AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Queen's Economics Department Working Paper No. 1321

Information Sharing and Incentives in Organizations

Jean-Etienne de Bettignies
Queen

Jan Zabojnik
Queen

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

12-2013

Information Sharing and Incentives in Organizations

Jean-Etienne de Bettignies and Jan Zabochnik^{*†}

This version: December 28, 2013

Abstract

We examine optimal information flows between a manager and a worker who is in charge of evaluating a parameter of interest, e.g. the value of a project. The manager may possess information about the parameter, and, if informed, may divulge her information to the worker. We show that information sharing may weaken the worker's incentives and that, consequently, the manager may find it optimal to conceal her information from the worker. Moreover, the manager faces a time-inconsistency problem, which leads her to conceal her information more often than she would if she could commit to an information sharing policy. We build on these results to address issues related to authority in organizations.

Keywords: Information non-disclosure, expert evaluation, agency costs, authority.

JEL classifications: D21, D82, L23.

^{*}A previous version of this paper was circulated under the title “On the Diffusion of Information in Organizations.” We thank Ricardo Alonso, Jim Brander, Tom Davidoff, Anne Duchêne, April Franco, Alberto Galasso, Denis Gromb, Thomas Hellmann, Ig Hortsman, Thor Koepl, Jeff McGill, Matt Mitchell, Mikhail Nediak, Tom Ross, Nate Schiff, Kostas Serfes, Peter Thompson, Mihkel Tombak, Ralph Winter, and seminar participants at Drexel University, Queen's University, University of British Columbia, and University of Toronto, for helpful comments. Bulat Gainullin and Jonathan Lee provided excellent research assistance.

[†]Queen's University, Kingston, ON, K7L 3N6, Canada. Email: jdebettignies@business.queensu.ca; zabochnik@econ.queensu.ca

1 Introduction

Management scientists have long recognized that one of the main challenges faced by any organization is to ensure that, whenever needed, knowledge flows from employees who generate it to all other employees for whom it may be production-relevant. In the economics literature on the topic, information sharing is usually viewed as desirable, and the research has focused on how incentives can be structured to promote information sharing (e.g., Snyder and Levitt, 1997) or on what the consequences are of imperfect information sharing (e.g., Dessein, 2002).

We show that information sharing has a downside because it can be detrimental to employees' incentives to generate additional information. We develop the idea in a setting where a manager must decide whether or not to undertake a project, which could be an intended merger, the launch of a new product, entry into a new market, etc. Before launching the project, the firm needs to undertake some preliminary preparations and investments, which are guided by the manager's information about the available project and about the economic environment that will affect its profitability, say, her belief about the nature of synergies created by a merger. We capture all this information through the manager's estimate of a single parameter, η , that can be thought of as the value of the project.

The manager may or may not have observed a private signal about this value. In either case, she tasks a subordinate (worker) with collecting additional information about the project, to further improve her estimate. Information collection requires effort and the more effort the worker puts in, the more accurate is the information he collects. But effort is costly and difficult to measure, which leads to agency concerns. We show that these agency concerns are exacerbated when the manager shares her information with the worker; in fact, even the very possibility that the manager is informed dampens the worker's incentives, because managerial information reduces the marginal impact of the worker's effort on the probability that the project goes ahead, and in turn on his expected payoff.

On the other hand, learning the manager's information makes the worker more productive (which we model as a lower marginal cost of effort) because it allows him to better direct his search for new information and to avoid wasting his effort on rediscovering information already known to the manager. The central question we address in this setting is whether the manager should share her own information with the worker before the latter decides on how much effort

to devote to the task.

If the worker's effort were verifiable, resolving the above trade-off would be easy. The manager would always share her information with the worker in order to increase his productivity and then simply tell the worker how much effort to exert, paying him just enough to compensate him for his cost of effort. However, when effort is not contractible and the worker is protected by limited liability, this solution is not feasible. The worker's incentives to provide effort then depend on the likelihood the project will be undertaken and on the reward he receives if the project goes ahead. For the most part of the paper, we focus on the simpler case in which the worker's reward takes the form of a private benefit from a project that is implemented, but we also demonstrate that our main results go through if we allow for the reward to take the form of a bonus in a contingent contract.

