

DISCUSSION PAPER SERIES

No. 4907

**A THEORY OF INFLUENCE:
THE STRATEGIC VALUE OF
PUBLIC IGNORANCE**

Isabelle Brocas and Juan D Carrillo

***INDUSTRIAL ORGANIZATION
and PUBLIC POLICY***



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP4907.asp

A THEORY OF INFLUENCE: THE STRATEGIC VALUE OF PUBLIC IGNORANCE

Isabelle Brocas, University of Southern California and CEPR
Juan D Carrillo, University of Southern California and CEPR

Discussion Paper No. 4907
February 2005

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

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

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

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

Copyright: Isabelle Brocas and Juan D Carrillo

ABSTRACT

A Theory of Influence: The Strategic Value of Public Ignorance*

We analyse an agency model where one individual decides how much evidence he collects. We assume that he has free access to information, but all the news acquired becomes automatically public. Conditional on the information disclosed, a second individual with conflicting preferences undertakes an action that affects the payoff of both agents. In this game of incomplete but symmetric information, we show that the first individual obtains rents due to his superior ability to decide whether to collect or forego evidence, i.e., due to his control in the generation of (public) information. We provide an analytical characterization of these rents, that we label 'rents of public ignorance'. They can be interpreted as, for example, the degree of influence that a chairman can exert on a committee due exclusively to his capacity to decide whether to keep discussions alive or terminate them and call a vote. Last, we show that similar insights are obtained if the agent decides first how much private information he collects and then how much of this information he transmits to the other agent.

JEL Classification: D82

Keywords: experimentation, incomplete and symmetric information, information control, informational rents, learning, optimal stopping rule, principal-agent and public ignorance

Isabelle Brocas
Department of Economics
University of Southern California
3022 S. Vermont Ave
Los Angeles, CA 90089-0253
USA
Tel: (1 213) 7408 842
Email: brocas@usc.edu

Juan D Carrillo
Department of Economics
University of Southern California
3022 S. Vermont Ave.
Los Angeles, CA 90089-0253
USA
Tel: (1 213) 740 3526
Fax: (1 213) 740 8543
Email: juandc@usc.edu

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

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

*We thank J Matsusaka, M Riordan, F Squintani, seminar participants at Columbia (Economics and Business School), Duke, Hunter College, U Michigan, U Rochester, UW Seattle, SUNY Buffalo, UC Los Angeles, UC Riverside and U Southern California (Economics and Business School) for helpful comments.

Submitted 29 January 2005

1 Motivation

How can I induce a rational individual whose preferences are different from mine to take actions that are close to my own interests? One possibility builds on the (by now classical) asymmetric information paradigm, where an agent strategically uses his private knowledge to obtain some “rents” (possibly at the expense of other agents).¹ Another possibility builds on an information control paradigm, where an individual without private information influences other people’s decision by strategically choosing the amount of evidence that becomes available to every agent in the economy (including himself). The goal of this paper is to analyze this second possibility. In other words, the paper focuses not on the rents that an agent may get from his *possession of private information* as is already well-known, but rather on the rents arising from his mere ability to *control the flow of public information*. To the best of our knowledge, this paradigm (complementary but different from the asymmetric information one) has never been explored in the economics literature before. Note that there is a fundamental difference between knowing that the other party has some information that he does not want to share and knowing that the other party has decided not to look for information: the act of no-transmission has some signalling value whereas the act of no-acquisition has no signalling value.

Situations in which one party implicitly or explicitly chooses the amount of information publicly available can be of very different nature. The archetypical example is agenda setting. Usually, one of the main roles of the chairman in a council (board of directors, Parliament, faculty recruiting committee) is to decide the moment at which consultations must stop and a decision has to be reached. If preferences among members are not fully congruent, his ability to keep the discussion alive and let members acquire and share information about the issue at stake or, on the contrary, terminate the debate and call a vote allows him to bias the final decision towards his preferred alternative. A similar effect operates in electoral contexts. Often (although not in US Presidential elections), the incumbent has the ability to decide the date of the election, reappointment or confidence

¹The effect of private information on the payoffs of agents in strategic games has been a cornerstone of the incentive theory literature since the 1970s. Authors have studied many classes of games (adverse selection, signaling, cheap-talk, bargaining, etc.) under different contexts (collusion, dynamics, multi-agency, common agency, etc.) and with applications to a wide range of situations (employer/employee, buyer/seller, regulator/firm, government/tax payer etc.). For an introductory exposition, see Laffont and Martimort (2002).

vote. The capacity to bring forward or postpone an election depending on whether the last measure was well received or harshly criticized by the population will be used to affect the perception of citizens and therefore to increase the likelihood of an electoral success. Some other examples are discussed in the paper.

In order to study this strategic game, we propose a model in which one agent with free access to information (the “leader”, he) decides whether to collect evidence about the state of the economy. Every piece of news is automatically shared with a second agent (the “decision-maker”, she). Therefore, agents play a game of imperfect but symmetric information, and have common beliefs at every point in time. Information acquisition is sequential so, conditional on the content of the information accumulated so far, the leader decides whether to keep adding evidence or stop the learning process. As a first step, we assume that the decision-maker has no ability to acquire information by herself. Once the leader decides to stop the generation of information, she takes an action that affects the payoff of both agents. Since the preferences of the two individuals are not congruent, the leader will use to his own advantage his privileged access to the generation of information. More specifically, given that information is costless for him, his incentives to acquire or forego information depend exclusively on the likelihood that new evidence will move the belief of the decision-maker further in the “right direction” (that is, towards the action preferred by the leader) vs. further in the “wrong direction”.² Stated differently, the costs and benefits of information are two sides of the same coin: news which turn out to be good increase the payoff of the leader but those which turn out to be bad decrease it. The paper characterizes analytically the leader’s optimal stopping rule for the generation of information and his equilibrium rents, defined as the difference in his expected utility between the case in which he decides and the case in which the decision-maker decides how much information is collected (Proposition 1). Naturally, the extent of this influence process will crucially depend on how congruent the preferences of the leader and the decision-maker are: the smaller the conflict, the easier will be for the leader to induce the decision-maker to undertake the action optimal for him (Proposition 2). We also offer some numerical examples that illustrate the power of this public ignorance mechanism as a tool for influence.

²These calculations are made in a probabilistic sense, since the leader has no better knowledge than the decision-maker about which information is likely to come next.

In a second stage, we analyze the case in which the decision-maker can also acquire information, although at a positive cost. We show that her ability to obtain information induces the leader to provide news up to the point where she does not have an incentive to incur the cost of re-initiating the learning process. Overall, introducing this possibility reduces the degree of influence that the leader can have on the decision-maker's choice, but it does not eliminate it (Proposition 3). The leader's rents are proportional to the decision-maker's cost of acquiring information and they vanish when her cost is nil.

The reader might object that, in some situations, the agent first decides whether to privately collect information (and how much) and then, conditional on its content, whether to make it publicly available. In other words, the game may sometimes have two different stages: first, the leader collects/foregoes information (as in the analysis so far, except that news are now private), and second, the leader transmits/withholds his information to the decision-maker. For example, a prosecutor first chooses how deep to privately interrogate his own witness and then which questions to ask in front of the jury.³ Similarly, newspapers decide whether to send a reporter to investigate an affair, and then whether to publish the results of the investigation. This extension is considered in section 4. Interestingly, if we assume that the information released by the leader is verifiable, Milgrom and Roberts (1986) show that the leader will not be able to make a strategic use of his private knowledge. The basic idea is that the decision-maker will adopt a skeptical, no-news-is-bad-news position: any information transmitted is verified and any information withheld is interpreted as negative for her interests. As a result, the leader will not be able to get rents by hiding some of the acquired information. Formally, the game is "as if" the news collected were automatically shared with the decision-maker; therefore, we can apply the same methodology and obtain the same results at the collection of information stage (Proposition 4). More generally, even in contexts where rents can be obtained through information withholding, this does not invalidate the message of the paper: some other rents can still be obtained through information avoidance.

