entirely as the experiment progresses.

[14] In other words, a decrease in sampling from 1.00 to 0.97 (as between $(3, 3)$ and $(13, 13)$ for $n = 0$) is not coded as a change even if the 0.03 difference is significant at the 1% level. In our view, such a case is closer to no change than to 1.0 decrease.

We then move on to the more interesting case of subjects who are behind in the game. A simple approach to determine changes in behavior is to divide the sample into early sequences (1 and 2) and late sequences (3 and 4) or into inexperienced (first 5 matches within a cost pair) and experienced (last 5 matches). We then determine the proportion of Nash equilibrium play in each subsample. The results are complied in Tables 8 and 9 (as in previous tables, standard errors clustered at the individual level are reported in parentheses and number of observations are reported in brackets).

	all states behind		marginal states	
	seq. 1 & 2	seq. 3 & 4	seq. 1 & 2	seq. 3 & 4
Theory is DRAW	.915 (.009) [$N = 3064$]	.925 (.008) [$N = 3066$]	.861 (.018) [$N = 1203$]	.879 (.014) [$N = 1201$]
Theory is NO DRAW	.571 (.036) [$N = 972$]	.653 (.029) [$N = 847$]	.506 (.040) [$N = 571$]	.591 (.035) [$N = 585$]
All	.832 (.011) [$N = 4036$]	.866 (.008) [$N = 3913$]	.747 (.014) [$N = 1774$]	.785 (.013) [$N = 1786$]

Table 8: Proportion of Nash behavior by sequence.

	all states behind		marginal states	
	inexperienced	experienced	inexperienced	experienced
Theory is DRAW	.915 (.008) [$N = 3112$]	.925 (.008) [$N = 3018$]	.860 (.015) [$N = 1239$]	.881 (.015) [$N = 1165$]
Theory is NO DRAW	.609 (.033) [$N = 896$]	.609 (.028) [$N = 923$]	.557 (.035) [$N = 580$]	.542 (.034) [$N = 576$]
All	.847 (.009) [$N = 4008$]	.851 (.009) [$N = 3941$]	.764 (.011) [$N = 1819$]	.769 (.013) [$N = 1741$]

Table 9: Proportion of Nash behavior by level of experience.

From Table 8, we notice that over-sampling both in the marginal states and in all states taken together decreases by roughly 8% when the cost treatment under consideration is played late in the experiment. Under-sampling remains mostly unaffected, possibly because it is quite low to start with. In all four cases, mistakes are reduced. By contrast, Table 9 suggests that experience within a cost treatment has virtually no effect on the behavior of subjects (no more than a 2% change).

A more rigorous look at the data consists in studying significance of the 'sequence' and 'experience' variables in the OLS regression presented in Table 7. The sequencing effect is significant for both players when their own cost is high. The positive coefficient of 'seq. late' and negative coefficient when combined with the number of draws behind suggests that, when that particular cost pair comes late, subjects sample more if they are behind by few draws and less if they are behind by many draws, as learning would predict.[15] This effect is not present in any of the other six subsamples. The effect of experience is only marginally significant in one of the eight subsamples. Overall, the OLS regression provides some limited evidence of learning due to sequencing and none due to experience.

Our next step is to determine the probability of playing the Nash equilibrium (whether it prescribes sampling or not) as a function of the different variables of the game. This Probit regression is presented in Table 10. Despite some potential specification problems, the results are instructive.

	BLUE		RED	
constant	2.17**	(.096)	2.48**	(.102)
draws behind	-.581**	(.029)	-.596**	(.028)
seq. late	.040	(.065)	.152*	(.068)
exp.	.023	(.056)	.024	(.046)
high own c	-.293**	(.080)	-.757**	(.075)
high rival c	.149*	(.063)	.071	(.049)
adj. R^2	0.16		0.18	

Table 10: Probit regression on the probability of Nash choice.

The variables that affect more significantly the probability of playing the Nash equilibrium are the state and the player's own cost. As we saw in section 2.1, over-sampling is more severe than under-sampling. The negative coefficient in these two variables reflect this asymmetry in the deviations. The coefficients for the learning variables are very much in line with the results of the descriptive summaries of Tables 8 and 9, and the OLS

[15]The p-value of 'seq. late' for the blue player with high cost is 0.054. The other three are below 5%.

regression of Table 10: experience is not significant whereas sequencing is only significant for the red player.

