

# DISCUSSION PAPER SERIES

No. 7457

**INFORMATION GATEKEEPERS:  
THEORY AND EXPERIMENTAL  
EVIDENCE**

Isabelle Brocas, Juan D Carrillo  
and Thomas R Palfrey

***INDUSTRIAL ORGANIZATION and  
PUBLIC POLICY***



**Centre for Economic Policy Research**

**[www.cepr.org](http://www.cepr.org)**

Available online at:

**[www.cepr.org/pubs/dps/DP7457.asp](http://www.cepr.org/pubs/dps/DP7457.asp)**

# 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

Discussion Paper No. 7457  
September 2009

Centre for Economic Policy Research  
53–56 Gt Sutton St, London EC1V 0DG, UK  
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820  
Email: [cepr@cepr.org](mailto:cepr@cepr.org), Website: [www.cepr.org](http://www.cepr.org)

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION and PUBLIC POLICY**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Isabelle Brocas, Juan D Carrillo and Thomas R Palfrey

## ABSTRACT

### Information Gatekeepers: Theory and Experimental Evidence\*

We consider a model where two adversaries can spend resources in acquiring public information about the unknown state of the world in order to influence the choice of a decision maker. We characterize the sampling strategies of the adversaries in the equilibrium of the game. We show that, as the cost of information acquisition for one adversary increases, that person collects less evidence whereas the other adversary 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 adversaries sampling more when their own cost is low and when the other adversary's cost is high. Finally, there is little evidence of learning over the 40 matches of the experiment.

JEL Classification: C91 and D83

Keywords: adversarial system, experiment, information acquisition and search

Isabelle Brocas  
Department of Economics  
University of Southern California  
3620 S. Vermont Avenue,  
Los Angeles, CA 90089-0253  
USA  
Email: [brocas@usc.edu](mailto:brocas@usc.edu)

Juan D Carrillo  
Department of Economics  
University of Southern California  
3620 S. Vermont Avenue,  
Los Angeles, CA 90089-0253  
USA  
Email: [juandc@usc.edu](mailto:juandc@usc.edu)

For further Discussion Papers by this author see:  
[www.cepr.org/pubs/new-dps/dplist.asp?authorid=138720](http://www.cepr.org/pubs/new-dps/dplist.asp?authorid=138720)

For further Discussion Papers by this author see:  
[www.cepr.org/pubs/new-dps/dplist.asp?authorid=131469](http://www.cepr.org/pubs/new-dps/dplist.asp?authorid=131469)

Thomas R Palfrey  
Humanities and Social Sciences  
California Institute of Technology  
Pasadena, CA 91125  
USA  
Email: [trp@hss.caltech.edu](mailto:trp@hss.caltech.edu)

For further Discussion Papers by this author see:  
[www.cepr.org/pubs/new-dps/dplist.asp?authorid=155219](http://www.cepr.org/pubs/new-dps/dplist.asp?authorid=155219)

\* We thank M. Castro, S.H. Chew, G. Frechette, A. Gaduh, D. Houser, E. Kamenica, S. Singhal, and the audiences at the CIRANO Political Economy workshop in Montreal, the Hong Kong University of Science and Technology, the National University of Singapore, the North American meeting of the European Science Association and the Stanford Institute for Theoretical Economics for comments, Dustin Beckett for research assistance, and Chris Crabbe for developing the software. We also thank the financial support of the National Science Foundation (SES-0617820), The Social Science Experimental Laboratory at Caltech, The Gordon and Betty Moore Foundation, the LUSK Center, the Office of the Provost at the University of Southern California and the Microsoft Corporation.

Submitted 03 September 2009

# 1 Motivation

The literature on the economics of information has devoted considerable effort to understand the strategic use of private information by agents in the economy. However, less is known about the *strategic collection of information*, 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 a similar vein, firms acquire and reveal through advertising information about the desirable attributes of their products and the undesirable attributes of the competitors' products. 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 agents with opposite objectives (the adversaries) can acquire costly evidence. When both adversaries choose to stop the acquisition of information, a third agent (the decision maker) makes a binary choice. Formally, there are two possible events. Nature draws one event from a common prior distribution. The adversaries can then acquire signals that are imperfectly correlated with the true event. Information affects the belief about the true event and is public, in the sense that all the news collected by one adversary are automatically shared with the other adversary and the decision maker. The mapping between the information and the decision maker's choice is deterministic and known: it favors the interests of one adversary if the belief about the relative likelihood of states is below a certain threshold, and it favors the interests of the other adversary if the belief is above that threshold.

There are two reasons to assume public rather than private information: realism and simplicity. In the case of jury trials, for example, it is explicitly forbidden that prosecutors and defense attorneys withhold information from other parties. Also, with private acquisition of information, the incentives to acquire and transmit information are inter-related. This complicates both the theoretical and the experimental analyses. Since the existing literature has extensively studied the transmission of information, we choose to focus instead on the acquisition of information.

Opposite incentives implies that adversaries never acquire costly information simultaneously. Indeed, when the current evidence implies that the decision maker will favor the interests of one adversary, extra evidence can only hurt him so he will not have incentives to acquire further information. However, if the current evidence favors the other adversary, then he must trade-off the cost of acquiring more information with the likelihood that such information will revert the decision. As the belief becomes more and more adverse, the likelihood of reverting 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 adversary, that adversary acquires information. He keeps sampling up to a point where either the belief is reversed, in which case the other adversary starts the sampling process over, or else it has become so unfavorable that it is preferable to give up. Solving this problem analytically is non-trivial, since the value function of each adversary depends on the sampling strategy of both adversaries. Indeed, the value of sampling for information in order to ‘send the belief to the other camp’ depends on how intensely the other adversary 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 adversary as a function of the common belief about the state and the cost of sampling for each adversary. We provide an analytical characterization of the Markov equilibrium, and perform some comparative statics: when the cost of news acquisition for an adversary increases, that adversary has fewer incentives to collect evidence, whereas the other adversary has more incentives to collect evidence.

