

A Theory of Influence *

Isabelle Brocas

USC and CEPR

Juan D. Carrillo

USC and CEPR

PRELIMINARY AND INCOMPLETE

Abstract

We analyze a game where one agent (the “leader”) decides first how much free information he collects, and then how much of this information he transmits to another agent (the “decision-maker”). Conditional on the information disclosed, the decision-maker undertakes an action that affects the payoff of both agents. We assume a conflict of preferences between the two agents and rely on previous research to argue that, in this context and with verifiable information, the leader will not obtain rents from his *possession of private information*. We then show that the leader will nonetheless obtain some rents due to his ability to decide whether to collect or forego information, that is due to his *control of the generation of information*. The paper provides an analytical characterization of these rents. We interpret them as the degree of *influence* that, e.g., mass media and defending counsels can exert on citizens and juries thanks to their discretion in the selection of news covered and their ability to proceed or stop interrogations, respectively.

*Correspondence address: Isabelle Brocas or Juan D. Carrillo, Department of Economics, University of Southern California, 3620 S. Vermont Ave., Los Angeles, CA - 90089, e-mail: <brocas@usc.edu> or <juandc@usc.edu>. We thank seminar participants at Columbia, Rochester, Michigan, Duke, UCLA and Seattle for helpful comments.

1 Motivation

Can I persuade an individual whose objectives are different from mine to take the action I like most? Do I need to manipulate that person's beliefs or can I affect his choice even if he is rational? The goal of this paper is to study different ways in which an individual can influence other people's decision to his own advantage. One possibility builds on the (by now classical) paradigm of agency theory with asymmetric information, whereby agents make a strategic use of private information in their interpersonal relations.¹ Another route is to study the choice of an individual to collect or forego evidence, anticipating a future strategic use. This second alternative, which to our knowledge has never been explored in the economics literature, is at the heart of this paper. It shows that an agent can derive rents not from his *possession of private information* as is already well-known, but rather from his mere ability to *control the generation of information*. Indeed, 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 such 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 decides first whether to acquire and then whether to reveal information can be of very different nature. One example is media coverage. Obviously, the media decides if the news gathered are presented to the citizens or not. However, a more basic and powerful tool to influence the public opinion is the discretion in the selection of the events that are covered and those that are neglected. Consider for example the strategy of a newspaper with some monopoly power and a vested interest in a specific issue. If the prevailing public opinion is sufficiently congruent with its own interest, it will be optimal to stop acquiring and providing further information. By contrast, when the newspaper is currently dissatisfied with the public opinion, collecting and revealing information with the hope that it will change their mind in the desired direction is the best option. A similar analysis can be applied to trials, where the defending counsel can influence the jury's final appraisal with his decision to reveal his information but also with his ability to proceed or stop an interrogation depending on the turn it takes.

In order to study this type of game, we present a model in which one agent (the "leader", he) decides first whether to collect free information and then whether to transmit the acquired

¹The effect of private information on the payoffs of all agents in the economy has been extensively studied in many different games (adverse selection, signaling, cheap-talk, etc.) under very different contexts (collusion, dynamics and renegotiation, multi-agent, common agency, etc.) and with applications to a wide range of situations (employer/employee, buyer/seller, regulator/firm, etc.) since the 1970s.

evidence to another agent (the “decision-maker”, she) who then takes an action that affects the payoff of both individuals. Since the preferences of both agents are not congruent, the leader will use to his own advantage his privileged position in the acquisition and transmission of information. If we assume that the information transmitted is verifiable and that the decision-maker knows the payoff structure of the leader, then we can rely on previous research to determine the outcome of the second stage of the game. In fact, in a seminal paper, Milgrom and Roberts (1986) show that the leader will not be able to make a strategic use of his private information. The basic idea is that the decision-maker will adopt a skeptical, no-news-is-bad-news position. That is, any information transmitted is verified and any information not transmitted is interpreted as negative for her interests. As a result, the leader in our game will not be able to get rents by hiding some of the acquired information; formally, it is “as if” the news collected were automatically shared with the decision-maker.