All in all, there is limited evidence of changes in sampling behavior over trials. One possible explanation is simply that subjects play relatively close to equilibrium right from the outset, so there is little loss in payoffs and also little room for learning. The result is summarized as follows.

Result 5 *Subjects ahead in the game learn to avoid sampling mistakes almost entirely. Subjects behind in the game exhibit limited learning over the course of the experiment.*

5 Conclusion

In this paper, we have analyzed a model of information acquisition by agents with opposite interests. We have characterized the Nash equilibrium of the game and shown that the choice variables are strategic substitutes: if the incentives to collect information of one individual increase, then the incentives of his rival decrease. We have tested the theory in a controlled environment. Behavior of subjects is close to Nash predictions, except that the decline in sampling around the equilibrium is smoother than predicted by theory. Mistakes are relatively infrequent, and take more often the form of over-sampling than under-sampling. Comparative statics on the subject's own cost and on the rival's cost generally follow the predictions of the theory both at the aggregate level and state-by-state. Finally, there is little evidence of learning.

The study can be extended in several directions. Form a theoretical viewpoint, it would be interesting to combine the acquisition of information and the revelation of information paradigms. In particular, one could extend the literature on games of persuasion to incorporate a sequential process of acquisition of private pieces of evidence. This would allow us to determine the optimal stopping rule given the anticipated future use of private information. From an experimental viewpoint, the similarity between empirical behavior and theoretical predictions is intriguing. It would be interesting to study behavior in yet more sophisticated environments. One possibility would be to consider three players. When the evidence favors one agent, which of the other two will be more likely to acquire information and which one will be more tempted to free-ride? Another possibility would be to allow subjects to engage in agreements with collusive side transfers that replace information acquisition. Because paying for information is inefficient from their joint viewpoint, the theory would predict always agreement and no sampling. In the experiment, will these agreements happen frequently? When they occur, are the payoffs of each player above their expected return in the non-cooperative Nash equilibrium with sampling? These and other related questions are left for future research.

Appendix A1: proof of Proposition 1

It is immediate to notice that the blue advocate will never sample if $n \geq 0$ and the red advocate will never sample if $n \leq -1$. Also, if at some stage no advocate finds it optimal to sample, no information is accumulated so it cannot be optimal to restart sampling. Suppose now that the event is $S = B$ and the state (the difference between blue and red draws) is $n \in \{0, \ldots, h-1\}$, where h is the value where the red advocate gives up sampling (we will determine this optimal value below). The value function of the red advocate, denoted $g_B^r(n)$, satisfies the following second-order difference equation with constant term:

$$g_B^r(n) = \theta\, g_B^r(n+1) + (1-\theta) g_B^r(n-1) - c_R.$$

where θ $(1-\theta)$ is the probability of receiving signal β (ρ) given that the event is B, thereby moving the state to $n+1$ $(n-1)$. Applying standard methods to solve for the generic term of this equation, we get:

$$g_B^r(n) = y_1 + y_2\, \lambda^n + F_R\, n \tag{1}$$

where $\lambda = (1-\theta)/\theta$ and $F_R = c_R/(2\theta - 1)$. In order to determine the constants (y_1, y_2), we need to use the two terminal conditions. By definition, we know that at $n = h$ the red advocates gives up and gets 0. Therefore, $g_B^r(h) = 0$. The lower terminal condition is more intricate. We have: $g_B^r(-1) = q_B^b \pi_R + (1-q_B^b) g_B^r(0)$, where q_S^b is the probability that the blue advocate reaches $n = -l$ before reaching $n = 0$ given event $S \in \{R, B\}$ and state $n = -1$. In other words, the red advocate knows that when $n = -1$, the blue advocate will restart sampling (thus the red advocate will stop paying costs). With probability q_B^b, the belief will reach $n = -l$. The blue advocate will stop at that point and the red advocate will obtain the payoff π_R. With probability $1 - q_B^b$, the belief will go back to $n = 0$. The value function of the red advocate will then be $g_B^r(0)$ and he will have to start sampling again. For the time being, let's take q_B^b as exogenous (naturally, we will need to determine later on what this value is). Using (1) and the two terminal conditions, we obtain a system of two equations ($g_B^r(h)$ and $g_B^r(-1)$) with two unknowns (y_1 and y_2). Solving this system, we can determine the values (y_1, y_2) which, once they are plugged back into (1), yield:

$$g_B^r(n) = \Big(\pi_R\, q_B^b + F_R(1 + h\, q_B^b)\Big) \left[\frac{\lambda^{n+1} - \lambda^{h+1}}{1 - \lambda + \lambda(1-\lambda^h) q_B^b}\right] - F_R(h - n) \tag{2}$$