We then report an experiment that analyzes behavior in this information acquisition game with two adversaries. We study variations of the game where each adversary 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 are not. Our first and main result, is that the empirical behavior in all treatments is close to the predictions of the theory both in action space (Result 1) and in payoff space (Result 2). This conclusion is remarkable given that the optimal stopping rule of an adversary is fairly sophisticated, it involves strategic considerations about the other adversary’s choice, and it prescribes corner choices (for a given belief, either never sample or always sample). To be more precise, the optimal action of an adversary who is currently disfavored by the existing evidence depends on whether the common 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 adversary takes the decision predicted by 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 other adversary is to play the Markov equilibrium, which reinforces the idea that deviations from equilibrium play are small. Similar results are obtained when we analyze choices in payoff space: given their empirical behavior, an adversary loses less than 5% of the payoff he would obtain if he best responded to the strategy of the other adversary.

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 adversary’s own cost is low and over-sampling occurs relatively more often when the adversary’s own 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 adversary increases, that adversary is less likely to acquire information whereas the other adversary 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 adversary 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 adversary 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 and the cost-benefit trade-off. It shows how an individual can affect the choices of others by selectively deciding whether to acquire or avoid

public information.<sup>1</sup> 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 firms to acquire verifiable information given the ex-post willingness to reveal it to consumers depending on its content.

On the experimental side, there is an extensive literature on search for payoffs in an individual decision making context (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. Risk aversion may account for the observed insufficient experimentation. Our paper extends that literature in two directions: *search for information* and *search in strategic contexts*. As pointed by Camerer (1995, p.673), a search for information could lead to different conclusions than a search for payoffs, even though both are formally quite similar. This does not seem to be the case in our setting. More importantly, the strategic nature of adversarial search substantially increases the complexity of the decision making problem relative to the individual decision making counterpart. Surprisingly, we still observe a behavior that is close to theory. The main difference is that, with strategic sampling, excessive experimentation occurs more frequently than insufficient experimentation.<sup>2</sup>

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 action space and payoff space, deviations from equilibrium as a function of the costs of both adversaries, comparative statics (aggregate and state-by-state), and learning. In section 5, we provide some concluding

---

<sup>1</sup>Recently, Kamenica and Gentzkow (2009) have approached the one-agent, no-cost model of Brocas and Carrillo (2007) from a mechanism design perspective and determined general conditions on the preferences of players such that the agent with the capacity to collect information can benefit from this option.

<sup>2</sup>Games of persuasion have also been studied in the laboratory. Forsythe, Isaac and Palfrey (1989) find that behavior in the laboratory approaches the prediction of the sequential equilibrium: there is an unraveling process by which sellers are trapped into fully revealing their information because buyers assume the worst whenever sellers withhold it.

remarks. The proof of the proposition is relegated to the appendix.

## 2 The model

### 2.1 The game

Consider a game with three agents. Agent 0 is a *decision maker* (congress, judge, consumer) who must undertake an action that affects the payoff of all three agents. Agents 1 and 2 are *adversaries* (lobbies, advocates, firms) who can collect costly evidence about an event that has realized in order to affect the belief (hence, the action) of the decision maker. We assume that all the information collected by adversaries becomes publicly available. Thus, at any point in time, the decision maker and adversaries share the same belief about which event has realized. However, because adversaries 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. As mentioned before, one reason to choose this assumption is to isolate the incentives for information gathering. In that respect, adding private information would only pollute the analysis.

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 decision maker must choose between two actions,  $a \in \{b, r\}$ . Both the payoff of the decision maker and the payoff of adversary  $i \in \{1, 2\}$  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 adversary  $i$  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 true event. 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 agents 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  $n_\beta$ , the number of  $\beta$ -signals, and  $n_\rho$ , the number of  $\rho$ -signals accumulated by adversaries. 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. From now on, we will refer to  $n$  as the *state*. It is immediate that  $\mu(n+1) > \mu(n)$  for all  $n$ ,  $\lim_{n \rightarrow -\infty} \mu(n) = 0$  and  $\lim_{n \rightarrow +\infty} \mu(n) = 1$ . We assume that from the decision maker'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  $v(b) \geq v(r)$  if and only if  $\mu \geq \mu^*$ . This can be equivalently expressed in terms of the state: there will always exist a state  $n^* \in \mathbb{Z}$  such that  $v(b) \geq v(r)$  if and only if  $n \geq n^*$ .

## 2.2 Optimal stopping rule with two adversaries

Suppose that two adversaries can collect public evidence. For simplicity, suppose that one adversary wants the decision maker to take action  $b$  independently of the event realized, and the other adversary wants the decision maker to take action  $r$  also independently of the event realized.<sup>3</sup> From now on, we call them the blue adversary and the red adversary, respectively. Formally, the payoffs of the blue and red adversaries 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  $\pi_R > 0$ .

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

---

<sup>3</sup>This assumption is without loss of generality. What we need for the theory is a vector of preferences such that the decision maker has conflicting interests with one adversary for beliefs in one compact set and conflicting interests with the other adversary for beliefs in another compact set.