With these premises in mind, we then move to the main goal of the paper: to determine the degree of influence that can be exerted with the ability to acquire or forego information. Anticipating that no evidence can be strategically hidden, the incentives of the leader to keep collecting information depend on the likelihood that new information will move the belief of the decision-maker further in the “right direction” (that is in the direction where she will take the action preferred by the leader) vs. further in the “wrong direction.” The paper characterizes analytically the optimal stopping rule for the generation of information and the rents obtained in equilibrium by the leader (Proposition 1). Naturally, the extent of this process of influence will crucially depend on the degree of conflict of interests between the leader and the decision-maker: the smaller the conflict, the easier the ability to induce the decision-maker to undertake the action that is optimal from the leader’s perspective. Yet, our comparative statics suggest that the power of influence with this procedure can be quite important (Proposition 2). In a second stage, we analyze the case in which the decision-maker can also acquire information, although at a positive cost. The capacity of the decision-maker to obtain information by herself induces the leader to provide news up to the point where she does not have an incentive to incur the cost of re-initiating learning. Overall, introducing this possibility reduces the leverage of the leader on the decision-maker’s choice but it does not eliminate it (Proposition 3). Naturally, the amount of influence will be proportional to the decision-maker’s cost of acquiring information.

Before presenting the model, we would like to mention some papers related to ours. The optimality of ignorance in multi-agent contexts has been analyzed in the incentives literature. Crémer (1995) shows that a principal may optimally choose an inefficient monitoring

technology. This creates an ex-post inefficient situation because the principal will not know the reasons for a poor performance of its agent. However, such commitment to ignorance is ex-ante optimal because it alleviates a moral hazard problem. Closer to our work are Crémer and Khalil (1992) and Crémer, Khalil and Rochet (1998). These papers study the optimal contract designed by a principal when the agent can spend resources to acquire information before signing the contract and before being offered the contract, respectively. In both papers, the trade-off from the agent’s perspective is a costly acquisition of information vs. a private strategic use of such news.² The principal anticipates the incentives of the agent and designs the optimal contract accordingly. Our work differs in that we do not allow the two parties to sign a contract. More importantly, the collection of information by the leader is costless and, in equilibrium, no acquired information can successfully be withheld to the other party. His decision to forego evidence is based exclusively on the likelihood that news will move the beliefs of the other party towards the “desired” direction vs. the “undesired” one. Hence, even if all these papers share the result that, in equilibrium, an individual foregoes evidence and gets some rents out of that decision, the reasons are of very different nature. Last, the importance of the work by Milgrom and Roberts (1986) has already been discussed.

2 A model of influence

We consider the following game. One agent (the “leader”, he) decides whether to collect some information about the state of the economy and, conditional on its content, whether to make that information public. Based on the news released, 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 agents, the leader might use to his own advantage the ability both to collect and to transmit information.

We assume that any information transmitted is verifiable and that the two agents cannot sign a contract stipulating the amount of information to be collected and/or revealed by the leader. In this setting, there are two *potential* sources of rents for the leader. The first one comes from the leader’s ability to collect or refrain from acquiring new evidence; the second one is given by his capacity to communicate or hide any privately learned information. While the second source relates to the classical asymmetric information paradigm, the first source –which as we will see below constitutes the focus of this paper– builds on a new paradigm, namely the control of information flows.

²In particular, there is nothing interesting to say when the agent has free access to information.

2.1 Preliminaries

The model we use to analyze the above mentioned game is the following. There are two possible states in the economy, $s \in \{A, B\}$. Agents 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 collected and 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 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 decision-maker's posterior belief that the true state is A , conditional on the information collected and eventually transmitted by the leader. 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, given (2), the leader prefers that the decision-maker take action a rather than o and action o rather than b , independently of the true state of the economy. This conflict of interests is graphically represented in Figure 1.

[INSERT FIGURE 1 HERE]

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

Assumption 1 $1 - \beta < \alpha$.³

The structure of information acquisition is the simplest one that we can think of. At each moment in time, the leader decides whether to collect one extra signal $\sigma \in \{a', b'\}$ or not.⁴ 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

³We use the convention that in case of payoff-indifference the decision-maker takes the action most preferred by the leader. Note that if $\alpha < 1 - \beta$, then only actions a and b could be optimal. Last, $\alpha = \beta \in (1/2, 1)$ corresponds to the pure symmetric case, in which the payoff when the "correct" action from the decision-maker's perspective is taken (a if $s = A$ and b if $s = B$) is the same in both cases.

⁴For a model of optimal delegation in the acquisition of information, see Aghion and Tirole (1997).

perfectly informs the leader about the true state. We assume that collecting information is neither costly nor generates a delay. So, in particular, the leader can decide to obtain an infinite number of signals, in which case he learns the true state almost with certainty.⁵ Naturally, the leader updates his belief using Bayes rule, and decides whether to keep accumulating evidence depending on the realization of past signals.

