

Information Gatekeepers: Theory and Experimental Evidence

Isabelle Brocas* Juan D. Carrillo† Thomas R. Palfrey‡

Abstract

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 a higher rate of 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.

Keywords: experimental design, search, information acquisition, adversarial system.

JEL Classification: C91, D83.

*Department of Economics, University of Southern California, 3620 S. Vermont Ave., Los Angeles, CA 90089 and CEPR.

†Department of Economics, University of Southern California, 3620 S. Vermont Ave., Los Angeles, CA 90089 and CEPR.

‡(Corresponding author) Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, email: <trp@hss.caltech.edu>, Phone: 626-395-4088, Fax: 626-432-1726.

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. Finally, firms reveal through advertising the characteristics of their products. This information affects the fit between the product and the preferences of consumers. In this paper, we build a theoretical model to understand the incentives of individuals to collect information in strategic contexts. We then test in a controlled laboratory setting whether subjects play according to the predictions of the theory.

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.

The main reason to assume public information is simplicity. Indeed, 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. Also, for some applications such as product advertising, one can argue that firms and consumers learn concurrently the *fit* between the characteristics of products revealed through advertising and the tastes of consumers.

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 show that the actions of adversaries are strategic substitutes: when the cost of news acquisition for an adversary increases, that adversary has fewer incentives to collect evidence, which in turn implies that the other adversary has more incentives to collect evidence.

We then report an experiment that analyzes behavior in this information acquisition game. 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 adversaries take the decision predicted by theory 85% of the time (92% of the time when the theory prescribes sampling and 61% of the time when the theory prescribes no sampling). Furthermore, the best response to the empirical strategy of the other adversary is to play the equilibrium strategy, 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 are more instances of 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. First, the amount of sampling is decreasing in the adversary’s cost of information acquisition. More interestingly, sampling by one adversary is (weakly) increasing in the other adversary’s cost. This means that subjects do not consider this game as an individual decision-making problem; they realize the strategic substitutability of actions and play accordingly. These comparative statics hold 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) is to our knowledge the first study

that analyzes how an individual can affect the choices of others by selectively deciding whether to acquire or avoid public information. The paper however focuses on a very simple one agent model with free information. It thus ignores the strategic component of optimal sampling and the cost-benefit trade-off.¹ In an independent and concurrent research, Gul and Pesendorfer (2009) also extend the setting of Brocas and Carrillo (2007) to include two agents with competing interests and a positive cost of sampling. As in our case, the authors show the optimality of a cutoff strategy and the strategic substitutability of the sampling cutoffs. The model is specified using a more general and elegant continuous time Brownian motion process with unknown drift. This formalization, however, is substantially more difficult both to implement experimentally (choices must necessarily be revised at discrete intervals of finite length) and to explain to subjects than ours.² Gentzkow and Kamenica (2011) use their previously developed methodology to analyze a still more general information gathering and transmission problem by multiple agents with arbitrary states, signals and preferences.³ Our paper shares also some similarities with the literature on technological races with strategic interactions and uncertainty (Harris and Vickers 1987; Horner 2004; Konrad and Kovenock 2009). Finally, there is also an older literature on games of persuasion (Matthews and Postlewaite 1985; Milgrom and Roberts 1986) 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 setting (see e.g. Schotter and Braunstein (1981) in a labor market context, Banks et al. (1997) in a two-arm bandit problem and the surveys by Camerer (1995) and Cox and Oaxaca (2008)). A main finding in this literature 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 to account for *search in strategic contexts*. The strategic, multi-person nature of adversarial

¹Kamenica and Gentzkow (2010) approach 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.

²Gul and Pesendorfer (2009) are able to show uniqueness of the equilibrium under the assumption that only the subject behind in the game can acquire information. The paper is more thorough in that it also studies the case of private acquisition of non-verifiable information, an issue that we ignore both in our theory and our experimental setting.

³Arbitrary preferences is especially interesting because it allows the authors to compare cooperative and non-cooperative outcomes.