In this setting, adversaries have opposite incentives and compete to provide information. Remember that, given the decision maker's utility described in section 2.1, there is a state  $n^*$  such that  $v(b) > v(r)$  if  $n \geq n^*$  and  $v(b) < v(r)$  if  $n \leq n^* - 1$ . We normalize his payoffs in such a way that  $n^* = 0$ .<sup>4</sup> It is then immediate that the blue adversary will never collect evidence if  $n \geq 0$ , as the current belief implies the optimal action from his viewpoint. For identical reasons, the red adversary will never collect evidence if  $n \leq -1$  (from now on, we will say that the blue adversary is “ahead” if  $n \geq 0$  and “behind” if  $n \leq -1$ ). Define  $\lambda \equiv \frac{1-\theta}{\theta}$  ( $< 1$ ),  $F_B \equiv \frac{c_B(1+\lambda)}{1-\lambda}$  and  $F_R \equiv \frac{c_R(1+\lambda)}{1-\lambda}$ . Although technically non-trivial, it is possible to characterize analytically the optimal sampling strategies under competing adversaries. We focus on Markov equilibria where the state variable is  $n$ , the difference between the number of  $\rho$  and  $\beta$  signals.

**Proposition 1** *The red adversary samples if and only if  $n \in \{0, \dots, h^* - 1\}$  and the blue adversary samples if and only if  $n \in \{-l^* + 1, \dots, -1\}$ . The equilibrium cutoffs are  $h^* = \arg \max_h \Pi_n^r(l^*, h)$  and  $l^* = \arg \max_l \Pi_n^b(l, h^*)$ , where:*

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

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

*Adversaries sample more if their cost is lower and their benefit is higher. Also, the stopping thresholds of adversaries are strategic substitutes, so an adversary samples more if the cost of the other adversary is higher and the benefit of the other adversary is smaller.*<sup>5</sup>

Proof: see Appendix.

The idea is simple. Two adversaries with conflicting goals will never accumulate evidence simultaneously. Indeed, for any given belief, one of the adversaries will be ahead

<sup>4</sup>It could be that  $v(b) = v(r)$  for  $n = n^*$ . We assume that a strict inequality holds. This way, we do not need to impose an ad-hoc tie-breaking rule (this point is more important for the experiment than for the theory).

<sup>5</sup>These comparative statics are determined by taking derivatives in the profit functions  $\Pi_n^r(h, l)$  and  $\Pi_n^b(h, l)$  (see Appendix). 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  $\Pi_n^r(l, h)$  and for each integer  $h$  we find the integer  $l$  that maximizes  $\Pi_n^b(l, h)$  and use these values to find the equilibrium. Naturally, the same comparative statics hold.

and therefore will not have incentives to collect information as it can only hurt his interests. Suppose now that  $n \geq 0$ . The red adversary (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 adversary samples or action  $r$  is undertaken yielding a payoff  $\pi_R$ . 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 adversary is stopped. A symmetric reasoning when  $n \leq -1$  implies a lower cutoff  $-l^*$  where sampling by the blue adversary is stopped. Overall, when the event is very likely to be  $B$  the red adversary gives up sampling, and when the event is very likely to be  $R$  the blue adversary gives up sampling. For beliefs in between, the adversary currently behind acquires evidence while the other does not. The strategies are graphically illustrated in Figure 1.

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

### 3 Experimental design and procedures

We conducted 8 sessions of the two-adversaries game with a total of 78 subjects. Subjects were recruited by email solicitation. Sessions were conducted at The Social Science Experimental Laboratory (SSEL) at the California Institute of Technology. All interaction between subjects was computerized, using an extension of the open source software package ‘Multistage Games.’<sup>6</sup> No subject participated in more than one session. In each

---

<sup>6</sup>Documentation and instructions for downloading the software can be found at <http://multistage.ssel.caltech.edu>.

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. At the beginning of each match, each subject in a pair was randomly assigned a role as either red or blue (from now on, we call them ‘red adversary’ and ‘blue adversary’ respectively).<sup>7</sup> The event was represented to the subject as an urn, red or blue, drawn by the computer with equal probability. 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 which urn was selected by the computer. That is, the true event remained unknown to subjects.

Each adversary 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 adversaries,  $c_R$  and  $c_B$  respectively, was known but varied on a match-by-match basis as detailed below. If one or both adversaries drew a ball, then *both adversaries* observed the color(s) of the ball(s) drawn. The ball was then replaced in the urn.<sup>8</sup> If at least one adversary 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 payoff to each adversary which depended exclusively on the color of the balls drawn by both adversaries.<sup>9</sup> More precisely, if the difference between the number of blue and the number of red balls drawn was 0 or greater, then the blue adversary earned a high payoff and the red adversary earned a low payoff. From now on, we will say that the blue adversary “won” the match and the red adversary “lost” the match. If the difference was -1 or smaller, then the blue adversary lost the match and earned a low payoff, whereas the

---

<sup>7</sup>In the experiment, we used neutral terminology: participant in the ‘blue’ role, participant in the ‘red’ role, etc.

<sup>8</sup>Even though the decision of drawing a ball within a round was taken simultaneously, the balls were drawn with replacement. That is, adversaries always had 3 balls to draw from (this point was clearly spelled out in the instructions).