Suppose for the time being that the information gathered by the leader is truthfully transmitted to the decision-maker. If a number n_a of signals a' and a number n_b of signals b' have been disclosed then, using standard statistical methods, it is possible to compute the posterior belief shared by the two agents:

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

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 $\mu(n) \equiv \Pr(A | n_a, n_b)$.⁶ 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 since we know that $n \in \mathbb{Z}$. We make this assumption only to avoid dealing with integer problems in the technical resolution of the model. As long as the amount of information of each individual signal is sufficiently small (technically, θ sufficiently close to $1/2$), this assumption implies little loss of generality.

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 $\mu \in [\pi_a, 1]$.

2.2 Examples

Before starting the formal analysis of the game, it will be useful to provide some concrete examples of the type of situations we have in mind. As described earlier, the main elements

⁵This assumption is clearly unrealistic: information is always costly. Yet, it allows us to isolate the net benefit from the leader's perspective of acquiring an extra piece of news. In Section 4.1 we extend the model to include a cost of acquiring information.

⁶Some properties of $\mu(n)$ are: (i) $\lim_{n \rightarrow -\infty} \mu(n) = 0$, (ii) $\lim_{n \rightarrow +\infty} \mu(n) = 1$, and (iii) $\mu(n+1) > \mu(n) \quad \forall n$.

of our model are: (i) a conflict of preferences between the two agents as modelled by (1) and (2), (ii) the ability of one and only one agent to both collect and transmit information to the other, and (iii) no possibility of contracting on the amount of information to be revealed. The examples presented below are only suggestive.

1. *Media.* Media professionals usually have an easy access to news but also a private interest, opinion or tendency in each issue (we often refer to a conservative newspaper, a liberal TV channel, etc.). The population is interested in knowing the truth but is forced to rely mostly on the information revealed by the media. Suppose that the action of a government can be either adequate (A) or inadequate (B). The public opinion about the government will be either good a , medium o or bad b , depending on their belief μ about the quality of the government's action. If the newspaper supports the government, it will try to influence the opinion towards a rather than o or towards o rather than b (independently of the actual quality of the decision) by strategically deciding whether to continue or stop both collecting and providing information. For example, depending on the turn of events following a decision concerning a foreign affair, the head of the newspaper will ask to the journalist working in that country whether to report the information gathered or not and also whether to keep accumulating news (i.e. whether to stay in the heat of the action or come back).⁷

2. *Trial.* A jury d has to decide between releasing (a), minor conviction (o) or major incarceration (b) of a suspect. The prisoner can be either innocent (A) or guilty (B) of a major offense. The payoff of the jury is highest when a guilty prisoner is convicted and an innocent one is released (b if B and a if A), as modelled by (1). By contrast, the defending counsel l only cares about minimizing the burden of his client independently of his culpability, as reflected by (2). The counsel has two different tools to influence the final appraisal of the jury. First, he will call only the witnesses whose statements are likely to favor his client. Second, when the prosecutor calls a witness unknown to him, he will proceed or stop with his interrogation depending on the information that is being revealed by that witness.

3. *Sales.* A client d decides whether to buy good a , good o or good b . She does not know their characteristics, only that either a or b is closest to her own preferences (state A or state B). A seller l receives commissions from the producers. The producer of good a pays the highest commissions and the producer of b the lowest ones. The seller can check out the characteristics of the goods to determine which one fits best the preferences of the client.

⁷Andrew Ang suggested to us the following related piece of anecdotal evidence. CNN has a special TV channel run exclusively in airports. When the AA plane crashed in New York on November 12, 2001, CNN simply shut-down this channel so that no TV information was available in US airports.

However, he will use his discretion both to investigate and to transmit this information in order to induce her to choose a rather than o and o rather than b .

Summing up, in all these cases, the leader (media, counsel, salesperson) will use to his own advantage two tools. First, his ability to acquire or stay away from information (i.e. the possibility to send the correspondent back, the ability to stop the interrogation, the capacity not to learn the characteristics of the product). Second, his ability to reveal or hide the information obtained. Naturally, the decision-maker (citizen, jury, client) anticipates this strategic behavior. This is why, knowing that the leader has some information that has not been released is fundamentally different from knowing that the leader has refrained from acquiring evidence. Also, in many contexts, several l -agents with opposite goals will compete to collect and provide information (newspaper and TV channel, defending counsel and public prosecutor, etc.). This issue is briefly explored in Section 4.2. Last, note that in the previous cases (specially examples 1 and 2), contracts between l and d specifying the amount of information to be revealed by the leader are simply not feasible.

2.3 Information transmission

As we have seen above, the leader in this game has to make two interrelated decisions. First how much information to acquire, and second how much information to transmit.