Before presenting the model, we would like to mention different areas of research related

³When the prosecutor faces a witness of the defense unknown to him, the game is closer to the one presented previously: (i) no private information is acquired by the prosecutor; (ii) the interrogation reveals information to the jury and the prosecutor simultaneously; and (iii) the prosecutor still has the ability to proceed or stop the interrogation. Note however, that this example has two leaders with conflicting goals (prosecutor and defense attorney), an extension mentioned but not formally analyzed in the paper.

to our paper. First, there exists a literature that studies the optimal delegation of decision rights by a principal to one or several agents when the latter are better informed than the former (or, equivalently, have an easier access to information) and the interests of the two parties are not perfectly congruent (see e.g., Gilligan and Krehbiel (1987), Aghion and Tirole (1997), Marino and Matsusaka (2004) and Aghion, Dewatripont and Rey (2004)). Contrary to this literature, we exogenously assume that the agent has control over actions and, instead, we endogenize the decision to generate or avoid information. Second, the optimality of ignorance in multi-agent contexts has also been widely analyzed in incentive theory. For example, Crémer and Khalil (1992) study the optimal contract designed by a principal when the agent can spend resources to obtain information before signing it. The trade-off from the agent’s perspective is the cost of information vs. the private strategic use of the news.⁴ Anticipating the incentives of the agent, the principal designs the optimal contract in a way that, in equilibrium, the agent does not obtain any information. In our paper, information is costless for the leader but becomes publicly available. His decision to forego evidence is based exclusively on the likelihood that news will move the beliefs of the other party towards the “desired” vs. the “undesired” direction. Hence, even if we share the result that an individual obtains some rents from his decision to forego evidence, the reasons are of very different nature. Last, since the model is based on an optimal stopping rule in the acquisition of information, it shares some characteristics with the multi-armed bandit and learning by experimentation literatures (see e.g., Aghion et al. (1991), Bolton and Harris (1999) or Keller and Rady (1999)).

The plan of the paper is as follows. In section 2, we present the model and some motivating examples. In section 3, we characterize the rents of public ignorance when the decision-maker can and cannot obtain information by herself. In section 4, we introduce the possibility of private acquisition of information by the leader. In section 5, we conclude and suggest directions for future research.

2 A model of influence

We consider the following game. One agent (the “leader”, he) has free access to information about the state of the economy but every piece of evidence collected becomes automatically

⁴In particular, there is nothing interesting to say when the agent has free access to information.

public. Based on the news obtained by the leader and observed by everyone, a second agent (the “decision-maker”, she) undertakes an action that affects the payoff of both individuals. Since there is a conflict of preferences between the two parties, the leader might use his ability to collect or forego information to his own advantage.

A crucial feature that differentiates our setting from most of the standard agency literature is that, in our framework, there is *no asymmetric information*. The leader controls the flow of news but, in case of deciding to obtain some pieces of information, these are revealed to both players simultaneously. Technically, agents play a game of incomplete but symmetric information. Thus, contrary to the usual hidden information literature in which agents may derive rents from their *possession of private information*, in this paper the leader will eventually get some rents due exclusively to his *control of the flow of public information*. Furthermore, the two agents cannot contract on the information to be revealed during the game. Again, this allows us to better isolate the effect of information control instead of information possession.⁵

2.1 Preliminaries

We use the following model to analyze this game. There are two possible states in the economy, $s \in \{A, B\}$. Agents have imperfect but symmetric information about states. They share a prior belief p that the true state is A , that is, $\Pr(A) = p$ and $\Pr(B) = 1 - p$. Conditional on the information released by the leader during the game, the decision-maker will choose among three possible actions, $\gamma \in \{a, o, b\}$. We denote by $u_i(\gamma)$ the utility of agent $i \in \{l, d\}$ (where l stands for leader and d for decision-maker). The conflict of preferences between l and d is modelled as follows:

$$u_d(a) = \begin{cases} 1/\alpha & \text{if } s = A \\ 0 & \text{if } s = B \end{cases}, \quad u_d(o) = 1, \quad u_d(b) = \begin{cases} 0 & \text{if } s = A \\ 1/\beta & \text{if } s = B \end{cases} \quad \text{with } (\alpha, \beta) \in (0, 1)^2 \quad (1)$$

$$u_l(a) = x, \quad u_l(o) = y, \quad u_l(b) = 0 \quad \forall s \text{ and } x > y > 0 \quad (2)$$

Denote by μ the posterior belief that the true state is A conditional on the information

⁵In section 4, we show that, as long as information is verifiable, similar results are obtained in a two-stage model where the leader first decides how much private information he collects and then how much of this information he reveals to the decision-maker.

transmitted during the game. Given (1), the decision-maker maximizes her expected payoff by taking action a when $\mu \geq \alpha$, action o when $\mu \in [1 - \beta, \alpha)$, and action b when $\mu < 1 - \beta$. By contrast and given (2), the leader prefers the decision-maker to take action a rather than o and action o rather b , independently of the true state of the economy. This conflict of interests is graphically represented in Figure 1.

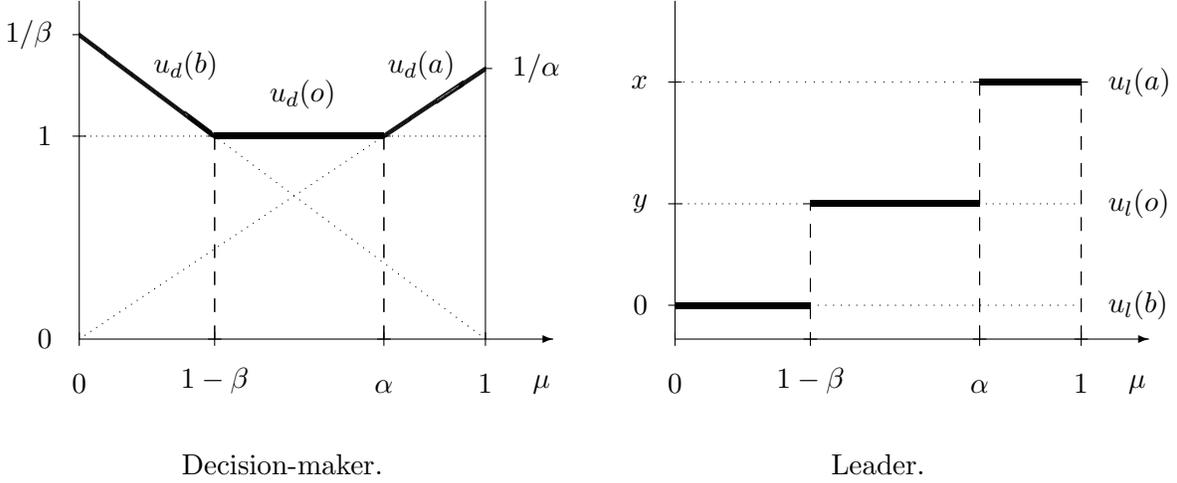


Figure 1. Payoff of Decision-maker and Leader.

Remark 1. It is important to notice that there would be no gain in generality if we assumed that the payoffs of the leader also depend on the true state. The theory presented below relies only on two elements: a conflict of interests between the two agents and a limited action space for the decision-maker (see also remark 3 in section 3.2).

The following assumption ensures that each action is potentially optimal depending on the beliefs of the decision-maker.

Assumption 1 $1 - \beta < \alpha$.⁶

The structure of information acquisition is the simplest one that we can think of. At each moment in time, the leader decides whether to make available one extra signal

⁶In case of payoff-indifference d takes the action preferred by l . Note that if $\alpha < 1 - \beta$, then only a and b can be optimal. Last, $\alpha = \beta \in (1/2, 1)$ corresponds to the pure symmetric case, in which the payoff when the “correct” action from d ’s perspective is taken (a if $s = A$ and b if $s = B$) is the same in both cases.