search substantially increases the complexity of the decision making problem relative to the individual decision making counterpart. In particular, subjects are expected to modify their stopping rule in response to a change in the rival’s payoff. Surprisingly, we still observe a behavior that is close to the theory, and the comparative statics with respect to the opponent’s cost also follow the predictions of the theory. The main difference is that, with strategic sampling, there is an increase in the rate of excessive experimentation, balancing the frequency of under-sampling and over-sampling.⁴ Technological races have also been studied experimentally. Zizzo (2002) tests the model of Harris and Vickers (1987) in the laboratory and finds substantial departures from the theoretical predictions. By contrast, Breitmoser et al. (2010) argue that a Quantal Response extension of Markov Perfect Equilibrium explains rather well the behavior of players in the infinite time horizon model of Horner (2004). Last, the difficulty of individuals to perform Bayesian updating has been long noted both in psychology and in economics. In individual decision contexts, Kahneman, Slovic and Tversky (1982) emphasize the mistakes in probabilistic assessments due to insufficient sensitivity to priors, sample size and accuracy of information among other factors. Charness and Levin (2005) show that subjects are less likely to follow the Bayesian updating strategy when it is consistent with a non-intuitive “switch when you win - stay when you lose” heuristic than when it is consistent with a more natural “switch when you lose - stay when you win” heuristic. Two reasons can explain why the empirical behavior of subjects is close to the theory in our experiment (and in the search problems discussed above) but not in other choice under uncertainty contexts. First, subjects may be employing a simple heuristic which, for our particular setting, *coincides* with the optimal choice. Second, search is ubiquitous in our everyday lives, so individuals have developed intuitive but accurate ways of solving this class of problems. We favor the second explanation, especially because whichever method they are using applies both to individual situations and strategic games.

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, compar-

⁴Another difference, which is probably of second-order importance, is that our paper deals with search for information rather than search for payoffs. According to Camerer (1995, p.673), this could lead to different conclusions even though both are formally similar. It does not seem to be the case in our setting.

ative statics (aggregate and state-by-state), and learning. In section 5, we provide some concluding remarks. The proof of the proposition is relegated to the appendix.

2 The model

2.1 The game

Consider a game with three agents. One agent is a *decision maker* (congress, judge, consumer) who must undertake an action that affects the payoff of all three agents. The other two agents 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, that is, agents play a game of imperfect but symmetric information. Therefore, at any point in time, decision maker and adversaries share the same belief about which event was 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, $a \in \{b, r\}$. His payoff depends on the action he takes and the event realized. Formally, his expected payoff is:

$$v(a) \equiv \sum_S \Pr(S)v(a|S)$$

To preserve symmetry, we assume that the common prior belief is $\Pr(S) = 1/2$. At each stage, each adversary simultaneously decides whether to pay a (strictly positive) cost 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_β , the number of β -signals, and n_ρ , the number of ρ -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 the 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.⁵ From now on, we call them the blue adversary and the red adversary, respectively. Without loss of generality, we normalize the payoffs of the blue and red adversaries to be 1 and 0 when their most preferred and least preferred action is taken by the decision maker.

Adversaries can acquire as many signals $s \in \{\beta, \rho\}$ as they wish. Asymmetries in the payoffs of adversaries are captured via the cost of a signal. Formally, the cost of each signal is c_B for the blue adversary and c_R for the red adversary, with $c_B \geq c_R$. 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

⁵This assumption is excessively restrictive. 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.

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.

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 $n^* = 0$.⁶ It is then immediate that the blue adversary will never collect information if $n \geq 0$, as evidence is costly and the current belief already implies the optimal action from his viewpoint. For identical reasons, the red adversary will never collect information 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, we can 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 ρ and β 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:*

$$\begin{aligned} \Pi_n^r(l, h) &= \frac{1}{1 + \lambda^n} \left[\left(1 + \lambda^l - F_R(h + 1)(1 - \lambda^l) \right) \left[\frac{\lambda^n - \lambda^h}{1 - \lambda^{h+l}} \right] - F_R(h - n)(1 - \lambda^n) \right], \\ \Pi_n^b(l, h) &= \frac{1}{1 + \lambda^n} \left[\left(1 + \lambda^h - F_B(l - 1)(1 - \lambda^h) \right) \left[\frac{1 - \lambda^{n+l}}{1 - \lambda^{h+l}} \right] + F_B(n + l)(1 - \lambda^n) \right]. \end{aligned}$$

*Adversaries sample more if their cost is lower. Also, the stopping thresholds are strategic substitutes, so adversaries sample more if the cost of their rival is higher.*⁷

Proof: see Appendix.

⁶It 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).

⁷These 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.

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 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 of 1. 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 with respect to the adversary's own payoffs are simple: a lower cost implies a higher incentive to sample. 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 means that if the cost of one adversary decreases, then the other adversary will engage in less sampling. A main contribution of the experimental study will be to test empirically this strategic substitutability of thresholds.

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.’⁸ 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. 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).⁹ 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.¹⁰ 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.¹¹ 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

⁸Documentation and instructions for downloading the software can be found at <http://multistage.ssel.caltech.edu>.

⁹In the experiment, we used neutral terminology: participant in the ‘blue’ role, participant in the ‘red’ role, etc.

¹⁰Even 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).

¹¹As shown in Proposition 1, if the adversary 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 adversary draws a ball in a given round shortens the duration of the experiment without, in principle, affecting the outcome.

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 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.¹² 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.¹³ 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.