We first start with the decision to transmit information. Suppose that the leader stops acquiring pieces of information when his posterior belief is $\mu \in [0, 1]$. He then provides a report to the decision-maker that will be denoted by $r(\mu) \subset [0, 1]$. Given our information verifiability assumption, we necessarily have the constraint $\mu \in r(\mu)$. It turns out that the sequential equilibrium of the information transmission part of this game has already been studied in a similar context by Milgrom and Roberts (1986). The conclusion reached by the authors can be restated as follows.⁸

Lemma 1 (Milgrom and Roberts, 1986 - Propositions 1 and 2)

Consider the continuation game where the leader stops acquiring evidence at a posterior μ . At every sequential equilibrium of this game, the decision-maker takes action b if $\mu \in [0, \pi_b)$, action o if $\mu \in [\pi_b, \pi_a)$ and action a if $\mu \in [\pi_a, 1]$.

In words, Lemma 1 states that the leader cannot influence the decision-maker by strategically manipulating the transmission of information. The idea is simple and intuitive. Since

⁸The proof follows directly from that of their paper.

the decision-maker knows the incentives of the leader to overstate the likelihood of state A , she will adopt what Milgrom and Roberts coin as a “skeptical position”: always assume the worst, which in our case corresponds to the lower bound of the report set.⁹ Under these circumstances, the leader will never include a posterior below his own belief. Furthermore, he will not be able to instil on the decision-maker a posterior above his own belief, even if he tries to. Note also that the equilibrium in this continuation game is not unique: if the leader’s belief is μ , then we can only say 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 just μ), and so are the equilibrium payoffs of the two agents.

How realistic is this idea that the leader cannot influence the decision-maker’s choice? An obvious way to restore some benefits of private information is to modify the ingredients of our model. For instance, we could assume that information is not verifiable, that the decision-maker is not sophisticated and anticipates only partially the incentives of the leader to misreport information, that she has imperfect knowledge of the leader’s payoff, etc.

However, instead of studying these variations, we want to concentrate on a fundamentally different way to persuade the decision-maker. Given Lemma 1, the leader anticipates that the information acquired will never remain private. Thus, a way to affect the decision-maker’s choice is simply to avoid information. In other words, we now focus on the paradigm of controlling the flow of information and, more specifically, on strategic ignorance as a way to influence the behavior of others.

2.4 Information acquisition

For the rest of the paper, we will say that any information acquired by the leader is automatically transmitted to the decision-maker. Naturally, this is formally equivalent to say that the leader chooses how much verifiable information is transmitted but that, in equilibrium, he cannot mislead her (as proved in Lemma 1).

A first step to characterize the rents obtained due to the 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 “a belief p ”). Suppose also that he stops acquiring information when he reaches a belief \bar{p} ($> p$) or a belief \underline{p} ($< p$). What is the probability of reaching one posterior before the other? Naturally, this

⁹So, if for example the leader reports a set $[\mu_1, \mu_2] \subset [0, 1]$, the decision-maker acts as if the report was α_1 .

will crucially depend on whether the true state is A or B . Formally, denote by $q^s(p; \underline{p}, \bar{p})$ the probability of reaching \bar{p} before reaching \underline{p} when the initial belief is $p \in (\underline{p}, \bar{p})$ and the true state is s . Similarly, $q(p; \underline{p}, \bar{p})$ is the unconditional probability of reaching \bar{p} before \underline{p} (i.e. the probability given the existing uncertainty about the true state of the world). By definition, we have $q^s(\underline{p}; \underline{p}, \bar{p}) = 0$ and $q^s(\bar{p}; \underline{p}, \bar{p}) = 1$ 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 2 and they are key for our entire analysis.

Lemma 2 $q^A(p; \underline{p}, \bar{p}) = \frac{p - \underline{p}}{\bar{p} - \underline{p}} \times \frac{\bar{p}}{p}$ and $q^B(p; \underline{p}, \bar{p}) = \frac{p - \underline{p}}{\bar{p} - \underline{p}} \times \frac{1 - \bar{p}}{1 - p}$.

Moreover, $q(p; \underline{p}, \bar{p}) \equiv p \times q^A(p; \underline{p}, \bar{p}) + (1 - p) \times q^B(p; \underline{p}, \bar{p}) = \frac{p - \underline{p}}{\bar{p} - \underline{p}}$.

Proof. See Appendix A1. □

Basically, Lemma 2 states that the probability of reaching a posterior upper bound \bar{p} before a posterior lower bound \underline{p} is proportional to the distance between the upper bound and the prior $(\bar{p} - p)$ relative to the distance between the prior and the lower bound $(p - \underline{p})$. For our model, it means that as the decision-maker's payoff of 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 between the prior and the posterior above which the decision-maker takes action a shrinks $(\alpha - p)$ so, other things equal, she is more likely to end up with a posterior belief in which that action is the optimal one (the opposite is true with β and b).