<sup>9</sup>As shown in Proposition 1, if the adversary unfavorably affected 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 adversary draws a ball in a given round shortens the duration of the experiment without, in principle, affecting the outcome.

red adversary won the match and earned a high payoff. From these earnings, adversaries 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 minimize (though not necessarily eliminate at all cost) the likelihood that an adversary earned a negative payoff in a given match once the costs were subtracted, because this could result in loss aversion effects. We therefore set the payoffs of winning and losing a match at 150 points and 50 points respectively, with the costs of sampling being 3 or 13 for each adversary.<sup>10</sup> Second, as in the theory section, roles were not symmetric. We gave an initial advantage to the blue adversary in order to implement a simple, deterministic and objective rule for the case  $n = 0$ . Finally, we computerized the role of the decision maker to make sure that sampling did not depend on (possibly incorrect) beliefs about the decision maker'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 adversary) 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.

---

<sup>10</sup>The exchange rate was 200 points = \$1.00. Notice that in the theoretical analysis, we allowed for the payoff of winning to be different for the blue and red adversaries ( $\pi_R \gtrless \pi_B$ ). By contrast, the payoff of losing was normalized to zero in all cases. Naturally, re-scaling payoffs has no consequences for the theory as long as the difference between winning and losing remains constant.

## 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 adversaries. 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 adversary stop sampling, respectively (see Figure 1). These equilibrium cutoffs are reported in Table 2.

The first cut at the data consists of comparing the empirical probabilities of sampling by the blue and red adversaries as a function of the state  $n$ , the difference between the number of blue draws and the number of red draws. Table 3 shows the empirical sampling frequencies and the 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  $n \in \{-4, \dots, 4\}$ , because there are few observations for choices in states outside this range (between 0 and 15 depending on the cost treatment).

Despite the data being rather coarse, it allows us to draw two main conclusions. First, adversaries understand the fundamentals of the game. Indeed, the theory predicts that both adversaries 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 adversaries had to simultaneously choose whether to sample, only in 4.8% of the cases the adversary ahead in the game did draw a ball. Furthermore, 3.2% correspond to a blue adversary drawing when  $n = 0$ . These small mistakes may be partly due to a misunderstanding of the tie-breaking rule, since the red adversary was significantly less likely to draw when  $n = -1$ . Furthermore and as we will see in section 4.5, these mistakes were greatly reduced over the course of the experiment. For the rest of the analysis and except otherwise noted, we will focus on the sampling strategy of the adversary behind in the game (red when  $n \geq 0$  and blue when  $n \leq -1$ ).

Second, sampling behavior is reasonably close to equilibrium predictions. Using Table 3, we can determine the number of instances where the adversary behind in the game played according to the predictions of theory. We separate the analysis in two groups. First, the aggregate data. These include all the observations of the adversary behind in

the game, separated into the cases where theory predicts draw and the cases where theory predicts no draw (the data is then pooled across roles). Second, the ‘marginal states.’ These include the observations in the last state where theory predicts that the adversary behind in the game should draw ( $h^* - 1$  for red and  $l^* + 1$  for blue), and the observations in the first state where theory predicts that the adversary behind in the game should not draw ( $h^*$  for red and  $l^*$  for blue). The data is compiled in Table 4.

The aggregate data reveals that the proportion of total observations consistent with the theoretical predictions is high, 85%, especially given that we only consider the choices of the adversary behind in the game. Also, there is a substantially lower proportion of under-sampling (in 8% of the observations there is no draw when there should be) than over-sampling (in 39% of the observations there is draw when there should not be). 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, adversaries almost invariably follow the theoretical prediction and sample (between 92% and 100% of the time). Perhaps more subtly, in equilibrium, an adversary can only under-sample once in each match, unless the other adversary 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.

The study of the marginal states allows us to investigate under- and over-sampling in more detail (i.e., when adversaries 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 4 shows, there are 8% fewer observations consistent with 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.<sup>11</sup> Therefore, although the fraction of mistakes is non-negligible (13% and 45%), behavior is still reasonably consistent with theory, especially for the ‘no draw’ case. Notice that the same remark applies as previously: for example, a red adversary who decides to stop in state  $x + 1$  when theory predicts that

---

<sup>11</sup>Although the differences are small in magnitude, they are all statistically significant given the large number of observations.

he should stop in  $x$  will, on average, produce several observations of over-sampling in  $x$  in every match.

Finally, we can also determine the optimal strategy of an adversary who knows the empirical sampling frequencies of the population. The problem turns out to be challenging because, contrary to the theoretical model, both adversaries 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 the adversaries in each role and in each cost treatment. In all eight cases, the best response coincides with the Markov equilibrium play described in Table 2. This result provides further support to the idea that adversaries' choices are close to the theoretical predictions. Indeed, if the strategies of an adversary were to depart systematically and substantially from equilibrium, the best responses of the other adversary would also be to depart from the Markov equilibrium. The results of this section are summarized as follows.

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

## 4.2 Aggregate payoffs

The next step consists in determining the expected payoffs of adversaries in the states where they 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 adversaries; (2) the expected payoffs if both adversaries played according to the Markov equilibrium; and (3) the expected payoff of an adversary who best responds to the empirical strategy of the other adversary which, given our previous result, coincides with the 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 5.<sup>12</sup>

Comparing (1) and (3), we notice that by deviating from the best response strategy, adversaries lose at most 3.9 points if their drawing cost is low, and at most 5.0 points

---

<sup>12</sup>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). We perform below what we think is a more informative comparison for the marginal states.