We assume that the project is undertaken – and the worker receives his private benefit or his bonus – only if the firm's preliminary preparations and investments were adequate, that is, if the manager's estimate was sufficiently close to the true value of the project, η , which is revealed shortly before the firm has to decide whether or not the project will be undertaken. To illustrate with a concrete example, think of a pharmaceutical firm developing a new drug. Before the firm can market it, the drug must be approved by the FDA, for which the firm has to demonstrate through clinical trials that the drug is safe and effective in its proposed use. Here, η represents the characteristics (say, the chemical composition) of a perfectly safe and effective drug, whereas the firm's estimate of η represents its belief about how to design a drug that is safe and effective. The clinical trials then reveal whether the properties of the new drug are sufficiently close to its targeted properties. Specifically, if the drug passes the clinical trials and gets approved by the FDA, the firm knows that the new drug's properties are close to the properties of a perfectly safe and effective drug even if they don't learn from the trials the exact chemical composition of such a perfect drug.

We show that, holding fixed the worker's belief about the probability that the manager is informed, the informed manager's decision to share her information with the worker can either strengthen or weaken the worker's incentives, depending on how much the manager's information decreases the worker's marginal costs of collecting additional information. If the impact of the manager's information on the worker's marginal cost is sufficiently large, his incentives to collect more information are strengthened if the manager's information is divulged to him.

However, if the manager's information has relatively little effect on the worker's marginal cost, then the net effect of information sharing is to dampen the worker's incentives. Consequently, the informed manager may find it optimal to conceal her information from the worker and pretend that she is uninformed.

After establishing this basic result, we show that the inefficiency in information sharing that we uncover is at least partly due to a time-inconsistency problem faced by the manager. This time-inconsistency problem stems from an externality that the informed type of the manager imposes upon the uninformed type: If the worker knew for sure that the manager did not receive a private signal about the project, he would have a strong incentive to exert effort to collect information about the project. However, the worker understands that the informed manager has a tendency to withhold information from him for fear of negatively affecting the worker's incentives. He therefore holds back on his effort even if the manager claims to be uninformed. But the informed manager does not internalize this effect when deciding whether to share her information with the worker. We show that if the manager could commit *ex ante* to an information sharing strategy, she would commit to share her information with the worker more often. In practice, such a commitment may take the form of the "open book" management approach adopted by some firms, which consists of disclosing to the firm's employees detailed operating information, such as its financial records and the sources of its profits (Davis, 1997).

Commitment improves efficiency, but we also show that even if the firm could commit to an information sharing policy, it would not always commit to the one that leads to maximum welfare. Specifically, when the *ex ante* probability that the manager will get to observe a private signal about the project is large, the externality explained above arises infrequently and the incentive benefits of hiding information from the worker prevail. In those cases, the manager does not want to commit to always reveal her information to the worker.

Building on these key results, we derive several other insights: First, we examine whether better informed managers share their information with subordinates more frequently. In our framework, there are two ways to measure how well a manager is informed. The first is through the precision of her signal. One might think that an increased accuracy of the manager's signal should favor information sharing, but we show that the opposite is true in our model: More precise information from the manager encourages the worker to further substitute information for effort, leading to a more imprecise final estimate. This makes divulging information relatively

less attractive and leads to less information sharing.

An alternative way to measure how well the manager is informed is through the probability that she will observe a private signal about the project. We show that using this measure leads to the opposite conclusion: A manager who is more likely to have information about the project is also more likely to share it with the worker. Moreover, this holds not only because information sharing is feasible only when the manager is informed but also because, once informed, the manager has a stronger incentive to share her information with the worker.

Second, we look at the connection between information sharing and authority. We introduce here the concept of *implicit* authority and relate it to the concept of real authority analyzed in Aghion and Tirole (1997). In our model, the manager always retains formal authority; moreover, by sharing her information with the worker the manager ensures that the final decision will be made solely based on the worker's report (which will combine the information shared by the manager with the new information he collected). So, divulging information to the worker gives him more real authority à la Aghion and Tirole (1997), in the sense that the manager more often rubberstamps the worker's report. However, we show that information sharing may lead to less delegation of *implicit* authority, which we define as the weight that the manager places on the worker's personal opinion in her decision whether or not to undertake the project. In particular, the worker's implicit authority declines with information sharing if divulging information reduces his effort, because a lower effort leads to a less precise personal estimate, which in turn reduces the weight placed on the worker's opinion in the manager's decision.

Related literature

Our paper is closely related to Prendergast's (1993) theory of "yes men:" In both papers, a principal enlists the help of an agent to obtain - through Bayesian updating - an estimate of some parameter of interest. However, in Prendergast's model, the worker is rewarded when his report is close enough to the *manager's estimate* and the main focus of the analysis is on the worker's incentive to bias his report towards the manager's estimate. In contrast, the worker in our model is rewarded if the manager's final estimate (based on the worker's report) is close enough to the project's true value and the main question we study is whether the manager will find it optimal to conceal her information from the worker. Thus, the two papers differ significantly in terms of both the incentives environment and the research question.