Note also that $q^A(\cdot) > q(\cdot) > q^B(\cdot)$ for all p, \underline{p} and \bar{p} . This is simply because, 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 and b' signals move it downwards, then for any starting belief p , it is more likely to reach an upper bound \bar{p} before a lower bound \underline{p} if the state is A than if the state is B . Last, $q^A(p; 0, \bar{p}) \rightarrow 1$, and $q^B(p; \underline{p}, 1) \rightarrow 0$. In words, the agent can never believe almost 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 , q and q^B do not depend on θ the informational content of each signal.

This result is obtained by direct inspection of the analytical expressions derived in Lemma 2, 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 θ the informational content of each signal

affects the *speed* at which one of the posteriors is reached but not the *relative probabilities* of attaining each of them. Roughly speaking, a more accurate information implies higher chances of receiving the ‘correct’ signal (a' if $s = A$ and b' if $s = B$) but also that an ‘incorrect’ signal will move the posterior farther away in the opposite direction. These two effects perfectly compensate each other.¹⁰

3 The optimal control of information generation

3.1 Influence through ignorance

Given the conflict of preferences between the two actors (see (1) and (2)), the leader will use his control of the flow of information to “induce” or “persuade” the decision-maker to undertake action a rather than o and action o rather than b . As information cannot be kept secret, 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 $(\underline{p}, \bar{p}) \in (0, 1)^2$ such that l does not collect extra information whenever one of these posterior beliefs is reached.¹¹ Using Lemma 2, we get our main result of the paper.

Proposition 1 *Suppose that $p \in (\pi_b, \pi_a)$. Two cases are possible:*¹²

- (i) *If $y/x < \pi_b/\pi_a$, then $\underline{p} = 0$ and $\bar{p} = \pi_a$. 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 $\underline{p} = \pi_b$ and $\bar{p} = \pi_a$. 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. □

¹⁰We have two remarks on this result. First, integer problems are important only when θ is large because then, after just one signal, the posterior is close to 0 or to 1. We do not take this case into account and simply assume that the information provided by each individual signal is small enough for Assumption 2 not to be too restrictive (that is, θ close enough to $1/2$). Second, if it were costly for the leader to obtain signals, then his willingness 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 or the other would still remain unaffected.

¹¹Again by abuse of notation, we will refer to $\underline{p} = 0$ (resp. $\bar{p} = 1$) the case in which the lower (resp. upper) bound is arbitrarily close to 0 (resp. 1), that is when the leader stops the collection of evidence only if he knows almost with certainty that the true state is B (resp. A).

¹²For the sake of completeness we have the following. If $p > \pi_a$ then $\underline{p} = \bar{p} = p$; no information is ever provided, the decision-maker takes always action a and the leader gets utility $U = x$. If $p < \pi_b$ and $x > \pi_a/\pi_b$, then $\underline{p} = 0$ and $\bar{p} = \pi_a$ and the leader gets utility $U = x \times p/\pi_a$. If $p < \pi_b$ and $x < \pi_a/\pi_b$, then $\underline{p} = 0$ and $\bar{p} = \pi_b$ and the leader gets utility $U = y \times p/\pi_b$.

As Proposition 1 shows, the leader derives rents from his ability to control the generation of 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 acquire pieces of news up to the point of learning with almost certainty whether the true state is A or B (formally, $\underline{p} = 0$ and $\bar{p} = 1$). Her expected payoff and that of the leader would be respectively $V_{(0,1)} = \frac{p}{\pi_a} + \frac{1-p}{1-\pi_b}$ (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 $U_{(\pi_b,\pi_a)}$).

The keys to determine the leader's optimal stopping rule are the following. First of all, 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 providing evidence if $\mu \in (0, \pi_b)$ or if $\mu \in (\pi_b, \pi_a)$: in the first case there is still a chance of hitting π_b and at least induce d to take action o rather than b , and in the second case the leader can at least wait until either π_b or π_a is hit. So, the only remaining question is whether to stop at π_b and obtain a payoff y with certainty or keep providing evidence. In the latter case, π_a is hit before 0 with probability π_b/π_a (and the leader obtains a payoff x) and 0 is hit before π_a with probability $1 - \pi_b/\pi_a$ (and the leader obtains a payoff of 0).