When the event is $S = R$, the second-order difference equation for the red advocate is:

$$g_R^r(n) = (1-\theta) g_R^r(n+1) + \theta\, g_R^r(n-1) - c_R$$

where the only difference is that the likelihood of moving the state to $n+1$ $(n-1)$ is now $1-\theta$ (θ). Solving in an analogous fashion, we get:

$$g_R^r(n) = \Big(\pi_R\, q_R^b - F_R(1 + h\, q_R^b)\Big) \left[\frac{1-\lambda^{h-n}}{\lambda^h(1-\lambda) + (1-\lambda^h) q_R^b}\right] + F_R(h-n) \tag{3}$$

At this point, we need to determine q_S^b. Recall that the blue advocate gives up at $n = -l$ (where $-l$ will be determined below). Let $h_S^b(n)$ denote the blue advocate's probability of reaching $n = -l$ before $n = 0$ given event S and a starting state n. Using the by now familiar second-order difference equation method, we have:

$$h_B^b(n) = \theta\, h_B^b(n+1) + (1-\theta) h_B^b(n-1) \quad \text{with} \quad h_B^b(-l) = 1 \quad \text{and} \quad h_B^b(0) = 0$$

and

$$h_R^b(n) = (1-\theta) h_R^b(n+1) + \theta\, h_R^b(n-1) \quad \text{with} \quad h_R^b(-l) = 1 \quad \text{and} \quad h_R^b(0) = 0$$

Note that $h_S^b(\cdot)$ captures exclusively the blue advocate's likelihood of reaching each stopping point ($-l$ or 0), that is, it does not take costs into consideration. This is the case because in the red advocate's calculation only the probabilities matter (not the net utility of the blue advocate). Solving for the generic term in a similar way as before, we now get:

$$h_B^b(n) = \frac{\lambda^{l+n} - \lambda^l}{1-\lambda^l} \quad \text{and} \quad h_R^b(n) = \frac{1-\lambda^{-n}}{1-\lambda^l}.$$

This implies that:

$$q_B^b \equiv h_B^b(-1) = \frac{\lambda^{l-1} - \lambda^l}{1-\lambda^l} \quad \text{and} \quad q_R^b \equiv h_R^b(-1) = \frac{1-\lambda}{1-\lambda^l}$$

Inserting the expressions of q_B^b in (2) and q_R^b in (3), we can finally determine $g_B^r(n)$ and $g_R^r(n)$ as a function of the parameters of the model.

Note that $\Pr(B \,|\, n) = \mu(n) = \frac{1}{1+\lambda^n}$ and $\Pr(R \,|\, n) = 1 - \mu(n) = \frac{\lambda^n}{1+\lambda^n}$. The expected payoff of the red advocate given state $n \in \{0, ..., h-1\}$, is then:

$$\begin{aligned} \Pi_n^r(l,h) &= \Pr(B \,|\, n)\, g_B^r(n) + \Pr(R \,|\, n)\, g_R^r(n) \qquad (4) \\ &= \frac{1}{1+\lambda^n}\left[\Big(\pi_R(1+\lambda^l) - F_R(h+1)(1-\lambda^l)\Big)\left[\frac{\lambda^n - \lambda^h}{1-\lambda^{h+l}}\right] - F_R(h-n)(1-\lambda^n)\right] \end{aligned}$$

A similar method can be used to determine the expected payoff of the blue advocate when the state is $n \in \{-l+1, ..., -1\}$, with the only exception that sampling is stopped at $n = -1$ rather than at $n = 0$. We then get:

$$\begin{aligned} \Pi_n^b(l,h) &= \Pr(B \mid n)\, g_B^b(n) + \Pr(R \mid n)\, g_R^b(n) \\ &= \frac{1}{1+\lambda^n}\left[\Big(\pi_B(1+\lambda^h) - F_B(l-1)(1-\lambda^h)\Big)\left[\frac{1-\lambda^{n+l}}{1-\lambda^{h+l}}\right] + F_B(n+l)(1-\lambda^n)\right] \end{aligned} \tag{5}$$