The issue of evaluation - information gathering - by the agent is central to the model, and

this theme is related to the growing literature on delegated expertise, which includes Lambert (1986), Demski and Sappington (1987), Core and Qian (2002), Gromb and Martimort (2007), and Malcomson (2009), among others. These models focus on environments in which an agent (or multiple agents) needs to be motivated to both collect information about available projects and to choose the project to be undertaken. In contrast, the agent in our model does not have formal decision-making authority with respect to the project. Furthermore, the literature on delegated expertise typically assumes that it is prohibitively costly for the agent to convey to the principal his information about the available projects, whereas communication between the principal and the agent is at the heart of our analysis.

In this respect, our paper is related to models in which an agent needs to communicate his private information to the principal. This literature includes the papers by Snyder and Levitt (1997) and Dessein (2002) that we have already mentioned, as well as Alonso and Matouschek (2007) and Alonso, Dessein and Matouschek (2008). The last two papers, like Dessein (2002), build on the cheap talk model of Crawford and Sobel (1982) to examine questions related to organizational design, such as the optimal delegation of authority and the choice between centralized and decentralized coordination. A key assumption underlying these models is that the sender's information is soft, i.e. the sender cannot certify his/her information. In our model, the sender's information (when divulged) can be perfectly verified by the receiver. More importantly, these papers are typically not concerned with the manager's incentive to communicate her information to the agent, which is of central importance to our theory.

One exception is Demski and Sappington (1986), who also model a principal with private information that can affect an agent's incentives. However, their principal never completely withholds her information from the agent, although she may find it optimal to delay its release when the information indicates that the agent's performance will be hard to measure.

Finally, the paper is also related to the line of research on delegation of authority in its various forms, e.g. informal (Baker, Gibbons and Murphy, 1999), formal (Zabojnik, 2002), and real (Aghion and Tirole, 1997) authority. The contribution here is dual: First, we introduce the concept of implicit authority and argue that in the context of evaluation exercises this concept may provide a more suitable measure of the agent's authority than the concept of real authority the literature has focused on following Aghion and Tirole (1997). Second, we show that by sharing her information with the worker, the manager can undermine the worker's

implicit authority, even though it may seem that the worker was given more “real authority.”

The paper proceeds as follows. In Section 2 we present the model and in Section 3 we establish the first-best benchmark. In Section 4, we analyze how the manager’s information affects the worker’s incentives, characterize the manager’s optimal strategy for sharing her information with the worker, and demonstrate that the manager faces a time-inconsistency problem. Section 5 provides comparative statics results with respect to the quality of the manager’s information, discusses the relationship between information sharing and authority, and extends the baseline model to allow for monetary incentives. In Section 6 we offer concluding remarks.

2 The Model

We consider an organization composed of a manager (she) and a worker/subordinate (he), both of them risk neutral. The manager’s goal is to estimate a parameter η . This parameter estimation could capture the evaluation of the true value of a project, as in the example used in the introduction, or more generally the evaluation of ideas, activities, employees, etc.

A key feature of the model is that the firm’s profits, gross of the cost of compensating the worker, are an increasing function of the accuracy of the estimate of the parameter. Specifically, the principal has access to a project which can be successful only if her posterior estimate of η is within the distance q of the true value of η : If the firm’s best estimate of η turns out to be from the interval $[\eta - q, \eta + q]$, the project is successful and yields a gross payoff to the firm of Π ; otherwise, the project is not undertaken and the firm’s payoff is zero.

We think of this specification, which is similar to the one used in Prendergast (1993), as a reduced-form way of capturing the idea that managerial decisions based on accurate information are superior to, and lead to higher profits than, decisions based on imprecise information. As mentioned in the introduction, we imagine η as representing information about the available project (merger, launch of a new product, entry into a new market, etc) and about all the factors that will affect its profitability. The project can be successful only if ahead of launching it the firm undertakes some preliminary preparations and investments, which are based on the manager’s estimate of η . Once η becomes known, the firm learns whether its preliminary preparations were adequate or not, which is modeled as the firm’s estimate of η being close to the true value of η . If the preliminary preparations were adequate (the distance between the

estimate and η is less than q), the firm goes ahead with the project, otherwise the project is abandoned.

Note that we do not require that the true value of η is observed by the firm; it is enough if the manager observes whether the true value of η is within the distance q of her estimate.