$\sigma \in \{a', b'\}$ or not. Signals are imperfectly correlated with the true state. Formally:

$$\Pr[a' | A] = \Pr[b' | B] = \theta \quad \text{and} \quad \Pr[a' | B] = \Pr[b' | A] = 1 - \theta,$$

where $\theta \in (1/2, 1)$. Note that as θ increases, the informational content of each signal σ also increases. When $\theta \rightarrow 1/2$ signals are completely uninformative and when $\theta \rightarrow 1$ one signal perfectly informs the leader about the true state. We assume that generating information is neither costly nor produces a delay. So, in particular, the leader can decide to disclose a countless number of signals, in which case agents learn the true state almost with certainty.⁷ Moreover, since it is a game of imperfect but symmetric information (both the leader and the decision-maker observe the signals), then the two parties (i) simultaneously update their belief using Bayes rule and (ii) share a common posterior belief at every stage of the game. Naturally, the leader's decision whether to keep accumulating evidence will be contingent on the realization of past signals.⁸

Suppose that the information generated is such that a number n_a of signals a' and a number n_b of signals b' have been released. Using standard statistical techniques, it is possible to compute the posterior belief shared by the two agents:

$$\begin{aligned} \Pr(A | n_a, n_b) &= \frac{\Pr(n_a, n_b | A) \Pr(A)}{\Pr(n_a, n_b | A) \Pr(A) + \Pr(n_a, n_b | B) \Pr(B)} \\ &= \frac{\theta^{n_a - n_b} \cdot p}{\theta^{n_a - n_b} \cdot p + (1 - \theta)^{n_a - n_b} \cdot (1 - p)}. \end{aligned}$$

The relevant variable which will be used from now on is $n \equiv n_a - n_b \in \mathbb{Z}$, that is the difference between the number of signals a' and the number of signals b' . Besides, we define the posterior $\mu(n) \equiv \Pr(A | n_a, n_b)$.⁹

Remark 2. From a modelling viewpoint, it is equivalent to assume that the leader sequentially chooses the number of pieces of information (e.g., the time spent by a reporter to cover a certain event) or that information exogenously arrives and the leader chooses when to stop its flow (e.g., finish the debate and call a vote).

⁷This assumption is clearly unrealistic: information is always costly. Yet, it allows us to isolate the net benefit from the leader's perspective of disclosing an extra piece of news.

⁸For recent principal-agent models that compare simultaneous vs. sequential acquisition of information, see Gromb and Martimort (2004) and Li (2004). These papers, however, are still based on the asymmetric information paradigm (signals are privately observed).

⁹Some properties of $\mu(n)$ are: (i) $\lim_{n \rightarrow -\infty} \mu(n) = 0$, (ii) $\lim_{n \rightarrow +\infty} \mu(n) = 1$, and (iii) $\mu(n+1) > \mu(n) \quad \forall n$.

An important technical assumption that will be maintained throughout the paper is the following.

Assumption 2 *We treat n as a real number.*

Obviously, this is a strong mathematical abuse. We make this assumption only to avoid dealing with integer problems in the technical resolution of the model. Note that the difference between prior and posterior after one signal is increasing in θ . Therefore, as long as the informational content of each signal is sufficiently small (technically, θ close to $1/2$), this assumption implies little loss of generality.

2.2 Examples

Before starting the formal analysis of the game, it will be useful to briefly discuss the kind of situations our model is designed to capture. As described above, the main elements of our model are: (i) a conflict of preferences between the two agents modelled by (1) and (2), (ii) a flow of information controlled by one of the agents, and (iii) no possibility of contracting on the amount of information to be revealed. The examples presented below are only suggestive.

The most common example is decision-making in small committees. Discussions among hiring committee members in Economics departments elicit valuable public information about the quality of the potential candidates. At the same time, differences in tastes for fields, teaching needs, etc. imply that preferences among members are not fully congruent. This paper studies how the chair can influence the final decision only with his capacity to decide whether to keep the discussion alive (e.g. by asking for another reference letter that every member will have access to, scheduling a new meeting to further discuss the case, proposing an alternative candidate, etc.) or terminate the debate and call a vote. Obviously, agenda setters in parliamentary debates, syndicate meetings and board of directors reunions have the power to influence choices in a similar manner.

A related effect occurs in politics. In many countries, the Constitution imposes an upper limit in the duration of electoral mandates. However, within that limit, the President decides when to hold the election. Depending on how the current (publicly observed) performance of the executive committee is perceived by the electorate, the President may choose to call an election immediately or postpone it and make a few more decisions.

Last, more controversial but also suggestive is the case of media coverage. Frequently, citizens are forced to rely exclusively on the news provided by media professionals in order to form their view about political events. A newspaper may influence the public opinion concerning, for example, US foreign policy in a given country without falsifying information. All it needs to do is, depending on the turn of events, strategically decide whether to maintain its reporters in that country and keep reporting the news as they occur or bring them back and stop covering the case.¹⁰

Summing up, in all these examples, the leader (chairperson, President, media) will use his ability to acquire or forego information (call a vote, hold the election, send the correspondent back) to his own advantage. Naturally, the decision-maker (committee member, citizen, public) realizes this strategic decision. However, since knowledge is symmetric, there is no information or signal she can deduce from his behavior. This is why, knowing that the leader has some information that has not been released (the classical signaling effect under asymmetric information) is fundamentally different from knowing that the leader has refrained from acquiring evidence (the public ignorance effect under symmetric information of this paper).

2.3 Information acquisition and posterior beliefs

A first step to characterize the rents obtained due to the leader’s ability to acquire or forego information is to determine the likelihood of reaching different beliefs conditional on the information currently available. More specifically, suppose that l currently believes that A is the true state with probability p (from now on, we will for short say that l holds “belief p ”). Suppose also that he stops acquiring information when he reaches a belief p_H ($> p$) or a belief p_L ($< p$). These cutoffs are exogenously given for the time being. What is the probability of reaching one posterior before the other? Naturally, this will crucially depend on whether the true state is A or B . Formally, denote by $q^s(p; p_L, p_H)$ the probability of reaching p_H before reaching p_L when the initial belief is $p \in (p_L, p_H)$ and the true state is s . Similarly, $q(p; p_L, p_H)$ is the unconditional probability of reaching p_H before p_L (i.e. the probability given the existing uncertainty about the true state of

¹⁰Andrew Ang suggested to us the following related piece of anecdotal evidence. CNN has a special TV channel run in airports. When the AA plane crashed in New York on November 12, 2001, CNN shut-down the channel until the case was clarified.

the world). By definition, we have $q^s(p_L; p_L, p_H) = 0$ and $q^s(p_H; p_L, p_H) = 1$ for all s . Interestingly, given our simple information acquisition game, it is possible to obtain analytical expressions of these probabilities. These are gathered in Lemma 1 and they are key for our analysis.¹¹

Lemma 1 $q^A(p; p_L, p_H) = \frac{p - p_L}{p_H - p_L} \times \frac{p_H}{p}$ and $q^B(p; p_L, p_H) = \frac{p - p_L}{p_H - p_L} \times \frac{1 - p_H}{1 - p}$.
 Moreover, $q(p; p_L, p_H) \equiv p \times q^A(p; p_L, p_H) + (1 - p) \times q^B(p; p_L, p_H) = \frac{p - p_L}{p_H - p_L}$.

Proof. See Appendix A1. □

Basically, Lemma 1 states that the probability of reaching a posterior upper bound p_H before a posterior lower bound p_L is proportional to the distance between the upper bound and the prior ($p_H - p$) relative to the distance between the prior and the lower bound ($p - p_L$). For our model, it means that as the decision-maker's payoff of taking action a under state A increases (i.e. as α decreases, see (1)), the cutoff above which the decision-maker is willing to take action a decreases (see Figure 1). This in turn implies that the distance ($\alpha - p$) between the prior and that cutoff shrinks. As a result and other things being equal, the decision-maker is more likely to end up with a posterior belief in which she finds it optimal to take that action (a similar argument occurs with β and b).