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 adversaries are not far from best responding to the strategy of their rivals. Comparing (1) and (2), we notice that the empirical choices of adversaries translate into net gains relative to the Markov 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 is constant across matches. Joint under-sampling is likely to result in lower costs and therefore higher average payoffs for both adversaries whereas joint over-sampling is likely to result in higher costs and therefore lower average payoffs (of course, this heuristic argument 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 adversary follows the empirical strategy and that the second adversary best responds to it (which, remember, also corresponds to the Markov equilibrium) at all states but  $n$ . We then determine the expected payoff in state  $n$  of the second adversary if he also plays the equilibrium strategy at  $n$  and if he plays the alternative strategy.<sup>13</sup> This exercise captures how much is lost by deviating from best response in one and only one state. The results are summarized in Table 6. We highlight in bold the payoffs given equilibrium play at all states. So, for example, since  $h^* = 3$  for the red adversary 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.

From this table, we can determine the utility loss of under-sampling and over-sampling in the marginal states, for each pair of costs and each role. We 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

---

<sup>13</sup>Notice that he may reach state  $n$  several times. The assumption is that he either always or never plays the equilibrium strategy.

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 partly due to the integer nature of the sampling strategies considered. Indeed, when the optimal stopping point is somewhere between  $x - 1$  and  $x$ , the adversary obtains a similar payoff when he stops at either of these thresholds. In that respect, using a discrete information accumulation process 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 6 with one difference. Instead of best responding, we assumed that the second adversary 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. We summarize the findings of this section as follows.

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

### 4.3 Deviations

We now explore in more detail the deviations from equilibrium behavior observed in the data. We start with an analysis of the adversaries' actions. From inspection of Table 3 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 7 we separate the marginal states into the last state where adversaries are supposed to draw and the first state where adversaries are supposed to not draw (just like in Table 4). We then report the proportion of sampling in each of these two cases.

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 adversaries draw with a substantially higher probability in  $h^* - 1$  and

$l^* + 1$  when their cost is high than when it is low. Also, in three out of four cases, their percentage decrease is also greater. This suggests that an adversary with low cost is more likely both to under-sample and to exhibit a less steep decline in drawing around the equilibrium than an adversary with high cost.

To further explore how costs affect deviations from equilibrium, we perform the same analysis as in Table 4, except that we separate the proportion of equilibrium play according to the adversary’s own cost. The results are displayed in Table 8.

When we pool together all states where the adversary is behind, the results are similar for low and high costs, simply because in non-marginal states adversaries generally play close to the 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 5, we notice that for the (13,13) treatment, the equilibrium payoffs exceed the empirical payoffs of adversaries in all four cases. By contrast, for the (3,3) treatment the empirical payoffs exceed the equilibrium payoffs of adversaries 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 adversaries and similar benefits whereas joint over-sampling in the (13,13) treatment results in higher costs for both adversaries and similar benefits.<sup>14</sup> The result is confirmed if we compare Markov equilibrium and best response to empirical behavior. When the cost of the red adversary is low, the blue adversary gets a higher payoff in (3) than in (2) whereas when the cost of the red adversary is high, the blue adversary gets a higher payoff in (2) than in (3). Since in both cases the blue adversary is choosing the same (optimal) strategy, this reinforces the idea that the red adversary 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 adversary when the blue adversary has cost 3 but not when the blue adversary has cost 13 (in that case, payoffs are almost identical in all four cases). However and as previously noted, payoff differences are generally small.

---

<sup>14</sup>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.

Finally, it is also instructive to compare the utility loss incurred by deviating from best response for adversaries with high and low cost of sampling. Using Table 6, we notice that in 3 out of 4 observations, the utility loss for the low cost adversary is bigger with under-sampling than with over-sampling. Conversely, in 3 out of 4 observations, the utility loss for the high cost adversary 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 an adversary with a high cost than for an adversary 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 can be 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 adversaries with low cost and over-sampling by adversaries with high cost. In general, over-sampling is more pronounced than under-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 an adversary as a function of the state. We only include states where the adversary is behind to ensure a monotonic theoretical relation.<sup>15</sup> For each role, we perform the regression on four subsamples, taking either the adversary’s own cost or the other adversary’s cost as fixed. In the former case, we introduce a dummy variable that codes whether the other adversary’s cost was high (high other  $c$ ). In the latter case, we introduce a dummy variable that codes whether the adversary’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

---

<sup>15</sup>Also, we already know from the previous analysis that behavior is almost invariably in accordance with theory when the adversary is ahead.

same cost pairs. We study a simple version of experience effects by introducing a dummy variable that separates the first 5 matches from the last 5 matches within a given cost pair (exp). We also include interactions terms. The results are summarized in Table 9.

Not surprisingly, as the difference between unfavorable and favorable draws increases, adversaries are less inclined to sample. The effect is strong and highly significant in all eight subsamples. Similarly, as an adversary’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 on the behavior of an adversary of the other adversary’s cost is more involved. Proposition 1 states that thresholds are strategic substitutes, so a higher cost by one adversary 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 adversary should translate into a higher level of sampling by the blue adversary 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 other  $c$ ’ for the blue adversary being positive in both cases but significant only when  $c_B = 3$ . For the red adversary, the integer constraint implies no increase in sampling when the blue adversary’s cost increases both when  $c_R = 3$  and when  $c_R = 13$ . In the data, the coefficient is significant when the cost of the red adversary is high. Overall, all four coefficients for ‘high other  $c$ ’ 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 explore 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 10. The table can be read as follows. For each state  $n$ , we consider only the adversary behind in the game. 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 [+].

For each state, we then compare the empirical and theoretical change in sampling between cost pairs. Note that theory predicts either 0% or 100% probability of sampling in each state (so no change at all or a 100% change between the row and column treatments).