¹²The exchange rate was 200 points = \$1.00. Notice that in the theoretical analysis the payoff of losing was normalized to zero. Re-scaling payoffs has no consequences for the theory.

¹³A sample copy of the instructions can be found in the online supplementary material.

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 adversaries 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 2.¹⁴

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, two-thirds of these mistakes 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.

¹⁴Although the empirical state space is $n \in \{-6, \dots, 7\}$, in Table 3, Figure 2 and in the mean comparison in Table 10, we restrict the analysis to $n \in \{-4, \dots, 4\}$, because there are few observations (between 0 and 15) for choices in states outside this range. All the other tables and statistical analyses are based on the entire data set.

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 lower frequency of under-sampling than over-sampling: in 8% of the cases where subjects should draw they choose instead not to draw whereas in 39% of the cases where subjects should not draw they choose instead to draw. Note, however, two caveats when we attempt to compare these two types of mistakes. First, 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. In fact, due to the stochastic nature of the process, a red adversary who chooses a cutoff strategy with an incorrect stopping state (e.g., $h^* + 1$ instead of h^*) will, on average, over-sample more than once in each match. Second, the total number of observations is asymmetric. For instance, suppose that an adversary under-samples in match 1 and over-samples in match 2. For both matches there are observations where the adversary should sample (and for match 2 he always samples when he should). Conversely, only for match 2 there are observations where the adversary should not sample (and at least in one of these cases he mistakenly samples). These considerations suggest that, in the absence of a behavioral theory about under- and over-sampling, we cannot make a direct comparison of the proportions of mistakes in either direction.

One way to make over- and under-sampling more comparable (although the caveats will still apply) is to restrict attention to the marginal states, that is, the states where 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 small but statistically significant in both cases: around 5% for under-sampling and 6% for over-sampling. Therefore, although the fraction of mistakes is non-negligible, behavior is still reasonably consistent with theory, especially for the ‘no draw’ case. Naturally, the increase in mistakes when we consider only the marginal states is more salient for under-sampling if we use as a baseline the mistakes in all states (increase from 8% to 13% v. increase from 39% to 45%). This is in part explained by the second caveat mentioned above: the number of observations where theory predicts draw and no draw respectively is less dissimilar when we consider only the marginal states (2404 v. 1156) than when we consider all states behind (6130 v. 1819).

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 imply a departure 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 5.¹⁵

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 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.

The previous comparisons are suggestive but incomplete. Indeed, one concern in this type of games is that small payoff differences between predicted and empirical choices may be due to a flatness in the payoff functions. 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.¹⁶ 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

¹⁵For 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.

¹⁶Notice that he may reach state n several times. The assumption is that he either always or never plays the equilibrium strategy.

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 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 (17.1 vs. 1.9 points).¹⁷ Overall, the exercise suggests that payoff functions are not flat; the loss incurred by a mistake in only one state is sometimes small but some other times quite high (17 points out of a total stake of 100 points minus the cost of sampling).¹⁸ 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

¹⁷This is partly due to the integer nature of the sampling strategies. 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 affect the results.

¹⁸We also performed the same computations as in Table 6 except that, 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 and are not reported for the sake of brevity.

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 footnote 17). 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.¹⁹ 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

¹⁹The 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.

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.

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 incurred with each draw (13 and 3), 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.²⁰ For each role, we perform

²⁰Also, we already know from the previous analysis that behavior is almost invariably in accordance with theory when the adversary is ahead.

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 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 (weakly) 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 strategic substitutability is, if anything, stronger than predicted by the theory. 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, which 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. We perform a standard t-test of the difference and report in parentheses the p-value for the statistical significance of the average 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.²¹ 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. 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 to a large extent in accordance with theory.²² 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.*

²¹In 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.

²²We should mention as a caveat that we do not use the clustered standard errors when performing the t-test, which is somewhat unsatisfactory since the observations are not independent. Note, however, that similar results are obtained even if we strengthen the statistical significance (e.g., 1% level). Results are also similar if we use a different criterion for the magnitude of the change (e.g., at least 20% change).

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, 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.²³ 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

²³The p-value of ‘seq. late’ for the blue player with high cost is 0.054. The other three are below 5%.

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 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 predictive power of the theory in a controlled laboratory setting. Behavior of subjects is remarkably close to predictions by theory even if, relative to individual decision making problems, choices in our game are substantially more complex and involve strategic considerations. Mistakes are relatively infrequent and, to some extent, 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. First, one could consider richer signal structures. For example, having a third “null” signal that contains no information is equivalent to having a higher sampling cost. One could also have a larger number of signals or even a continuum of signals. These models would be more complicated to solve analytically, but we conjecture that the solutions would be stopping rules with similar characteristics and comparative statics as in our binary signal model.