It is essential to notice that even if the decision-maker knows that the leader is going to control the generation of information to his own advantage, she is never worse-off by accepting the news presented to her. In other words, the decision-maker can only decrease her payoff if she decides to refuse the information offered by the leader (in that case, and given that $p \in (\pi_b, \pi_a)$, she would take action o and get a payoff of 1).¹³ One can also see that the more information is provided, the greater is the expected payoff of the decision-maker; in our simple game, she gets a payoff greater than 1 (the payoff when no information is transmitted) only if the leader does not stop the acquisition of news when $\mu = \pi_b$.

Overall, the leader cannot affect the decision-maker's choice by strategically refraining from transmitting the information acquired (see Lemma 1). However, the mere control of the flow of news gives him a substantial power, which is used to "induce", "persuade", or "influence" her actions (see Proposition 1). In terms of our examples, it means that a newspaper can induce a public opinion favorable to its own ideas more frequently than it would be the case in the absence of an ideology bias and an active control of the acquisition of news (that is, a conscious decision on whether to keep informing on certain issues or rather

¹³Naturally, this is also true when $p \notin (\pi_b, \pi_a)$.

stop collecting evidence). Similarly, consider any probability p that a convict is innocent. The chances that the final interrogation conducted by the defending counsel will succeed in persuading the jury to release his client are greater than p , simply because he will keep asking questions or refrain from doing so depending on the information revealed by the witness. Last, a seller will succeed in placing the good that generates highest commissions more often than the preferences of his clients would require.

As the reader can notice, we have focused on the simplest conflict of preferences between our two actors to highlight in its crudest form the benefits of information control. Yet, the analytical characterization of the functions q , q^A and q^B is general enough to allow an application of this same information control principle to more complex settings like, for example, situations in which the leader's payoff also depends on the true state of the economy.

3.2 Comparative statics

Naturally, one might wonder whether the possibility of influencing the behavior of another agent can be of considerable importance or if it is always of second-order magnitude. Suppose that l 's difference in payoffs between actions a and o is big relative to the likelihood of hitting π_a starting at π_b (formally, $y/x < \pi_b/\pi_a$ so that $\underline{p} = 0$ and $\bar{p} = \pi_a (= \alpha)$). In that case, the decision-maker will never end up taking action b contrary to her best interest (formally, $\Pr(b | A) \equiv 1 - q^A(p; 0, \alpha) = 0$). However, the probability of inducing her to take the leader's optimal action a contrary to her own ex-post best interest is given by the following simple formula:

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

From this, we deduce for example that $q^B(1/2; 0, 3/4) = 1/3$ and $q^B(1/2; 0, 2/3) = 1/2$. Some other probabilities are represented in Figure 2.

[INSERT FIGURE 2 HERE]

One should realize that, for the leader to obtain rents from his control of the flow of information, the decision-maker must necessarily have a limited action space. For instance, if in our second example the jury decided between a continuum of possible sentences (even though the accused is either guilty or innocent of a major offense) and the utility of the counsel were linear in the severity of the sentence, then controlling the flow of news would not generate any extra benefit.¹⁴ However, we have two remarks on this. First, suppose that

¹⁴Stated differently, l gets rents because, for example, d does not make a distinction in her sentence between a suspect who is innocent with probability π_a and one who is innocent with probability 1: in both cases,

there is a cutoff in the probability of the prisoner’s being guilty below which the suspect is released and above which the sentence is proportional to the jury’s belief. Only as a result of this threshold, the counsel is already able to derive rents from his acquisition of news.¹⁵ Second, in most situations it is almost inevitable to make choices that fall in a limited set of categories.¹⁶ Yet, the greater the number of categories is, the smaller will be the leader’s ability to generate rents with the control of information.

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 those of the leader. In terms of the primitives of our model, 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, the defending counsel will be mostly interested in dealing with a jury who derives a much higher payoff from releasing an innocent than from convicting a guilty suspect ($1/\alpha \gg 1/\beta$).

Naturally, this type of decision-maker will not always exist. Often, her payoff will be either always highly sensitive or always weakly sensitive to a “correct” action, independently of which one is taken. That is, the jury will be either always strongly or always weakly concerned by making the right choice. Similarly, when listening to news provided by the media, individuals are either highly or mildly interested in adopting the objectively correct position (technically, the first case corresponds to both α and β low and the second one to both α and β high). In this context, does the leader prefer to face a decision-maker strongly concerned or weakly concerned by her action? 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.

$\gamma = a$.

¹⁵Technically, all we need for our theory is a non-linearity in the leader’s payoff. Restricting the decision-maker to three possible actions is the simplest but by no means the only way to achieve it.