In a Nash equilibrium, the best response functions of the red and blue advocates are

$$h^*(l) = \arg\max_h \Pi_n^r(l,h) \quad \text{and} \quad l^*(h) = \arg\max_l \Pi_n^b(l,h)$$

Taking first-order conditions in (4) and (5), we obtain:

$$-\lambda^{h^*(l)} \ln\lambda \Big[\pi_R(1+\lambda^l) - F_R(h^*(l)+1)(1-\lambda^l)\Big] = F_R(1-\lambda^{h^*(l)})(1-\lambda^{l+h^*(l)})$$

$$-\lambda^{l^*(h)} \ln\lambda \Big[\pi_B(1+\lambda^h) - F_B(l^*(h)-1)(1-\lambda^h)\Big] = F_B(1-\lambda^{l^*(h)})(1-\lambda^{l^*(h)+h})$$

As expected, h^* and l^* do not depend on n, that is, the optimal stopping rules of the two advocates *are not* revised with the realizations of the sampling process. Notice that $\frac{\partial^2\Pi_n^r}{\partial h^2}\Big|_{h^*} < 0$ and $\frac{\partial^2\Pi_n^b}{\partial l^2}\Big|_{l^*} < 0$, so h^* and l^* are indeed maxima. Note also that $\frac{\partial h^*}{\partial \pi_R} \propto \frac{\partial^2\Pi_n^r}{\partial h \partial \pi_R}\Big|_{h^*} > 0$, $\frac{\partial h^*}{\partial c_R} \propto \frac{\partial^2\Pi_n^r}{\partial h \partial c_R}\Big|_{h^*} < 0$, $\frac{\partial l^*}{\partial \pi_B} \propto \frac{\partial^2\Pi_n^b}{\partial l \partial \pi_B}\Big|_{l^*} > 0$, and $\frac{\partial l^*}{\partial c_B} \propto \frac{\partial^2\Pi_n^b}{\partial l \partial c_B}\Big|_{l^*} < 0$, Analogously, $\frac{\partial h^*}{\partial l} \propto \frac{\partial^2\Pi_n^r}{\partial h \partial l}\Big|_{h^*} < 0$ and $\frac{\partial l^*}{\partial h} \propto \frac{\partial^2\Pi_n^b}{\partial h \partial l}\Big|_{l^*} < 0$. □

References

1. Brocas, I. and J.D. Carrillo (2007) "Influence through Ignorance." *RAND Journal of Economics*, Vol.38, pp.931-947.

2. Camerer, C.F. (1995) "Individual Decision Making." in J.H. Kagel and A.E. Roth eds. *The Handbook of Experimental Economics*, pp.587-703.

3. Cox, J.C. and R.L. Oaxaca (2008) "Laboratory Tests of Job Search Models." in C. Plott and V. Smith eds. *Handbook of Experimental Economic Results, Vol.1*, pp.311-318.

4. Forsythe, R., M. Isaac and T.R. Palfrey (1989) "Theories and Tests of 'Blind Bidding' in Sealed-bid Auctions." *RAND Journal of Economics*, Vol.20, pp.214-238.

5. Gul, F. and W. Pesendorfer (2008) "The War of Information." *mimeo*, Princeton University.

6. Matthews, S. and A. Postlewaite (1985) "Quality Testing and Disclosure." *RAND Journal of Economics*, Vol.16, pp.328-340.

7. Milgrom, P. and J. Roberts (1986) "Relying on the Information of Interested Parties." *RAND Journal of Economics*, Vol.17, pp.18-32.

8. Schotter, A. and Y.M. Braunstein (1981) "Economic Search: an Experimental Study." *Economic Inquiry*, Vol.19, pp.1-25.