A second possible extension of the theory would be to combine the acquisition of information and the revelation of information paradigms as in the model of Gul and Pesendorfer (2009). 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 possibility 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: $g_B^r(-1) = q_B^b + (1 - q_B^b)g_B^r(0)$, where q_S^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 a payoff of 1. 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:

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

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 $1 - \theta$ (θ). Solving in an analogous fashion, we get:

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

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:

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

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:

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

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), the best response functions satisfy:

$$-\lambda^{h^*(l)} \ln \lambda \left[1 + \lambda^l - F_R(h^*(l) + 1)(1 - \lambda^l) \right] = F_R(1 - \lambda^{h^*(l)})(1 - \lambda^{l+h^*(l)}) \quad (6)$$

$$-\lambda^{l^*(h)} \ln \lambda \left[1 + \lambda^h - F_B(l^*(h) - 1)(1 - \lambda^h) \right] = F_B(1 - \lambda^{l^*(h)})(1 - \lambda^{l^*(h)+h}) \quad (7)$$

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. Also $\left. \frac{\partial^2 \Pi_n^r}{\partial h^2} \right|_{h^*(l)} < 0$ and $\left. \frac{\partial^2 \Pi_n^b}{\partial l^2} \right|_{l^*(h)} < 0$, so h^* and l^* are indeed maxima.

From (4) there exists $\bar{h}(l)$ such that $\Pi_n^r(l, h) < 0$ for all l and $h > \bar{h}(l)$. Similarly, from (5) there exists $\bar{l}(h)$ such that $\Pi_n^b(l, h) < 0$ for all h and $l > \bar{l}(h)$. It means that $h^*(l) < +\infty$ and $l^*(h) < +\infty$ (i.e., cutoffs are finite for all $c_R > 0$ and $c_B > 0$). This together with the continuity of the best response functions is sufficient to ensure that an equilibrium always exists. Note however that the first-order conditions (6) and (7) can have unique, multiple or no interior solution (in the last case, only a corner solution will exist with either one or both adversaries never sampling).

Also, $\left. \frac{\partial h^*}{\partial l} \right|_{h^*(l)} \propto \left. \frac{\partial^2 \Pi_n^r}{\partial h \partial l} \right|_{h^*(l)} < 0$ and $\left. \frac{\partial l^*}{\partial h} \right|_{l^*(h)} \propto \left. \frac{\partial^2 \Pi_n^b}{\partial h \partial l} \right|_{l^*(h)} < 0$, which means that h^* and l^* are strategic substitutes. Finally, $\left. \frac{\partial^2 \Pi_n^r}{\partial h \partial c_R} \right|_{h^*(l)} < 0$, so the reaction function $h^*(l)$ shifts downwards when c_R increases. Together with the strategic substitutability, it means that in any stable equilibrium h^* is non-increasing in c_R and l^* is non-decreasing in c_R . Similarly, $\left. \frac{\partial^2 \Pi_n^b}{\partial l \partial c_B} \right|_{l^*(h)} < 0$, so the reaction function $l^*(h)$ shifts downwards when c_B increases, which again means that in any stable equilibrium, l^* is non-increasing in c_B and h^* is non-decreasing in c_B . \square

Electronic supplementary material: sample copy of instructions

Thank you for agreeing to participate in this research experiment on group decision making. During the experiment we require your complete, undistracted attention. So we ask that you follow these instructions carefully. You may not use the computer except as specifically instructed. Do not chat with other students, or engage in activities such as using your cell phones or head phones, reading books, etc.

For your participation, you will be paid in cash, at the end of the experiment. Different participants may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. So it is important that you listen carefully, and fully understand the instructions before we begin. You will be asked some review questions after the instructions, which have to be answered correctly before we can begin the paid session.

The entire experiment will take place through computer terminals, and all interaction between you will take place through the computers. It is important that you not talk or in any way try to communicate with other participants during the experiment except according to the rules described in the instructions.

We will start with a brief instruction period in which you will be given a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, raise your hand and your question will be answered out loud so everyone can hear. If any questions arise after the experiment has begun, raise your hand, and an experimenter will come and assist you privately.

You will make choices over a sequence of 40 matches. In each match, you will be randomly paired with another participant. You will make a series of decisions and receive a payoff that will depend on your decisions in that match and on the decisions of the participant you are paired with. We will explain exactly how these payoffs are computed in a minute.

At the end of the session, you will be paid the sum of what you have earned in all matches, plus the show-up fee of \$10. Everyone will be paid in private and you are under no obligation to tell others how much you earned. Your earnings during the experiment are denominated in points. At the end of the experiment you will be paid \$1.00 for every 200 points you have earned.