¹⁶Although not modelled in this paper, one can also think of some simple reasons for which it might be desirable to restrict the number of choices: avoidance of ambiguous assessments, cost-savings, etc. Also, psychological evidence suggests that individuals use a limited number of categories to think about problems. For an interesting theoretical analysis of the implications of such cognitive bias, see Mullainathan (2001).

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

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

Remember that $\alpha (= \pi_a)$ 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 (resp. 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 (resp. harmful) for the leader.¹⁷ In other words, as the decision-maker becomes relatively less concerned by the quality of her choice, the leader has more leeway in the collection and provision of news because the posteriors where the flow of information is stopped (π_a and π_b) are farther away from the prior p . This increased freedom 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 freedom 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. Proposition 2 then suggests that a defending counsel should prefer a jury strongly concerned by the quality of the sentence only if they a priori sympathize with his client. Similarly, an ideological newspaper should prefer citizens who are easily influenced to take strong positions (in one direction or another) only if their initial beliefs are relatively congruent with its own interests.¹⁸

3.4 Discussion: is this “a rational approach to influence”? (to be done)

4 Extensions

4.1 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 acquire some news, although at a higher cost than the leader. For example, citizens can always acquire pieces of evidence 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 $c (> 0)$

¹⁷In the symmetric case $\alpha = \beta$, the leader prefers a strongly concerned decision-maker if and only if $p > 1/2$.

¹⁸The analysis assumes that d 's payoffs when she takes the “wrong” action (a if B and b if A) and the “neutral” action (o) are fixed and equal to 0 and 1, respectively. Hence, the comparative statics on α and β should be interpreted as the effect of changing the payoff difference between the correct and the wrong action *relative* to the payoff difference between the neutral and the wrong action.

(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:¹⁹

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

Given that the leader cannot succeed in strategically hiding information to the decision-maker, it is always optimal for the latter to wait for the former to stop providing pieces of news before deciding whether to pay the cost and become perfectly informed or not. Suppose that, for a given posterior μ , the leader stops the flow of information. The next Lemma characterizes the optimal continuation strategy of the decision-maker.

Lemma 3 *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. □

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 provision 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 (fixed) costs. By contrast, for intermediate beliefs (and, more precisely, when $\mu \in (\pi_a c, 1 - \pi_a c)$) news are extremely informative for optimal decision-making and therefore 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 that case, the leader knows that if he stops providing pieces of news when $\mu \in \{\pi_b, \pi_a\}$, as dictated by Proposition 1(ii), then the decision-maker will re-initialize the learning process and undertake her optimal action. The anticipation of this possibility induces the leader to modify his optimal information collection strategy.

¹⁹Technically, this all-or-nothing learning technology corresponds to the case in which d pays c for a signal $\sigma \in \{a'', b''\}$ such that $\Pr[a''|A] = \Pr[b''|B] = 1$ (i.e., there is perfect correlation between signal and state). This assumption is not necessary, but it greatly simplifies calculations. All the results also extend to asymmetric payoffs ($\pi_a \neq 1 - \pi_b$).

Proposition 3 *If $c < (1 - \pi_a)/\pi_a$, then $\underline{p} = 0$, $\bar{p} = 1 - \pi_a c$ and the decision-maker never re-initializes learning. The leader's utility is $U^c = x \times \frac{p}{1 - \pi_a c}$ ($< \min \{U_{(0, \pi_a)}, U_{(\pi_b, \pi_a)}\}$) and that of the decision-maker is $V^c = \frac{1 - \pi_a c(1 + p)}{\pi_a(1 - \pi_a c)}$ ($> \max \{V_{(0, \pi_a)}, V_{(\pi_b, \pi_a)}\}$).*

Proof. See Appendix A4. □

When the decision-maker has the ability to acquire some news and these are not excessively costly ($c < (1 - \pi_a)/\pi_a$), 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 the action 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 . This is obtained by setting $\underline{p} = 0$ and $\bar{p} = 1 - \pi_a c$.

Overall, the decision-maker's ability to collect information by herself reduces the leader's discretion in the acquisition and provision of news, and therefore his capacity to manipulate her choices. As the decision-maker's cost c to obtain news decreases, more information is disclosed, which implies that her expected welfare increases and that of the leader decreases (formally, $\partial V^c / \partial c < 0$ and $\partial U^c / \partial c > 0$). 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 $c > 0$). Similar qualitative results would apply if the costs for the leader to disclose news were positive but smaller than for the decision-maker.