n (blue draws - red draws)	-4	-3	-2	-1	0	1	2	3	4
$(c_R, c_B) = (3, 3)$									
# observations	51	158	266	439	791	407	245	98	18
Pr[red sampling – Nash]	.00	.00	.00	.00	1.00	1.00	1.00	.00	.00
Pr[red sampling – empirical]	.02	.03	.01	.12	1.00	.92	.62	.30	.33
(standard error)	(.02)	(.01)	(.01)	(.02)	(.00)	(.01)	(.03)	(.05)	(.11)
Pr[blue sampling – Nash]	.00	.00	1.00	1.00	.00	.00	.00	.00	.00
Pr[blue sampling – empirical]	.45	.45	.77	1.0	.14	.01	.00	.00	.00
(standard error)	(.07)	(.04)	(.03)	(.00)	(.01)	(.01)	(.00)	(.00)	(.00)
$(c_R, c_B) = (3, 13)$									
# observations	13	58	223	394	731	394	216	88	26
Pr[red sampling – Nash]	.00	.00	.00	.00	1.00	1.00	1.00	.00	.00
Pr[red sampling – empirical]	.15	.10	.02	.09	.99	.94	.67	.48	.39
(standard error)	(.10)	(.04)	(.01)	(.01)	(.00)	(.01)	(.03)	(.05)	(.10)
Pr[blue sampling – Nash]	.00	.00	.00	1.00	.00	.00	.00	.00	.00
Pr[blue sampling – empirical]	.54	.24	.31	.92	.06	.00	.00	.00	.00
(standard error)	(13.8)	(5.6)	(3.1)	(1.4)	(0.9)	(0.0)	(0.0)	(0.0)	(0.0)
$(c_R, c_B) = (13, 3)$									
# observations	43	124	228	363	624	287	94	7	1
Pr[red sampling – Nash]	.00	.00	.00	.00	1.00	.00	.00	.00	.00
Pr[red sampling – empirical]	.02	.01	.00	.04	.93	.52	.11	.14	.00
(standard error)	(.02)	(.01)	(.00)	(.01)	(.01)	(.03)	(.03)	(.14)	n/a
Pr[blue sampling – Nash]	.00	1.00	1.00	1.00	.00	.00	.00	.00	.00
Pr[blue sampling – empirical]	.37	.52	.87	1.00	.06	.00	.00	.00	.00
(standard error)	(.08)	(.05)	(.02)	(.00)	(.01)	(.00)	(.00)	(.00)	n/a
$(c_R, c_B) = (13, 13)$									
# observations	5	40	171	301	607	259	96	10	3
Pr[red sampling – Nash]	.00	.00	.00	.00	1.00	.00	.00	.00	.00
Pr[red sampling – empirical]	.00	.00	.00	.02	.97	.55	.15	.40	.67
(standard error)	(.00)	(.00)	(.00)	(.01)	(.01)	(.03)	(.04)	(.16)	(.33)
Pr[blue sampling – Nash]	.00	.00	.00	1.00	.00	.00	.00	.00	.00
Pr[blue sampling – empirical]	.80	.13	.33	.97	.09	.00	.00	.00	.00
(standard error)	(.20)	(.05)	(.04)	(.01)	(.01)	(.00)	(.00)	(.00)	(.00)

Table 11: Sampling frequencies.

n (blue - red draws)	-4	-3	-2	-1	0	1	2	3	4
$(c_R, c_B) = (3, 3)$									
RED draw	93.9	87.5	73.0	58.3	**38.5**	**16.1**	**3.4**	-2.6	-3.0
RED no draw	**99.7**	**98.0**	**88.7**	**66.6**	3.1	0.1	0.0	**0.0**	**0.0**
BLUE draw	-3.0	-1.1	**6.1**	**22.6**	43.6	61.1	79.1	91.8	94.5
BLUE no draw	**0.0**	**0.1**	0.1	1.7	**49.7**	**76.3**	**93.9**	**99.3**	**99.9**
$(c_R, c_B) = (3, 13)$									
RED draw	95.1	93.5	79.8	60.5	**44.7**	**19.5**	**4.8**	-1.8	-3.0
RED no draw	**99.9**	**99.7**	**96.9**	**75.8**	1.6	0.0	0.0	**0.0**	**0.0**
BLUE draw	-14.3	-14.1	-17.1	**3.3**	20.9	32.1	54.4	72.9	80.4
BLUE no draw	**0.0**	**0.0**	**0.0**	1.4	**36.6**	**68.9**	**91.0**	**98.3**	**99.8**
$(c_R, c_B) = (13, 3)$									
RED draw	76.8	64.7	44.7	33.4	**12.3**	-10.7	-13.0	-13.0	-13.0
RED no draw	**99.4**	**96.0**	**80.5**	**50.2**	1.3	**0.0**	**0.0**	**0.0**	**0.0**
BLUE draw	-4.1	**0.6**	**9.8**	**31.0**	47.1	70.1	91.7	96.4	96.9
BLUE no draw	**0.0**	0.0	0.0	0.7	**64.2**	**91.6**	**99.6**	**100.0**	**100.0**
$(c_R, c_B) = (13, 13)$									
RED draw	78.4	78.3	57.1	36.6	**18.4**	-6.4	-13.0	-13.0	-13.0
RED no draw	**99.9**	**99.8**	**94.9**	**62.2**	2.4	**0.0**	**0.0**	**0.0**	**0.0**
BLUE draw	-13.0	-13.0	-13.0	**9.4**	26.1	44.9	74.5	80.0	78.6
BLUE no draw	**0.0**	**0.0**	**0.0**	0.5	**50.3**	**87.6**	**99.3**	**99.9**	**100.0**