We code a (positive or negative) empirical change in probability as ‘significant’ when (i) the magnitude of the (positive or negative) change is at least 10%, and (ii) the change is statistically significant at the 5% level.<sup>16</sup> Using this criteria, we obtain that 23 out of 24 mean comparisons for the red adversary follow the patterns predicted by theory: no difference in 15 cases, and a statistically significant decrease in 8 cases. For the blue adversary, 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 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.<sup>17</sup> Notice that our method controls neither for joint correlation between tests (when one sampling departs significantly from theory, several comparisons are affected) nor for multiplicity of tests (we make 48 comparisons at a 5% significance level). However, 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 theory both in aggregate and state-by-state: an adversary samples more when his cost is low and when the cost of the other adversary is high.*

## 4.5 Learning

We now study whether subjects change their behavior over the course of the experiment. We know from section 4.1 that the proportion of mistakes by adversaries 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.

We then move on to the more interesting case of adversaries 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, that is, matches 1 to 20) and late sequences (3 and 4, that is,

---

<sup>16</sup>In other words, a decrease in sampling from 1.00 to 0.97 (as, for example, between (3, 3) and (13, 13) for  $n = 0$ ) is *not* coded as a change even if the 3% difference is statistically significant.

<sup>17</sup>Similar results are obtained if we used different criteria for the magnitude of the change (e.g., at least 20% change) and/or the statistical significance (e.g., 1% level).

matches 21 to 40) or into inexperienced (first 5 matches within a cost pair) and experienced (last 5 matches). We then determine the proportion of equilibrium play in each subsample. The results are compiled in Tables 11 and 12.

From Table 11, 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, partly because it is quite low to start with. In all four cases, mistakes are reduced. By contrast, Table 12 suggests that experience within a cost treatment has virtually no effect on the behavior of adversaries.

A more rigorous look at the data consists in studying significance of the ‘sequence’ and ‘experience’ variables in the Probit regression presented in Table 9. The sequencing effect is significant for both adversaries 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, adversaries sample more if they are behind by few draws and less if they are behind by many draws, as learning would predict.<sup>18</sup> 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 regression provides limited evidence of learning due to sequencing and none due to experience.

All in all, there is little evidence of changes in sampling behavior over trials. One possible explanation is that subjects had insufficient exposure to the game (40 matches under 4 different cost treatments). We tend to favor a simpler explanation: subjects play relatively close to equilibrium right from the outset, so there is little room for learning. The result is summarized as follows.

**Result 5** *Adversaries ahead in the game learn to avoid sampling mistakes almost entirely. Adversaries 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 adversaries with opposite interests. We have characterized the Markov equilibrium of the game and shown

---

<sup>18</sup>The p-value of ‘seq. late’ for the blue player with high cost is 0.054. The other three are below 5%.

that the choice variables are strategic substitutes: if the incentives to collect information of one adversary increase, then the incentives of the other adversary decrease. We have tested the theory in a controlled environment. As in the individual decision making literature, behavior of subjects is remarkably close to predictions by theory even if in our game the strategies are substantially more complex. Mistakes are relatively infrequent and, contrary to the individual decision making literature, take more often the form of over-sampling than under-sampling. Comparative statics on the adversary's own cost and the other adversary's cost generally follow the predictions of 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. From 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 non-verifiable information. 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 even more sophisticated environments. One possibility would be to consider three adversaries. When the evidence favors one adversary, 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 let adversaries choose the accuracy of information, that is, the correlation between event and signal. A different extension would be to allow adversaries to engage in agreements with collusive side transfers that would 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, will the payoffs of each adversary be above or below their expected return in the non-cooperative Markov equilibrium with sampling? A final option would be to use this framework to study bribery, for example by letting the decision maker play an active role and demand bribes from the adversaries in exchange of a certain action. Will he be able to extract the full surplus of the adversaries? These and other related questions are left for future research.

## Appendix: proof of Proposition 1

It is immediate that the blue adversary will never sample if  $n \geq 0$  and the red adversary will never sample if  $n \leq -1$ . Also, if at some stage no adversary 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 is  $n \in \{0, \dots, h - 1\}$ , where  $h$  is the value where the red adversary gives up sampling (we will determine this optimal value below). The value function of the red adversary, 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  $\theta(1 - \theta)$  is the probability of receiving signal  $\beta(\rho)$  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 adversary 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  $q_B^b$  is the probability that the blue adversary reaches  $n = -l$  before reaching  $n = 0$  given event  $S \in \{R, B\}$  and state  $n = -1$ . In other words, the red adversary knows that when  $n = -1$ , the blue adversary will restart sampling (thus the red adversary will stop paying costs). With probability  $q_B^b$ , the belief will reach  $n = -l$ . The blue adversary will stop at that point and the red adversary will obtain the payoff  $\pi_R$ . With probability  $1 - q_B^b$ , the belief will go back to  $n = 0$ . The value function of the red adversary 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:

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

When the event is  $S = R$ , the second-order difference equation for the red adversary 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  $1 - \theta$  ( $\theta$ ). Solving in an analogous fashion, we get:

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

At this point, we need to determine  $q_S^b$ . Recall that the blue adversary gives up at  $n = -l$  (where  $-l$  will be determined below). Let  $h_S^b(n)$  denote the blue adversary'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:

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

and

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

Note that  $h_S^b(\cdot)$  captures exclusively the blue adversary'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 adversary's calculation only the probabilities matter (not the net utility of the blue adversary). 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 adversary given state  $n \in \{0, \dots, h - 1\}$ , is then:

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

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

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

In a Markov equilibrium, the best response functions of the red and blue adversaries 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:

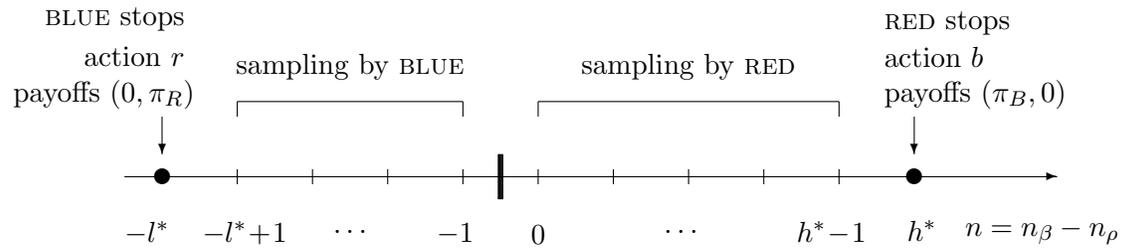
$$-\lambda^{h^*(l)} \ln \lambda \left[ \pi_R(1+\lambda^l) - F_R(h^*(l)+1)(1-\lambda^l) \right] = F_R(1-\lambda^{h^*(l)})(1-\lambda^{l+h^*(l)})$$

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

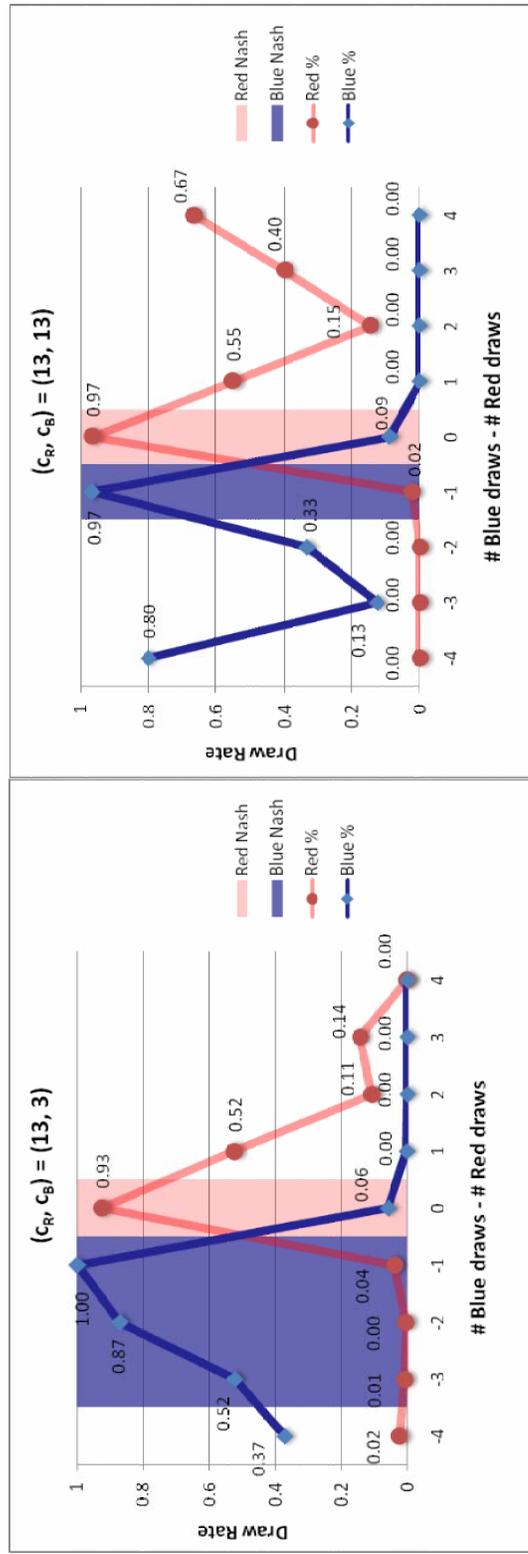
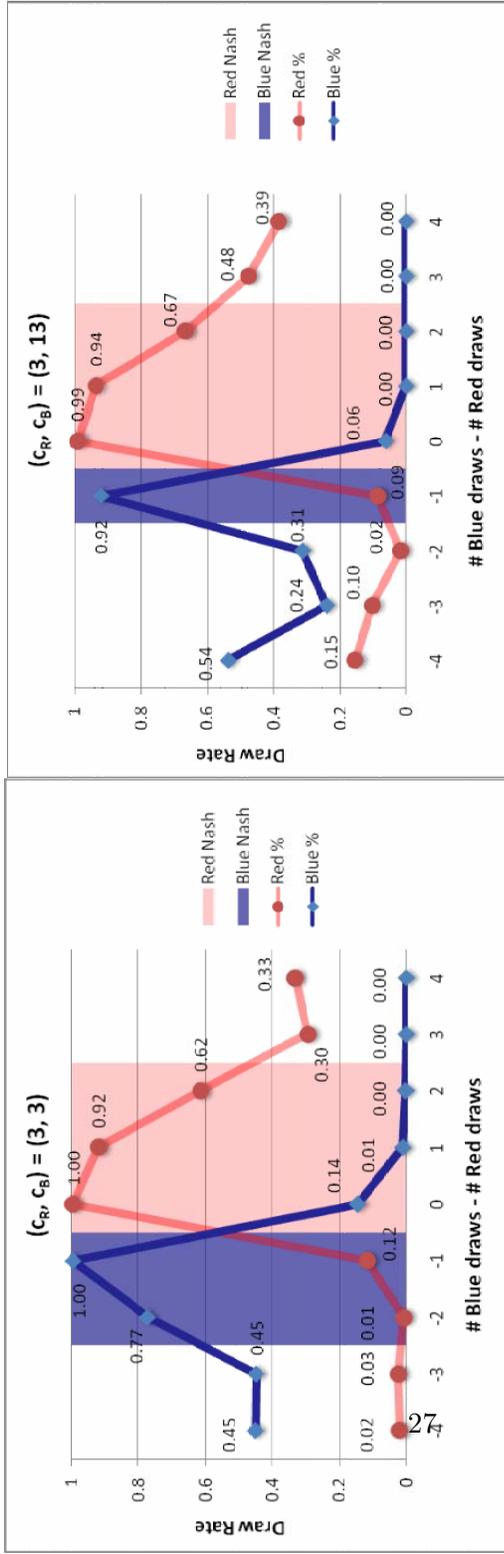
As expected,  $h^*$  and  $l^*$  do not depend on  $n$ , that is, the optimal stopping rules of the two adversaries *are not* revised with the realizations of the sampling process. Notice that  $\frac{\partial^2 \Pi_n^r}{\partial h^2} \Big|_{h^*} < 0$  and  $\frac{\partial^2 \Pi_n^b}{\partial l^2} \Big|_{l^*} < 0$ , 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  $\frac{\partial l^*}{\partial h} \propto \frac{\partial^2 \Pi_n^b}{\partial h \partial l} \Big|_{l^*} < 0$ .  $\square$