Interestingly, in our model the leader makes sure that the decision-maker never re-initiates the acquisition of news in equilibrium. 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 those papers, the optimal contracts designed by the principal are such that collusion, renegotiation, and acquisition of information do not occur in equilibrium. However, just as in our paper, the mere possibility of engaging in these practices has an impact on 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. The most obvious method to obtain this result would be to introduce some uncertainty, like a

random and privately learned cost for the decision-maker or a random positive cost for the leader to collect news.²⁰

4.2 Competition to provide information (to be done)

5 Conclusion

The paper has argued 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 choice to acquire or forego extra evidence. While the first mechanisms builds on the classical asymmetric information paradigm and has been extensively studied in the literature, the second one has not received any attention. The goal of our paper was to provide a first careful analysis of the rents that can be extracted via that mechanism in a particular game.

Obviously, the nature of the conflict of preferences between the different actors will crucially affect the decision to acquire new evidence and the equilibrium rents of the agent with privileged access to news. One interesting alley for future research would be to look at situations where the optimal action from the leader's viewpoint also depends on the state of the economy. For instance, Business Schools want to know the teaching skills of their professors. Naturally, the companies that elaborate school rankings are also interested in this information. If the policy of a school is to make any possessed information available both to the students and the public in general, it may be preferable not to have teaching evaluations at all, even if it comes at the expense of not being informed about the quality of each professor. Note that this type of conflicts could be easily analyzed in the context of our model since we have determined not only the unconditional probability of reaching a given posterior before another (see $q(\cdot)$), but also the probability of reaching a given posterior before another conditional on the true state (see $q^A(\cdot)$ and $q^B(\cdot)$).

A second alley for future research would be to enrich the model by combining the paradigms of information acquisition and information transmission. To better focus on the decision to acquire news, our model has been designed in a way that no piece of information collected by the leader can be strategically hidden to the decision-maker in equilibrium. In many situations however, not only the nature but even the amount of information is non-verifiable. Under these circumstances, private information can (and will) be used to one's

²⁰Note that a similar device is employed by Kofman and Lawarrée (1996) to show that collusion may be an equilibrium outcome in a model à la Tirole (1986).

own advantage. The question that naturally arises is how the ability to manipulate the transmission of information modifies the incentives to acquire news in a first place. In other words, if the counsel can obtain rents from his possession and non-disclosure of private information, will he still have incentives to forego free evidence or will he rather prefer to be fully informed?

References (to be completed)

1. Aghion, P. and J. Tirole (1997), "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), pp. 1-29.
2. Crémer, J. (1995), "Arm's Length Relationships," *Quarterly Journal of Economics*, 110(2), pp. 275-295.
3. Crémer, J. and F. Khalil (1992), "Gathering Information Before Signing a Contract," *American Economic Review*, 82(3), pp. 566-578.
4. Crémer, J., Khalil, F. and J.C. Rochet (1998), "Strategic Information Gathering Before a Contract is Offered," *Journal of Economic Theory*, 81(1), pp. 166-200.
5. Dewatripont, M. (1988), "Commitment through Renegotiation-Proof Contracts with Third Parties," *Review of Economic Studies*, 55(3), pp. 377-389.
6. Kofman and Lawarrée (1996), "On the Optimality of Allowing Collusion," *Journal of Public Economics*, 61(3), pp. 383-407.
7. Milgrom, P. and J. Roberts (1986), "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 17, pp. 18-32.
8. Mullainathan, S. (2001), "Thinking through Categories," *mimeo*, MIT.
9. 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 2

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 $\Pr(A \mid +t) = \pi_a (> p)$ and that the posterior belief when the difference of signals reaches $n = -k$ is $\Pr(A \mid -k) = \pi_b (< p)$. 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, $\lambda^s(0) \equiv q^s(p; \pi_b, \pi_a)$. 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 $1 - \theta$:

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

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

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

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

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

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

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

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 $\mu \in (0, \pi_b) \cup (\pi_b, \pi_a)$. The only issue left is then whether l will stop if $\mu = \pi_b$ or continue until $\mu \in \{0, \pi_a\}$.²¹ When $\mu = \pi_b$, the leader prefers to stop information (and get a payoff of y for sure) rather than continue until hitting $\mu \in \{0, \pi_a\}$ (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:

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

and that of the decision-maker is:

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

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.

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

Appendix A3: proof of Lemma 3

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 given 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 $\underline{p} = 0$ or $\bar{p} = 1 - \pi_a c$. 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 $\underline{p} = 0$ and $\bar{p} = 1 - \pi_a c$. For all $p \in (\pi_a c, 1 - \pi_a c)$, the utility of the leader and the decision-maker are then given by:

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

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




Figure 1. Payoff of Decision-maker and Leader.




Figure 2. Probability that d takes action a when $s = B$.