The experiment has 4 parts. Here is how each match in part 1 goes. First, the computer randomly matches you into pairs. Since there are 10 participants in today's session, this will result in 5 matched pairs. You are not told the identity of the participant you are matched with. Your payoff depends only on your decision and the decision of the one participant you are matched with. What happens in the other pairs has no effect on your payoff and vice versa. Your decisions are not revealed to participants in the other pairs.

Next, the computer randomly assigns a role to each member of a pair: one will have the "red" role and the other will have the "blue" role. The computer will show you two Boxes. Each Box contains red and blue balls. Box 1 contains 2 red balls and 1 blue ball while Box 2 contains 2 blue balls and 1 red ball. The computer randomly chooses one of these boxes for your pair. In each match, there is a 50% chance the computer selects Box 1 and a 50% chance the computer selects Box 2. This random selection of boxes is done separately and independently for each pair and each match. Neither you nor the participant you are paired with is told which Box has been selected for your pair until the end of the match. The computer next shuffles the balls in the selected Box and hides the colors of the balls on your computer screen so you cannot tell which Box it is.

You will then have an opportunity to pay a cost and observe the color of one ball in the box. Or, you may choose to not pay the cost and not observe a ball. To draw a ball, you simply click on one of the balls in the Box and its color will be revealed. The other participant is given the same opportunity. The Box number (1 or 2) is the same for them, but the computer shuffles the balls in the box on their screen independently of how it shuffled the balls on your screen. That is to say, all of these draws are done WITH replacement. If you both decide to draw a ball, then it is as if you drew one ball, then put it back in the box, the computer reshuffled the contents of the box for the other participant, who then drew their ball, and put it back in the box. The color of all balls that are drawn are shown to BOTH participants, but

you only pay a cost for the balls that YOU draw. The cost to draw a ball for the participant in the red role is 3 points. The cost to draw a ball for the participant in the blue role is 3 points.

Suppose you and/or the other participant decided to draw a ball. Once both participants have observed the color of the ball (or balls), we move to a second round of ball drawing. The Box stays the same, but the balls in the box are reshuffled, again independently for you and the other participant. You and the other participant will each then have another opportunity to pay a cost to draw another ball. Remember, all these draws are independent and with replacement. This means that no matter how many balls you and the other participant draw, you will never know for sure which Box it is. For example, suppose you had chosen to draw a ball in both round 1 and round 2. If both balls you drew were red, you might be tempted to think that the Box is for sure Box 1 since Box 2 has only one red ball, but this is logically incorrect. It could have been Box 2 and you just happened to pick the same ball in both rounds. The same logic applies if both you and the other participant draw two balls of the same color (in the same round or in different rounds).

This ball drawing process continues round after round until there is a round when neither of you decides to draw a ball. If you do not draw a ball but the other participant does, then your pair moves to the next round and BOTH of you have the opportunity to keep drawing balls. However, if there is a round when neither of you draws a ball, the match ends. Remember, balls are always drawn with replacement so every ball you or the other participant draws is placed back in the Box and the contents are reshuffled. Therefore, there are always 3 balls in the Box to draw from (2 red and 1 blue if the computer chose Box 1 or 1 red and 2 blue if the computer chose Box 2).

[SCREEN 1] - This slide summarizes the rules.

Your earnings are affected by the final outcome of how many red and blue balls were drawn in total by both you and the other participant. These outcome earnings are as follows:

- If more red balls were drawn than blue balls, then the outcome is red. The participant in the red role receives 150 points and the participant in the blue role receives 50 points.
- If more blue balls were drawn than red balls, then the outcome is blue. The participant in the red role receives 50 points and the participant in the blue role receives 150 points.
- Finally, if the same number of red and blue balls were drawn, then the outcome is also blue. The participant in the red role receives 50 points and the participant in the blue role receives 150 points. This includes the case where no balls were drawn by any participant.

From these outcome earnings, we subtract your ball drawing costs, equal to the number of balls you drew during the match times your cost of drawing each ball. Remember you only pay a cost for the balls you draw, not for the balls the other participant draws. Your payoff for the match is then your outcome earnings minus your ball drawing costs.

[SCREEN 2] - This slide shows a summary of how your Total earnings are computed in a match [explain the slide]

Note that payoffs depend only on the color of the balls drawn and are not directly affected by the Box that was randomly selected by the computer. However, while you are drawing balls, the likelihood of drawing red or blue balls depends on which Box was chosen by the computer for that match.

When all pairs have finished the match and have seen the results, we proceed to the next match. For the next match, the computer randomly reassigns all participants to a new pair, a new role as red or blue, and randomly selects a box for each pair. The new assignments do not depend in any way on the past decisions of any participant including you and are done completely randomly by the computer. The assignments are independent across pairs, across participants, and across matches. This second match then follows the same rules and payoffs as the first match.