Note also that $q^A(\cdot) > q(\cdot) > q^B(\cdot)$ for all p , p_L and p_H . By definition, the likelihood of obtaining a' rather than b' signals is higher when the state is A than when the state is B . Since a' signals move the belief upwards (towards state A) and b' signals move it downwards (towards state B), then for any starting belief p , it is more likely to reach an upper bound p_H before a lower bound p_L if the state is A than if the state is B . Last but not least, $q^A(p; 0, p_H) = 1$, and $q^B(p; p_L, 1) = 0$: the agent can never believe with certainty that one state is true when in fact it is not. An interesting corollary follows from the previous analysis.

Corollary 1 $q^A(\cdot)$, $q(\cdot)$ and $q^B(\cdot)$ do not depend on θ , the informational content of a signal.

¹¹Technically, the question amounts to determine the evolution of a stochastic process with two absorbing states. Although the case of one absorbing state is well-known in the literature, despite our efforts we have not been able to find in a Statistics manual the analytical characterization for the case of two absorbing states. We therefore prove it in Appendix A1. It can be trivially checked that if we set $p_H = 1$ or $p_L = 0$, then the results of the one absorbing state case are replicated.

This result is obtained by direct inspection of the analytical expressions derived in Lemma 1, and it may at first seem surprising. The idea is that, as long as we do not take into account integer problems (see assumption 2), then the informational content θ of each signal affects the *speed* at which one of the posteriors is reached but not the *relative probabilities* of reaching each of them. Roughly speaking, more accurate information (higher θ) implies greater chances of receiving the “correct” signal (a' if $s = A$ and b' if $s = B$) but also that an “incorrect” signal moves the posterior farther away in the opposite direction. These two effects perfectly compensate each other.¹²

3 The optimal control of information generation

3.1 Influence through public ignorance

Since the relevant cutoff beliefs that determine the optimal action to be taken by the decision-maker are α and $1 - \beta$, it will prove useful to define $\pi_a \equiv \alpha$ and $\pi_b \equiv 1 - \beta$. According to this new notation, for a given posterior μ that the true state is A , d strictly prefers to undertake action b if $\mu \in [0, \pi_b)$, action o if $\mu \in [\pi_b, \pi_a)$ and action a if $\mu \in [\pi_a, 1]$.

Given the conflict of preferences between the two agents, the leader will use his control of the flow of information to induce the decision-maker to undertake action a rather than o or action o rather than b . As information is revealed to both players simultaneously, this influence can only be achieved through the decision to stop collecting additional news. In this context, an optimal stopping rule is characterized by a pair of probabilities $(p_L, p_H) \in (0, 1)^2$ such that l does not collect extra information whenever one of these posterior beliefs is reached.¹³ Using Lemma 1, we get the main result of the paper.

Proposition 1 *Suppose that $p \in (\pi_b, \pi_a)$. Two cases are possible:*¹⁴

¹²Note that, if there was a positive fixed cost per signal, then the willingness of the leader to collect information would depend on θ (a higher θ would mean a more informative signal for the same cost). However, the relative probabilities of reaching one posterior vs. the other would still remain unaffected.

¹³We assume that $p_L \notin \{0, 1\}$ and $p_H \notin \{0, 1\}$ because the collection of signals by l eventually has to stop in order for d to choose an action. However, we allow p_L and p_H to be arbitrarily close to 0 and 1. This means that l can choose to stop the collection of information only when he knows the true state almost with certainty (and there is no cost associated to that).

¹⁴For the sake of completeness we have: (i) if $p > \pi_a$, then $p_L = p_H = p$ and l gets utility x ; (ii) if $p < \pi_b$ and $x > \pi_a/\pi_b$, then $p_L \rightarrow 0$, $p_H = \pi_a$ and l gets an expected utility $x p/\pi_a$; (iii) if $p < \pi_b$ and $x < \pi_a/\pi_b$, then $p_L \rightarrow 0$, $p_H = \pi_b$ and l gets an expected utility $y p/\pi_b$.

(i) If $y/x < \pi_b/\pi_a$, then $p_L \rightarrow 0$ and $p_H = \pi_a$. The leader's expected utility is $U_{(0,\pi_a)} = x \times \frac{p}{\pi_a}$ and that of the decision-maker is $V_{(0,\pi_a)} = \frac{\pi_a - p \pi_b}{\pi_a - \pi_a \pi_b} (> 1)$.

(ii) If $y/x > \pi_b/\pi_a$, then $p_L = \pi_b$ and $p_H = \pi_a$. The leader's expected utility is $U_{(\pi_b,\pi_a)} = x \times \frac{p - \pi_b}{\pi_a - \pi_b} + y \times \frac{\pi_a - p}{\pi_a - \pi_b}$ and that of the decision-maker is $V_{(\pi_b,\pi_a)} = 1$.

Proof. See Appendix A2. □

Proposition 1 shows that the leader derives rents from his ability to control the generation of public information or, put it differently, from the possibility to stop the flow of news. This comes at the expense of the decision-maker. Indeed, if the decision-maker could decide on the amount of information to be generated during the game, she would force the leader to learn with almost certainty whether the true state is A or B (formally, $p_L \rightarrow 0$ and $p_H \rightarrow 1$). Her expected payoff and that of the leader would be $V_{(0,1)} = \frac{p}{\pi_a} + \frac{1-p}{1-\pi_b}$ (greater than both $V_{(0,\pi_a)}$ and $V_{(\pi_b,\pi_a)}$) and $U_{(0,1)} = x \times p$ (smaller than both $U_{(0,\pi_a)}$ and $U_{(\pi_b,\pi_a)}$), respectively.

The keys to determine the leader's optimal stopping rule are the following. First, once the posterior π_a is reached, the leader has no further incentive to collect information (under this belief, d takes action a , which provides the greatest payoff to l). Second, the leader will never stop accumulating evidence if $\mu \in (\pi_b, \pi_a)$: his payoff is y so he can, at least, wait until either π_b or π_a is hit. Third, if $\mu < \pi_b$ his payoff of stopping is 0; again, he might as well keep sampling until he either learns almost with certainty that the true state is B ($\mu \rightarrow 0$) or hits π_b . Therefore, the only remaining question is whether to stop at $\mu = \pi_b$ and obtain a payoff y with certainty or keep providing evidence. In the latter case, with probability π_b/π_a the posterior $\mu = \pi_a$ is hit and the leader obtains a payoff x , and with probability $1 - \pi_b/\pi_a$ agents learn that the true state is B almost with certainty and the leader obtains a payoff of 0.

Note that, as the leader's payoff under action o gets close to his payoff under action a (i.e., as y/x increases), his willingness to stop at a posterior π_b and accept action o rather than gamble for a or b also increases. More generally, the fact that information is symmetric for both players implies that the costs and benefits of collecting information are two sides of the same coin: extra evidence may move the belief of the decision-maker towards

or against the direction preferred by the leader. Depending on the relative likelihood of these two forces, the leader will choose whether to keep generating information or not.

Overall, the ability to control the flow of news and remain publicly ignorant gives the leader some power, which is used to influence the actions of the decision-maker. It is essential to notice that the decision-maker realizes that the leader controls the generation of information to his own advantage. However, contrary to the asymmetric information models where no-news (i.e., no transmission of information) signals that the leader has bad-news, in this model no-news has no signaling value; it can only be interpreted as the leader being satisfied with the existing information or “fearing” what may come next. Also, and by the same token, it is never in the decision-maker’s best interest to refuse pieces of information presented to her, even if she could commit to.¹⁵

As the reader can notice, we have focused on the simplest conflict of preferences between our two agents to highlight the value of information control in its roughest form. Yet, the analytical characterization of the functions $q(\cdot)$, $q^A(\cdot)$ and $q^B(\cdot)$ provided in Lemma 1 is general enough to allow an application of this same public ignorance principle to more complex economic settings.