Table 12: Values to drawing and not drawing by state.

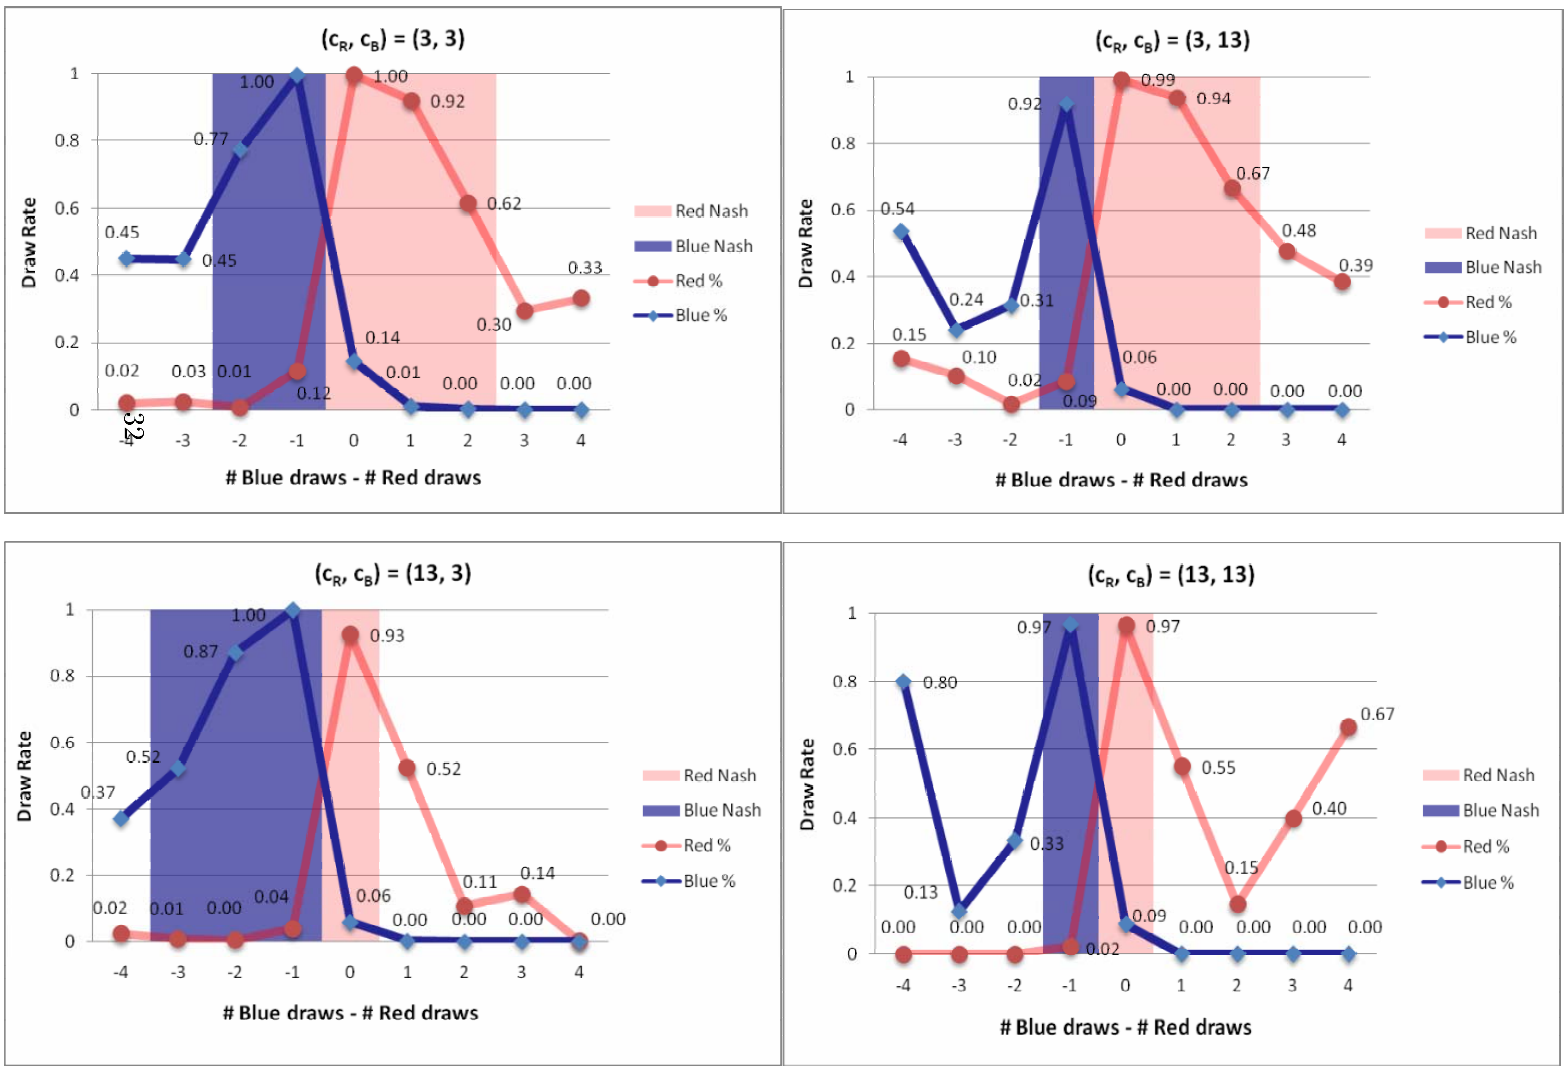

Figure 2. Sampling frequencies by state and cost treatment

RED	$n=0$			$n=1$			$n=2$			$n=3$		
(c_R, c_B)	$(3,13)$	$(13,3)$	$(13,13)$	$(3,13)$	$(13,3)$	$(13,13)$	$(3,13)$	$(13,3)$	$(13,13)$	$(3,13)$	$(13,3)$	$(13,13)$
$(3,3)$	-.004 [o] (.262)	-.076 [o] (.000)	-.031 [o] (.000)	.018 [o] (.326)	-.399 [-] (.000)	-.369 [-] (.000)	.050 [o] (.261)	-.510 [-] (.000)	-.470 [-] (.000)	.181 [o] (.011)	-.153 [o] (.386)	.104 [o] (.496)
$(3,13)$	—	-.067 [o] (.000)	-.026 [o] (.001)	—	-.416 [-] (.000)	-.387 [-] (.000)	—	-.560 [-] (.000)	-.521 [-] (.000)	—	-.334 [o] (.087)	-.077 [o] (.643)
$(13,3)$	—	—	.041 [o] (.002)	—	—	.029 [o] (.490)	—	—	.039 [o] (.413)	—	—	.257 [o] (.252)

BLUE	$n=-1$			$n=-2$			$n=-3$			$n=-4$		
(c_R, c_B)	$(3,13)$	$(13,3)$	$(13,13)$	$(3,13)$	$(13,3)$	$(13,13)$	$(3,13)$	$(13,3)$	$(13,13)$	$(3,13)$	$(13,3)$	$(13,13)$
$(3,3)$	-.074 [o] (.000)	.005 [o] (.198)	-.029 [o] (.002)	-.461 [-] (.000)	.098 [o] (.005)	-.441 [-] (.000)	-.208 [o] (.006)	.075 [+] (.212)	-.324 [o] (.000)	.087 [o] (.573)	-.079 [o] (.439)	-.349 [o] (.136)
$(3,13)$	—	.079 [o] (.000)	.045 [o] (.012)	—	.559 [+] (.000)	-.019 [o] (.683)	—	.283 [+] (.000)	-.116 [o] (.152)	—	-.166 [o] (.285)	.262 [o] (.308)
$(13,3)$	—	—	-.033 [o] (.000)	—	—	-.539 [-] (.000)	—	—	-.399 [-] (.000)	—	—	.428 [o] (.066)

Table 13: Comparison of sampling across treatments (p-values in parentheses)