This continues for 10 matches, at which point the first part of the experiment is over and you will receive new instructions for part 2.

We will now begin the Practice session and go through 2 practice matches to familiarize you with the computer interface and the procedures. During the practice matches, please do not hit any keys until you are asked to, and when you enter information, please do exactly as asked. Remember, you are not paid for

these 2 practice matches. At the end of the second practice match you will have to answer some review questions. Everyone must answer all the questions correctly before the experiment can begin.

[AUTHENTICATE CLIENTS]

Please double click on the icon on your desktop that says **BCP**. When the computer prompts you for your name, type your First and Last name. Then click SUBMIT and wait for further instructions.

[START GAME]

[SCREEN 3]

You now see the first screen of the experiment on your computer. It should look similar to this screen.

[Point to overhead screen as you explain this]

At the top left of the screen, you see your subject ID. Please record that ID in a piece of paper. You have been randomly matched by the computer with exactly one of the other participants. This pair assignment will change after each match.

You have been assigned a role as red or blue, which you can see in the top part of the screen [point on overhead]. The participant you are matched with has been assigned the opposite role (blue or red). Below this information you can see the two possible Boxes, one with more red balls and one with more blue balls. The number of red and blue balls in each Box is displayed on the left side [point on overhead]. Once you have examined the content of each Box, press "OK". When you and the participant you are matched with have both clicked OK, the computer will randomly pick either Box 1 or Box 2 to be your pair's box for the match.

[SCREEN 4]

You should now see the following screen on your computer. On the left side there is a reminder that the colors of the balls in the Box selected are reshuffled and hidden [point on overhead]. On the right side, you can see the Box selected by the computer. Notice that neither you nor the participant you are matched with can see the colors of the balls. Underneath the Box there is a button labeled "No draw" [point on overhead].

Below the Box, there is a summary of the information relevant for the game [point on overhead]: The cost of each draw for you and for the other participant; The number of red and blue balls that have been drawn so far; And your payoff if no participant draws a ball in which case the match ends.

Remember that if your role is red, then there must be more red balls drawn than blue balls in order for you to obtain 150 points in outcome earnings (otherwise you obtain 50 points). If your role is blue, then there must be at least as many blue balls drawn as red balls in order for you to obtain 150 points in outcome earnings (otherwise you obtain 50 points). Note that if there are an equal number of red and blue balls drawn (including the case of 0 balls drawn in the match) the participant in the blue role receives 150 points in outcome earnings and the participant in the red role receives 50 points in outcome earnings.

At this time, please click on one of the hidden balls. Of course during the real experiment you could have chosen "no draw" instead. You will then see a screen similar to this one.

[SCREEN 5]

We are now beginning round two of this practice match. The screen is similar to the previous one with some updated information. It shows the draws that you and the other participant have made. In this example, each of the participants drew a red ball [point to screen]. Remember, this does not imply that the Box must be Box 1 because each ball was effectively drawn with replacement. Every time a ball is drawn it is put back in the Box, and the contents are reshuffled before the next draw and the colors hidden [point to screen]. This screen also says what your outcome earnings will be if you and the other participant do not draw any more balls in the match. The screen also shows the number of draws you have made. [point to screen]

Please click "No draw" now. Since both you and the participant in your pair chose no draw, this will end the match. In the real experiment the match would continue to another round if either of you draws a ball in the current round.

[SCREEN 6]

At this point, the match has ended and you should all see a screen similar to this one. On the right side, you are now shown which Box the computer selected for your pair at the beginning of the match, along with the latest position of the 3 reshuffled balls [point to screen]. On the left side, there is a summary of the information relevant in this match: the number and colors of all the balls drawn by both participants, your number of draws, the outcome of the match, and your payoff for the match, that is, your outcome earnings minus your draw costs. [point to screen]

The bottom half of your screen contains a table summarizing the results for all matches you have participated in [point to screen]. This is called your history screen. It will be filled out as the experiment proceeds. Notice that it only shows the results from your pair, not the results from any of the other pairs.

[Describe the History Screen]

Here is a brief recap of the important things you will see and how they affect your payoffs.

[SCREEN 7 - read it]

Are there any questions? We now proceed to the second practice match. For this match you will be randomly re-matched into pairs, and the computer will randomly assign you to a red or blue role, with the other member having the opposite role.

[START next MATCH]

Please notice your new role. Then click "OK". Then wait. The computer has now randomly assigned the Box for your pair, which stays the same for the match. Click "No Draw".

[WAIT for them to complete match 2]

Practice match 2 is now over. Notice that if you are Blue, your payoff is 150 and if you are Red your payoff is 50, because there are the same number of red and blue balls drawn (zero, in this case).