3.2 Comparative statics

How much influence can be exerted through public ignorance is an empirical issue, largely beyond the scope of this paper. However, we can present some numerical examples to provide a quantitative sense of the effects described in the previous section. Suppose that l ’s difference in payoffs between actions a and o is large, so that $y/x < \pi_b/\pi_a$. In this case, $p_L \rightarrow 0$ and $p_H = \pi_a (= \alpha)$. Given these stopping rules, d will never end up taking action b contrary to her best interest (formally, $\Pr(b | A) = \lim_{p_L \rightarrow 0} 1 - q^A(p; p_L, \alpha) = 0$). The likelihood that the leader induces her to take his optimal action (a) contrary to her own best interest (when $s = B$) is given by the following simple formula:

$$\Pr(a | B) = \lim_{p_L \rightarrow 0} q^B(p; p_L, \alpha) = \frac{p}{1-p} \times \frac{1-\alpha}{\alpha}$$

¹⁵In our specific case and given that $p \in (\pi_b, \pi_a)$, if she refuses information then she takes action o and gets a payoff of 1 ($= V_{(\pi_b, \pi_a)} < V_{(0, \pi_a)}$). Also, her expected payoff increases with the amount of information provided (in our simple game, $V(\cdot)$ is greater than 1 only if the leader does not stop the acquisition of news when $\mu = \pi_b$). Naturally, the conclusions extend to more general settings.

Hence, starting from a prior $p = 1/2$, if $\alpha = 2/3$ this probability is as large as $1/2$. Some other probabilities are illustrated in Figure 2.

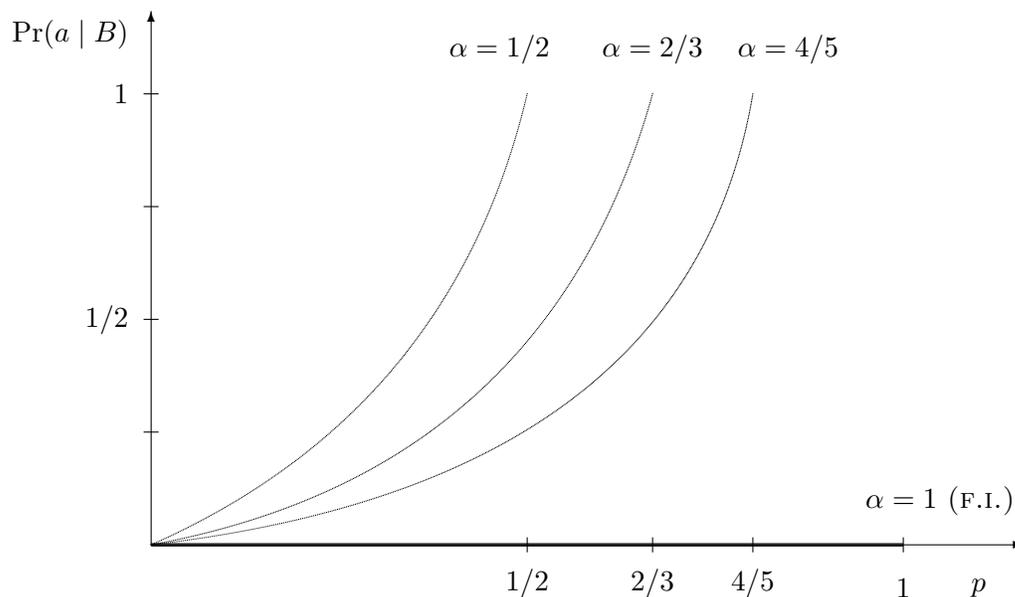


Figure 2. Probability that d takes action a when $s = B$ and $y/x < (1 - \beta)/\alpha$.

Remark 3. For the leader to obtain rents from his control of the flow of information, the decision-maker must necessarily have a restricted action space. For instance, suppose that the chairman wants to hire a candidate independently of his quality (recall that it does not matter in the model whether l 's payoff depends on p or not, see remark 1). If the probability that the committee votes for the candidate is linear in the quality revealed through discussions, then controlling the debate does not generate any extra benefit to the leader. Note however that if there is a threshold quality above which the committee agrees to vote for the candidate and below which the probability of voting for him is linear in quality, the chairman will again be able to derive rents from his control of news.¹⁶

¹⁶Technically, all we need is a non-linearity in l 's payoff as a function of the unknown parameter. Restricting d 's action space to 3 alternatives is the simplest but by no means the only way to achieve it.

3.3 Which decision-maker is more easily influenced?

There is a trivial answer to this question: the decision-maker whose interests are more congruent with the interests of the leader. This means, a decision-maker almost indifferent between actions b and o when $s = B$ (β close to 1) and highly interested in action a rather than o when $s = A$ (α close to 0). In words, a committee member who, just like the chairman, is very much inclined to hire the candidate independently of his quality.

Often, the decision-maker will be either strongly or weakly concerned by making the right choice (e.g., strongly or weakly concerned by hiring a candidate if and only if he is good). The question we ask is: does the leader prefer to face a decision-maker strongly affected or weakly affected by the quality of her decision? On the one hand it is easier to persuade the first type of decision-maker to take action a rather than action o . On the other hand, it is more difficult to persuade her to take action o rather than action b . From Proposition 1(i), we know that if l 's payoff under action o is sufficiently low relative to the payoff under action a (y/x small), then d 's concern for b relative to o is irrelevant. As a result, the leader strictly prefers to face a strongly concerned decision-maker, so as to maximize the likelihood of inducing her to take action a . The interesting situation arises when the leader wants to avoid action b . The next Proposition deals with this case.

Proposition 2 *If $y/x > \pi_b/\pi_a$, then $\frac{\partial U}{\partial \alpha} + \frac{\partial U}{\partial \beta} \propto \frac{\alpha + (1 - \beta)}{2} - p \quad \left(= \frac{\pi_a + \pi_b}{2} - p \right)$.*

Proof. Immediate if we take the derivative of $U_{(\pi_b, \pi_a)}$. □

Remember that α ($= \pi_a$) and $1 - \beta$ ($= \pi_b$) are the posterior beliefs where the leader stops the provision of information when $y/x > \pi_b/\pi_a$. Proposition 2 states that, if the prior belief p is closer (respectively, more distant) to the upper cutoff than to the lower one, then an increase in the decision-maker's payoff of taking her first-best action –i.e. a decrease in α and β – is beneficial (respectively, harmful) for the leader. The idea is that the leader has more flexibility in the collection of news when it is relatively less important for the decision-maker to make the correct choice, because the posteriors where the flow of information is stopped (π_a and π_b) are farther away from the prior p . This increased leeway is valuable if the decision-maker is initially biased against the leader's most preferred action, since it implies a greater chance to reverse the initial handicap. For

the same reason, more leeway is detrimental when the decision-maker is initially biased in favor of the leader's optimal action, since it implies greater chances of reversing the initial advantage. The result then suggests that a chairman prefers a committee member with strong preferences only if her initial beliefs are relatively congruent with his own interests.

3.4 Costly information available to the decision-maker

We have considered the most favorable situation for achieving influence: the leader has free access to information and the decision-maker has no access at all. Often, the decision-maker also has the possibility to affect the generation of public information, possibly at a higher cost than the leader. For example, committee members can spend part of their social capital in requesting more information to be accumulated before casting votes. Also, citizens can independently search for information about events of national interest. We formally introduce this possibility in our model by assuming that, at any moment, the decision-maker can become perfectly informed about the true state s by paying a cost c (> 0) (since only the decision-maker undertakes an action, it is irrelevant whether the state is also revealed to the leader or not). Furthermore, in order to avoid a multiplication of cases, we consider a symmetric payoff situation for the decision-maker:¹⁷

Assumption 3 $\alpha = \beta$ or, equivalently, $\pi_a = 1 - \pi_b$.

Given that information is costly for the decision-maker and that the leader does not have private access to it, the former will always find it optimal to wait until the latter stops providing pieces of news before deciding whether to pay the cost and become perfectly informed. Suppose that the leader stops the flow of information at a given posterior μ . It is possible to characterize the optimal continuation strategy of the decision-maker.