## 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. Kamenica, E. and M. Gentzkow (2009) "Bayesian Persuasion." *mimeo*, University of Chicago.
7. Matthews, S. and A. Postlewaite (1985) "Quality Testing and Disclosure." *RAND Journal of Economics*, Vol.16, pp.328-340.
8. Milgrom, P. and J. Roberts (1986) "Relying on the Information of Interested Parties." *RAND Journal of Economics*, Vol.17, pp.18-32.
9. Schotter, A. and Y.M. Braunstein (1981) "Economic Search: an Experimental Study." *Economic Inquiry*, Vol.19, pp.1-25.



**Figure 1.** Sampling strategies by BLUE and RED adversaries



**Figure 2.** Sampling frequencies by state and cost treatment

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: Markov equilibrium.

$n$ (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 – theory]	.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 – theory]	.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 – theory]	.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 – theory]	.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 – theory]	.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 – theory]	.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 – theory]	.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 – theory]	.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 3: Sampling frequencies.

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

Table 4: Proportion of equilibrium behavior when adversary is behind (standard errors clustered at the individual level in parentheses; number of observations in brackets).

state ( $c_R, c_B$ )	$n = -1$				$n = 0$			
	(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) Markov eq.	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) Markov eq.	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 5: Expected payoffs of BLUE and RED adversaries at  $n = -1$  and  $n = 0$ .

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

Table 6: Values to drawing and not drawing by state.

adversary $(c_R, c_B)$	RED				BLUE			
	(3, 3)	(3, 13)	(13, 3)	(13, 13)	(3, 3)	(3, 13)	(13, 3)	(13, 13)
<hr/> <hr/>								
marginal state								
Theory is DRAW	.62	.67	.93	.97	.77	.92	.52	.97
Theory is NO DRAW	.30	.48	.52	.55	.45	.31	.37	.33

Table 7: Empirical probabilities of drawing in marginal states.

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

Table 8: Proportion of equilibrium behavior by adversaries' own cost (standard errors clustered at the individual level in parentheses; number of observations in brackets).

	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** (.269)	1.58** (.389)	2.30** (.422)	3.00** (.347)	3.11** (.218)	1.78** (.340)	3.50** (.294)	3.10** (.344)
draws behind	-.710** (.144)	-.642* (.282)	-.641** (.216)	-.822** (.192)	-.990** (.090)	-.713** (.192)	-1.02** (.126)	-.779** (.148)
seq. late	.241 (.216)	.889 (.460)	.199 (.371)	.579 (.366)	-.261 (.335)	1.79** (.355)	-.163 (.345)	.251 (.272)
draw $\times$ seq.	-.121 (.108)	-.655* (.316)	-.140 (.237)	-.381 (.219)	.134 (.147)	-1.19** (.209)	.097 (.173)	-.179 (.142)
exp.	.011 (.219)	.541 (.410)	.388 (.296)	-.179 (.354)	.045 (.230)	.594* (.300)	.184 (.255)	.141 (.243)
draw $\times$ exp.	.041 (.114)	-.331 (.291)	-.187 (.186)	.125 (.202)	.006 (.098)	-.346 (.182)	-.096 (.131)	-.028 (.136)
high own $c$	—	—	-.836** (.165)	-1.17** (.164)	—	—	-1.26** (.118)	-1.08** (.157)
high other $c$	.244* (.110)	.022 (.100)	—	—	.150 (.096)	.248* (.110)	—	—
adj. $R^2$	0.28	0.27	0.25	0.32	0.35	0.36	0.35	0.33

Table 9: Probit regression on probability of sampling (standard errors clustered at individual level in parentheses; \* = significant at 5% level, \*\* = significant at 1% level).



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

Table 11: Proportion of equilibrium behavior by sequence (standard errors clustered at individual level in parentheses; number of observations in brackets).

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

Table 12: Proportion of equilibrium behavior by level of experience (standard errors clustered at individual level in parentheses; number of observations in brackets).