You must now answer the review questions correctly before we begin the paid session. Once you answer all the questions correctly, click submit. After both participants in your pair have correctly answered the first round of questions, a second round of questions will appear. After both participants in your pair have correctly answered the second round of questions, a third round of questions will appear. When you have answered these, click submit and the quiz will disappear from your screen.

[WAIT for everyone to finish the Quiz]

Are there any questions before we begin with the paid session? We will now begin with the 10 paid matches of part 1. Please pull out your dividers for the paid session of the experiment. If there are any problems or questions from this point on, raise your hand and an experimenter will come and assist you.

Acknowledgments

We thank M. Castro, S.H. Chew, G. Frechette, A. Gaduh, D. Houser, E. Kamenica, S. Singhal, the Associate Editor, two referees 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, the Stanford Institute for Theoretical Economics and the Conference on Empirical Legal Studies at USC 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 and SES-0962802), The Social Science Experimental Laboratory at Caltech, The Gordon and Betty Moore Foundation, the LUSK Center for Real Estate, the Office of the Provost at the University of Southern California and the Microsoft Corporation.

References

1. Banks, J., Olson, M., Porter, D.: An experimental analysis of the bandit problem. *Econ. Theory.* 10, 55-77 (1997)
2. Breitmoser, Y., Tan, J.H.W., Zizzo, D.J.: Understanding perpetual R&D races. *Econ. Theory.* 44, 445-467 (2010)
3. Brocas, I., Carrillo, J.D.: Influence through ignorance. *RAND J. Econ.* 38, 931-947 (2007)
4. Camerer, C.F.: Individual decision making. In Kagel, J., Roth, A. (eds.) *The Handbook of Experimental Economics*, pp. 587-703. Princeton, NJ (1995).
5. Charness, G., Levin, D.: When optimal choices feel wrong: a laboratory study of Bayesian updating, complexity, and affect. *Amer. Econ. Rev.* 95, 1300-1309 (2005)
6. Cox, J.C., Oaxaca, R.L.: Laboratory tests of job search models. In Plott, C., Smith, V. (eds.) *Handbook of Experimental Economic Results*, pp. 311-318. North Holland (2008)
7. Gentzkow, M., Kamenica, E.: Competition in persuasion. Unpublished manuscript (2011)
8. Gul, F., Pesendorfer, W.: The war of information. Unpublished manuscript (2009)
9. Harris, C., Vickers, J.: Racing with uncertainty. *Rev. Econ. Stud.* 54, 1-21 (1987)
10. Horner, J.: A perpetual race to stay ahead. *Rev. Econ. Stud.* 71, 1065-1088 (2004)
11. Kamenica, E., Gentzkow, M.: Bayesian persuasion. Forthcoming in *Amer. Econ. Rev.* (2010)
12. Kahneman, D., Slovic, P., Tversky, A.: *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge: UK (1982)
13. Konrad, K.A., Kovenock, D.: Multi-battle contests. *Games Econ. Beh.* 66, 256-274 (2009)

14. Matthews, S.; Postlewaite, A.: Quality testing and disclosure. *RAND J. Econ.* 16, 328-340 (1985)
15. Milgrom, P.; Roberts, J.: Relying on the information of interested parties. *RAND J. Econ.* 17, 18-32 (1986)
16. Schotter, A.; Braunstein, Y.M.: Economic search: an experimental study. *Econ. Inquiry.* 19, 1-25 (1981)
17. Zizzo, D.J.: Racing with uncertainty: a patent race experiment. *Int. J. Ind. Org.* 20, 877-902 (2002)