Lemma 2 *If $c \geq (1 - \pi_a)/\pi_a$, then d never learns the true state. If $c < (1 - \pi_a)/\pi_a$, then d optimally learns the true state if and only if her posterior belief is $\mu \in (\pi_a c, 1 - \pi_a c)$.*

Proof. See Appendix A3. □

¹⁷Technically, this technology corresponds to the case in which d pays c for a signal $\sigma \in \{a'', b''\}$ such that $\Pr[a''|A] = \Pr[b''|B] = \rho$ with $\rho = 1$ (perfect correlation between signal and state). The results presented below extend to imperfect correlation ($\rho \in (1/2, 1)$) and asymmetric payoffs ($\pi_a \neq 1 - \pi_b$) but the number of cases to be studied separately grow considerably.

When the acquisition of information is excessively costly (formally, $c \geq (1 - \pi_a)/\pi_a$), the decision-maker strictly prefers to rely on the news disclosed by the leader, even if she anticipates his strategic use of information. The leader anticipates her decision not to re-initiate learning, and therefore keeps the same optimal strategy as in Proposition 1. When information is not too costly (formally, $c < (1 - \pi_a)/\pi_a$), the decision-maker chooses whether to become fully informed or not depending on her current beliefs. Indeed, if she is sufficiently confident that the true state is either A or B (i.e., if $\mu \geq 1 - \pi_a c$ or $\mu \leq \pi_a c$), then the gains of perfect information are small relative to the cost. By contrast, for intermediate beliefs (and, more precisely, when $\mu \in (\pi_a c, 1 - \pi_a c)$) news are highly informative for optimal decision-making, and therefore they are acquired.

Note that if $c < (1 - \pi_a)/\pi_a$, then $\pi_a c < 1 - \pi_a (= \pi_b)$ and $1 - \pi_a c > \pi_a$. In words, the leader knows that if he stops collecting evidence when $\mu \in \{\pi_b, \pi_a\}$, as dictated by Proposition 1(ii), then the decision-maker will for sure re-initiate the learning process and undertake her optimal action. The anticipation of this possibility induces the leader to modify his optimal information gathering strategy.

Proposition 3 *If $c < (1 - \pi_a)/\pi_a$, then $p_L \rightarrow 0$, $p_H = 1 - \pi_a c$ and the decision-maker never re-initiates learning. The leader's utility is $U^c = x \frac{p}{1 - \pi_a c}$ ($< \min \{U_{(0, \pi_a)}, U_{(\pi_b, \pi_a)}\}$) and the decision-maker's utility is $V^c = \frac{1 - \pi_a c(1 + p)}{\pi_a(1 - \pi_a c)}$ ($> \max \{V_{(0, \pi_a)}, V_{(\pi_b, \pi_a)}\}$). Last, $\partial U^c / \partial c > 0$ and $\partial V^c / \partial c < 0$ for all $c > 0$.*

Proof. See Appendix A4. □

When the decision-maker has the ability to acquire some news and these are not excessively costly, the leader is forced to increase the amount of information released, otherwise he will not be able to influence her choices. In particular, o is never going to be selected by the decision-maker in equilibrium since, for any belief $\mu \in [\pi_b, \pi_a]$, she strictly prefers to pay the cost of learning the true state. Given that either a or b will be selected in equilibrium, the optimal way for the leader to influence the decision-maker's choice is to remove any lower bound at which information collection is stopped and push the upper bound up to the point where the decision-maker is willing to take action a . Formally, $p_L \rightarrow 0$ and $p_H = 1 - \pi_a c$.

Overall, the decision-maker's ability to collect information by herself reduces the leader's discretion in the provision of news, and therefore his capacity to manipulate her choices. As the decision-maker's cost c to obtain news decreases, more signals need to be disclosed to avoid re-initialization of the information gathering process, which implies that her expected welfare increases and that of the leader decreases. However, as long as this cost is positive, the leader will always derive rents from his free access to information ($U^c > U_{(0,1)}$ for all $c > 0$).

Interestingly, in our model the leader makes sure that, in equilibrium, the decision-maker never re-initiates the acquisition of information. Yet, her capacity to obtain information is enough to increase her welfare to the detriment of the leader's utility. This conclusion is similar to the results obtained in the contract theory literature on collusion (Tirole, 1986), renegotiation (Dewatripont, 1988), and information acquisition (Cr mer and Khalil, 1992) for example. In these papers, the optimal contract designed by the principal is such that collusion, renegotiation, and acquisition of information never occur in equilibrium. However, just as in our paper, the mere possibility of engaging in these practices affects the payoff of the different players. Last, it is possible to enrich our model in a way that the main conclusions hold and yet the decision-maker sometimes restarts the acquisition of news. One possibility would be to introduce some uncertainty on the decision maker's cost of gathering information.¹⁸

4 Information acquisition and information transmission

In the model presented in section 2, we assumed away any asymmetry of information in order to better isolate the rents extracted with the ability to forego evidence from the rents extracted with the private possession of evidence. However, it is often the case that the leader first decides whether to collect information and then, conditional on its content, whether to transmit it to the decision-maker. Following our previous example, newspapers decide whether to send reporters to cover a story but also whether to release or retain the information discovered. A similar case occurs in auditing. The auditor is concerned with his reputation and tries to provide an accurate judgment of the firm's

¹⁸A similar device is employed by Kofman and Lawarr e (1996) to show that collusion may be an equilibrium outcome in a model   la Tirole (1986).

value. The payoff of the manager (compensation, reappointment) depends on whether the auditor’s appraisal is high or low, independently of the firm’s true value. The manager has usually an easier access to information and will use his capacity to both collect/forego and transmit/withhold evidence to the auditor to his own advantage. As a third example, the role of a defense counsel is not to present to the jury an objective assessment of the case but rather to defend his client independently of his culpability. Once again, the attorney will keep or stop searching for new evidence (“I don’t want to know whether you did kill the person or not”) as a function of his current information. However, he will also choose to present some news to the jury and conceal some others.¹⁹ Note that, for practical or legal reasons, in the situations described above, parties cannot contract on the amount of information to be acquired and revealed.

To capture the type of situations described above, we propose in this section a two-stage variant of the model, where the leader first decides the amount of information privately obtained, and then the amount of information transmitted. Intuitively, the incentives to acquire information are now different: the act of “no-news transmitted” can and will be interpreted by the decision-maker as “bad news possessed”.

To make matters simple, we assume that the preferences of leader and decision-maker are publicly known and that information is verifiable. Suppose that the leader stops acquiring pieces of information when his posterior belief is $\mu \in [0, 1]$. He then provides a report $r(\mu)$ ($\subset [0, 1]$) to the decision-maker. Verifiability implies that the posterior belief must be contained in the report or, formally, that $\mu \in r(\mu)$. As it turns out, the properties of the sequential equilibrium of the information transmission stage of this game have already been studied in a classical paper by Milgrom and Roberts (1986). We can combine their result with our Proposition 1 to determine the equilibrium of our new game.

Proposition 4 *In the two-stage game with private acquisition of verifiable information, the optimal stopping rule of the leader is the same as in Proposition 1. Moreover, in every sequential equilibrium, the action taken by the decision-maker and the payoffs of both agents are also the same as in Proposition 1.*

¹⁹For an interesting comparison of the optimal incentives to acquire and transmit costly information when agents are partisan and non-partisan, see Dewatripont and Tirole (1999).

Proof. Suppose that l stops at a posterior belief μ that $s = A$ and reports $R \subset [0, 1]$ (with $\mu \in R$). Denote by λ the posterior belief inferred by the decision-maker given R .

Step 1 (Information transmission 1). Suppose that $\lambda = \tilde{\mu}$ iff $\nexists \hat{\mu} < \tilde{\mu}$ s.t. $\hat{\mu} \in R$. Milgrom and Roberts (1986, proposition 1) shows that, assuming this inference, then $\nexists \mu' < \mu$ s.t. $\mu' \in R$: the lower bound of the leader's report will always be his true belief.

Step 2 (Information transmission 2). Using this result, Milgrom and Roberts (1986, proposition 2) then shows that, at every sequential equilibrium, the decision-maker infers $\lambda = \tilde{\mu}$ iff $\nexists \hat{\mu} < \tilde{\mu}$ s.t. $\hat{\mu} \in R$.

Step 3 (Information acquisition). Since for any belief μ of the leader, the report R will be such that the decision-maker infers $\lambda = \mu$, the optimal acquisition of information rule is the same as in Proposition 1, and so are the payoffs of both agents. \square

Proposition 4 states that, in our particular game, the leader will not be able to influence the decision-maker by strategically manipulating the transmission of his private information. The idea is simple and intuitive. Suppose that the leader has stopped at a posterior belief μ that the true state is A . Since the decision-maker knows the incentives of the leader to overstate this belief, she will adopt what Milgrom and Roberts coin as a "skeptical posture": always assume the worst. In our case, this corresponds to the lower bound of the report set. Under these circumstances, the leader will never include in the report a posterior below his own belief μ . Furthermore, he will not be able to instil on the decision-maker a posterior above his own belief, even if he tries to. Note that the equilibrium in the information transmission game is not unique: if the leader's belief is μ , then we can only state that his equilibrium strategy will contain μ as the lower bound of the report set. By contrast, the response of the decision-maker is unique (act as if the report was μ), and so are the equilibrium payoffs of the two agents. Since private information cannot be withheld in equilibrium, the optimal information collection strategy of the leader will be the same as in section 3. Overall, under verifiable information and publicly known preferences, no influence can be exerted with the possession of private information but the same influence as before can still be exerted with the control of information generation.

The reader might find unrealistic the fact that information cannot be kept private, at least partially. There are several ways (some of them discussed in Milgrom and Roberts' paper) to restore benefits from possessing private information: forgeable information (i.e.,

relaxing verifiability), a decision-maker who anticipates only partially the incentives of the leader to misreport information (i.e., introducing bounded rationality), agents with imperfect knowledge of each other's preferences (i.e., including a second dimension of private information), etc. Adding these ingredients would invariably make the model more realistic, but they would not invalidate the result that public ignorance can be used to influence the behavior of others.

5 Conclusion

The starting point of the paper was to argue that an agent with privileged access to information can influence the decision of his peers in two fundamentally different ways. First, with his decision to hide or transmit his private information. Second, with his decision to acquire or forego new (private or public) information. While the first mechanism builds on the classical asymmetric information paradigm and has been extensively studied in the literature, to the best of our knowledge the second one has not received any attention. The goal of our paper was to provide a first careful look to the rents that can be extracted with this public ignorance mechanism.

Many other issues related to this paradigm deserve attention. For example, in our framework, the preferences of the different agents are common knowledge. Therefore, the incentives to obtain evidence depend on whether extra information is likely to push the decision-maker to behave in the leader's best or worst interest. It could be interesting to study this same trade-off with several sequential actions and imperfect information about each other's preferences. Our intuition is that the leader may find it optimal to stop accumulating evidence, observe the first action of the decision-maker, deduce (perfectly or imperfectly) her preferences from her behavior and decide whether to re-initiate the acquisition of information.²⁰ We may also learn new insights by combining the paradigms of information acquisition and information transmission, as suggested at the end of section 4. Suppose that the decision-maker cannot perfectly deduce the leader's information from his report, for example because information is not verifiable. How will the leader's ability to manipulate the transmission of information modify his willingness to acquire news in a

²⁰There would be a signaling problem from the decision-maker's viewpoint. However, our framework should be able to handle this case given that Lemma 1 characterizes the relative probabilities of reaching each posterior in a rather general way.

first place? Last, one would also like to determine under which situations will the leader lose part or all the rents of public ignorance. Again at an intuitive level, rents may shrink or even vanish if several leaders with conflicting goals (two newspapers, a defense counsel and prosecutor) compete to provide information. Also, if the relationship between the two agents is repeated, the decision-maker may credibly threaten to take the action least preferred by the leader unless full information is revealed. Our hope is that these questions will stimulate further research on the subject.

References

1. Aghion, P., Bolton, P., Harris, C. and B. Jullien (1991), "Optimal Learning by Experimentation," *Review of Economic Studies*, 58, 621-654.
2. Aghion, P., Dewatripont, M. and P. Rey (2004), "Transferable Control," *Journal of the European Economic Association*, 2(1), pp. 115-138.
3. Aghion, P. and J. Tirole (1997), "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), pp. 1-29.
4. Bolton, P. and C. Harris (1999), "Strategic Experimentation," *Econometrica*, 67(2), 349-74.
5. Crémer, J. and F. Khalil (1992), "Gathering Information Before Signing a Contract," *American Economic Review*, 82(3), pp. 566-578.
6. Dewatripont, M. (1988), "Commitment through Renegotiation-Proof Contracts with Third Parties," *Review of Economic Studies*, 55(3), pp. 377-389.
7. Dewatripont, M. and J. Tirole (1999), "Advocates," *Journal of Political Economy*, 107(1), pp. 1-39.
8. Gilligan, T. and K. Krehbiel (1987), "Collective Decision-making and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures," *Journal of Law, Economics, and Organization*, 3(2), pp. 287-335.
9. Gromb, D. and D. Martimort (2004), "The Organization of Delegated Expertise," *mimeo*, LBS and Toulouse.
10. Keller, G. and S. Rady (1999), "Optimal Experimentation in a Changing Environment," *Review of Economic Studies*, 66(3), 475-507.
11. Kofman and Lawarrée (1996), "On the Optimality of Allowing Collusion," *Journal of Public Economics*, 61(3), pp. 383-407.
12. Laffont, J.J. and D. Martimort (2002), *The Theory of Incentives*, Princeton University Press, Princeton, USA.

13. Li, W. (2004), "Mind Changes and the Design of Reporting Protocols," *mimeo*, UC Riverside.
14. Marino, A. and J. Matsusaka (2004), "Decision Processes, Agency Problems, and Information: An Economic Analysis of Capital Budgeting Procedures," forthcoming in *Review of Financial Studies*.
15. Milgrom, P. and J. Roberts (1986), "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 17, pp. 18-32.
16. Tirole, J. (1986), "Hierarchies and Bureaucracies: on the Role of Collusion in Organizations," *Journal of Law, Economics, and Organization*, 2, pp. 181-214.

Appendix

Appendix A1. Proof of Lemma 1

Recall that the prior belief of state A (i.e. when $n = 0$) is p . Suppose that the posterior belief when the difference of signals reaches $n = +t$ is $\Pr(A \mid +t) = p_H (> p)$ and that the posterior belief when the difference of signals reaches $n = -k$ is $\Pr(A \mid -k) = p_L (< p)$. Denote by $\lambda^s(n)$ the probability of reaching a difference of signals equal to $+t$ before reaching a difference of signals equal to $-k$ when the current difference of signals is n and the true state is s . By definition, we have $\lambda^s(-k) = 0$ and $\lambda^s(+t) = 1$. Besides, $\lambda^s(0) \equiv q^s(p; p_L, p_H)$. From the definition of the transmission of information we have:

$$\lambda^A(n) = \theta \cdot \lambda^A(n+1) + (1-\theta) \cdot \lambda^A(n-1) \quad \forall n \in \{-k+1, \dots, t-1\} \quad (3)$$

$$\lambda^B(n) = (1-\theta) \cdot \lambda^B(n+1) + \theta \cdot \lambda^B(n-1) \quad \forall n \in \{-k+1, \dots, t-1\} \quad (4)$$

From (3), we have:

$$\lambda^A(n+1) - \frac{1}{\theta} \lambda^A(n) + \frac{1-\theta}{\theta} \lambda^A(n-1) = 0.$$