Figure 1. Sampling strategies by BLUE and RED adversaries



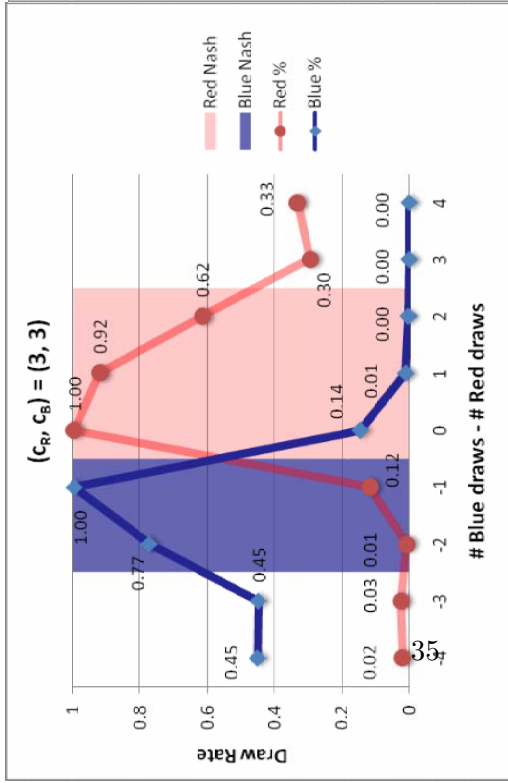
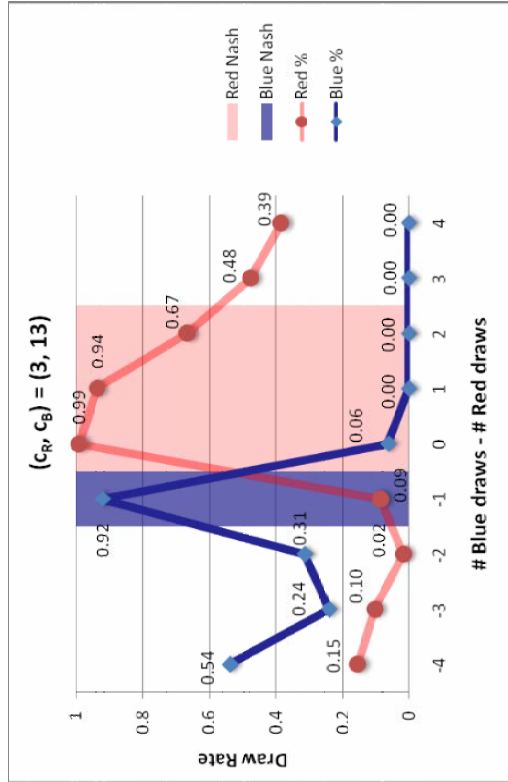


Figure 2. Sampling frequencies by state and cost treatment

| 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) | (.02) | (.01) | (.03) | (.00) | (.02) | (.04) | (.07) | (.13) |
| 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) | (.10) | (.06) | (.04) | (.00) | (.03) | (.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) | (.16) | (.08) | (.01) | (.02) | (.00) | (.02) | (.04) | (.07) | (.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) | (.21) | (.09) | (.05) | (.02) | (.02) | (.00) | (.00) | (.00) | (.00) |
| $(c_R, c_B) = (13, 3)$ | | | | | | | | | |
| # 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) | (.02) | (.02) | (.04) | (.04) | (.15) | 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) | (.11) | (.06) | (.03) | (.00) | (.02) | (.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) | (.05) | (.06) | (.20) | (.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) | (.25) | (.07) | (.05) | (.01) | (.03) | (.00) | (.00) | (.00) | (.00) |

Table 3: Sampling frequencies (standard errors clustered at the individual level in parentheses).

| | 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 |
|-------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------|--------------|
| $(c_R, c_B) = (3, 3)$ | | | | | | | | | |
| 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 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) |
| 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 sampling frequencies 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).

| RED | $n = 0$ | | $n = 1$ | | $n = 2$ | | $n = 3$ | | |
|---------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| | (3, 13) | (13, 3) | (3, 13) | (13, 13) | (3, 13) | (13, 3) | (3, 13) | (13, 13) | |
| (3, 3) | -0.004 [o] (.262) | -0.076 [o] (.000) | -0.031 [o] (.000) | -0.399 [-] (.000) | -0.369 [-] (.000) | -0.510 [-] (.000) | -0.470 [-] (.000) | -0.153 [o] (.386) | .104 [o] (.496) |
| (3, 13) | — | -0.067 [o] (.000) | -0.026 [o] (.001) | -0.416 [-] (.000) | -0.387 [-] (.000) | -0.560 [-] (.000) | -0.521 [-] (.000) | -0.334 [o] (.087) | -0.077 [o] (.643) |
| (13, 3) | — | — | 0.041 [o] (.002) | — | 0.029 [o] (.490) | — | 0.039 [o] (.413) | — | .257 [o] (.252) |
| BLUE | $n = -1$ | | $n = -2$ | | $n = -3$ | | $n = -4$ | | |
| (3, 13) | (3, 13) | (13, 3) | (13, 13) | (3, 13) | (13, 3) | (3, 13) | (13, 13) | (3, 13) | (13, 13) |
| (3, 3) | -0.074 [o] (.000) | .005 [o] (.198) | -0.029 [o] (.002) | -0.461 [-] (.000) | .098 [o] (.005) | -0.208 [o] (.006) | -0.324 [o] (.000) | .087 [o] (.573) | -0.079 [o] (.439) |
| (3, 13) | — | .079 [o] (.000) | .045 [o] (.012) | — | .559 [+] (.000) | — | -0.116 [o] (.152) | — | -0.166 [o] (.285) |
| (13, 3) | — | — | -0.033 [o] (.000) | — | -0.539 [-] (.000) | — | -0.399 [-] (.000) | — | .428 [o] (.066) |

Table 10: Comparison of sampling across treatments (p-values in parentheses)

| | 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).