The generic solution to this second-order difference equation is of the form:

$$\lambda^A(n) = \kappa_1 \cdot r_1^n + \kappa_2 \cdot r_2^n,$$

where (κ_1, κ_2) are constants and (r_1, r_2) are the roots of the following second order equation:

$$x^2 - \frac{1}{\theta} x + \frac{1-\theta}{\theta} = 0.$$

Simple calculations yield:

$$r_1 = \frac{1-\theta}{\theta} \quad \text{and} \quad r_2 = 1$$

In order to determine the values of (κ_1, κ_2) , we use the fact that $\lambda^A(-k) = 0$ and $\lambda^A(t) = 1$:

$$\lambda^A(-k) = 0 \Rightarrow \kappa_1 \left(\frac{1-\theta}{\theta} \right)^{-k} + \kappa_2 = 0 \quad \text{and} \quad \lambda^A(t) = 1 \Rightarrow \kappa_1 \left(\frac{1-\theta}{\theta} \right)^t + \kappa_2 = 1$$

Denoting $\Theta \equiv \frac{1-\theta}{\theta}$, one can easily see that $\kappa_1 = \frac{1}{\Theta^t - \Theta^{-k}}$ and $\kappa_2 = -\frac{\Theta^{-k}}{\Theta^t - \Theta^{-k}}$, and therefore the general solution is:

$$\lambda^A(n) = \frac{1 - \Theta^{n+k}}{1 - \Theta^{t+k}} \quad \forall n \in \{-k, \dots, t\} \quad (5)$$

Note from (3) and (4) that the case $s = B$ is obtained simply by switching θ and $1 - \theta$:

$$\lambda^B(n) = \frac{1 - (1/\Theta)^{n+k}}{1 - (1/\Theta)^{t+k}} \Leftrightarrow \lambda^B(n) = \frac{\Theta^{t-n} - \Theta^{t+k}}{1 - \Theta^{t+k}} \quad (6)$$

Obviously, $p \equiv \Pr(A | 0)$. From the definitions of $\Pr(A | n)$, p_H and p_L , and treating k and t as real numbers (see Assumption 2) we have:

$$p_H \equiv \Pr(A | +t) = \frac{1}{1 + \left(\frac{1-\theta}{\theta}\right)^t \frac{1-p}{p}} \Leftrightarrow \Theta^t = \frac{p}{1-p} \frac{1-p_H}{p_H} \quad (7)$$

$$p_L \equiv \Pr(A | -k) = \frac{1}{1 + \left(\frac{1-\theta}{\theta}\right)^{-k} \frac{1-p}{p}} \Leftrightarrow \Theta^{-k} = \frac{p}{1-p} \frac{1-p_L}{p_L} \quad (8)$$

Therefore, combining (5), (6), (7), and (8), we get:

$$\lambda^A(0) = \frac{p - p_L}{p_H - p_L} \times \frac{p_H}{p} \left(= q^A(p; p_L, p_H) \right), \quad \lambda^B(0) = \frac{p - p_L}{p_H - p_L} \times \frac{1 - p_H}{1 - p} \left(= q^B(p; p_L, p_H) \right),$$

$$\text{and } q(p; p_L, p_H) \equiv p \times \lambda^A(0) + (1 - p) \times \lambda^B(0) = \frac{p - p_L}{p_H - p_L}.$$

Appendix A2. Proof of Proposition 1

According to the payoff structure of l given by (2), it is obvious that l will never provide information whenever $\mu \geq \pi_a$. Furthermore, given Assumption 2, extra information cannot hurt l if $\mu \in (\pi_b, \pi_a)$ or if $\mu < \pi_b$. The only issue left is then whether l will stop if $\mu = \pi_b$ or continue until either $\mu \rightarrow 0$ or $\mu = \pi_a$.²¹ When $\mu = \pi_b$, the leader prefers to stop information and get a payoff of y for sure rather than continue and get either 0 or x if and only if:

$$y > x \times q(\pi_b; 0, \pi_a) + 0 \times \left[1 - q(\pi_b; 0, \pi_a) \right] \Leftrightarrow x/y < \pi_a/\pi_b.$$

For each of these cases, the utility of the leader is:

$$U_{(0, \pi_a)} = x \times q(p; 0, \pi_a) \quad \text{and} \quad U_{(\pi_b, \pi_a)} = x \times q(p; \pi_b, \pi_a) + y \times \left(1 - q(p; \pi_b, \pi_a) \right),$$

and that of the decision-maker is:

$$V_{(\tau, \pi_a)} = p \times \frac{1}{\pi_a} \times q^A(p; \tau, \pi_a) + (1 - p) \times \frac{1}{1 - \pi_b} \times \left(1 - q^B(p; \tau, \pi_a) \right)$$

where $\tau \in \{0, \pi_b\}$. Simple algebra gives the final outcome.

²¹Note that because there is symmetric information, d cannot infer from the decision of l to stop providing information anything about the true state of the world. It is therefore a dominant strategy for l to play according to this rule.

Appendix A3: proof of Lemma 2

Given Assumption 3, the decision-maker's payoff when she decides to pay the cost of becoming perfectly informed is:

$$V_L = \frac{1}{\pi_a} - c, \quad (9)$$

By contrast, the payoff of not getting informed depends on the current posterior belief μ (which determines the action to be taken). We have:

$$V_N(\mu) = \begin{cases} \mu/\pi_a & \text{if } \mu \geq \pi_a \\ 1 & \text{if } \mu \in (1 - \pi_a, \pi_a) \\ (1 - \mu)/\pi_a & \text{if } \mu \leq 1 - \pi_a \end{cases} \quad (10)$$

From (9) and (10) we get that:

$$V_L > V_N(\mu) \Leftrightarrow \begin{cases} \mu < 1 - \pi_a c & \text{if } \mu \geq \pi_a \\ c < (1 - \pi_a)/\pi_a & \text{if } \mu \in (1 - \pi_a, \pi_a) \\ \mu > \pi_a c & \text{if } \mu \leq 1 - \pi_a \end{cases}$$

Note that $\mu < 1 - \pi_a c$ and $\mu > \pi_a$ are compatible if and only if $c < (1 - \pi_a)/\pi_a$. Similarly, $\mu > \pi_a c$ and $\mu < 1 - \pi_a$ are also compatible if and only if $c < (1 - \pi_a)/\pi_a$.

Appendix A4: proof of Proposition 3

If $c < (1 - \pi_a)/\pi_a$, then $\pi_a c < 1 - \pi_a$ and $1 - \pi_a c > \pi_a$. Hence, the leader will never provide information if $\mu \geq 1 - \pi_a c$ (which guarantees a payoff of x) and he will always give further information if $\mu < \pi_a c$ (because otherwise he gets a payoff of 0).

When $\mu \in (\pi_a c, 1 - \pi_a c)$, the leader has two options. First, to keep providing news until $p_L = 0$ or $p_H = 1 - \pi_a c$. This implies a payoff $x \times q(\mu; 0, 1 - \pi_a c) = x \frac{\mu}{1 - \pi_a c}$. Second, to stop providing news. In this case learning is re-initialized by the decision-maker and the leader's payoff is $x \times q(\mu; 0, 1) = x\mu$. Since the first alternative is dominant, the overall optimal stopping rule is $p_L = 0$ and $p_H = 1 - \pi_a c$. For all $p \in (\pi_a c, 1 - \pi_a c)$, the utility of the leader and the decision-maker are then given by:

$$U^c = x \times q(p; 0, 1 - \pi_a c) = x \frac{p}{1 - \pi_a c}$$

$$V^c = p \times q^A(p; 0, 1 - \pi_a c) \frac{1}{\pi_a} + (1 - p) \times \left(1 - q^B(p; 0, 1 - \pi_a c)\right) \frac{1}{\pi_a} = \frac{1 - \pi_a c(1 + p)}{\pi_a(1 - \pi_a c)}.$$

It follows that $\partial U^c / \partial c > 0$ and $\partial V^c / \partial c < 0$.