Priors and signals. At the beginning of the game, both the manager and the worker have the same prior $p_0 = N(\eta_0, 1/h_0)$ about the distribution of η , where h_0 is the precision of the prior belief.¹ After the worker is hired, with probability β the manager observes an imperfect signal of η , $\eta_m = \eta + \varepsilon_m$, with $\varepsilon_m \sim N(0, 1/h_m)$. With probability $1 - \beta$, she does not observe any signal. Both the signal and whether or not she has received it are the manager's private information, but the worker knows the distribution of ε_m .

At a later point in the game, the worker can also obtain a signal, denoted η_w , with $\eta_w = \eta + \varepsilon_w$ and $\varepsilon_w \sim N(0, 1/h_w)$. The precision of the worker's signal increases in the effort he provides: $h_w = h_w(e)$, with $h'_w(\cdot) > 0$ and $h''_w(\cdot) < 0$. Effort is costly, with the cost being equal to $\delta C(e)$, where $C'(\cdot) > 0$, $C''(\cdot) > 0$, and $C(0) = 0$, and $\delta \in \{t, 1\}$ denotes a parameter that depends on whether the manager shares her information with the worker. In particular, $\delta = 1$ if the manager does not share her information with the worker, whereas $\delta = t \in (0, 1)$ if she does. This formalizes the idea that by sharing her information with the worker, the manager can point him in the right direction and thus save him the effort of exploring avenues that are unlikely to yield precise information about the firm's project. An alternative interpretation is that information sharing prevents duplication of effort in that a worker who is uninformed about what exactly the principal knows could spend his effort uncovering information that is already contained in the principal's signal.

Information sharing and updating. If informed, the manager decides whether to share her information with the worker before the worker exerts his effort. The manager's information is assumed to be hard; that is, the worker can verify its veracity.² If the manager shares her information, the worker uses it to update his prior and then collects additional information.

¹For expositional convenience, we will work with precisions rather than variances, where given a variance σ^2 , the corresponding precision is defined by $h = 1/\sigma^2$.

²For example, true information may come with supporting documents that can be requested and examined by the worker, while creating similar documents for "fake" information may be prohibitively costly. This assumption simplifies the analysis because it eliminates the possibility for the manager to strategically provide erroneous information to the worker. It is also one of the distinctions between this model and cheap talk models, where information is assumed to be soft (as discussed in the introduction).

Subsequently, he reports to the manager his posterior belief, which he has formed by combining his own signal with the information shared by the manager and with his prior belief p_0 .^{3,4} We denote the manager's posterior belief in this case (which is equal to the worker's report) by p_S , where S stands for "sharing". Given the properties of Bayesian updating under normal distributions, p_S represents a normal distribution: $p_S = N(\eta_S, 1/h_S)$.

If the manager does not share her information with the worker, either strategically or because she is uninformed, the worker simply reports his posterior distribution $p_N = N(\eta_N, 1/h_N)$, derived using only the prior p_0 and his signal η_w . Here, the subscript N stands for "no sharing". If the manager is uninformed, p_N becomes her posterior belief. If she is informed, she combines the worker's report with her own signal to arrive at the posterior belief $p_P = N(\eta_P, 1/h_P)$, where the subscript P indicates that this is the manager's "private" posterior.

To sum up, the manager's best estimate of the parameter η is the mean $\bar{\eta}$ of her posterior belief $\bar{p} = N(\bar{\eta}, 1/\bar{h})$, where $(\bar{\eta}, 1/\bar{h}) \in \{(\eta_S, 1/h_S), (\eta_N, 1/h_N), (\eta_P, 1/h_P)\}$.

Let us point out here that even though the assumption that the manager shares either all or none of her information may appear stylized, the model actually does allow for partial information sharing, through a mixed strategy. We will see that such partial information sharing indeed arises in our model endogenously, as an equilibrium outcome.

Contracting and the worker's payoff. The worker's payoff consists of a private benefit, B , that he receives if the project is undertaken, otherwise his payoff is zero. In order to better isolate the effects of information sharing, we abstract from monetary transfers throughout the majority of the paper. This implies both that contingent contracts are not feasible and that the manager cannot extract from the worker his private benefits through an up-front payment. The non-contractibility assumption is relaxed in section 5.3, where we allow the worker's pay to depend on whether the project goes ahead. We show there that if the worker is protected by limited liability (so that agency concerns are not assumed away) the main results of our model continue to hold.

Timing of the game. At date 0, Nature draws the parameter η to be estimated, after which the worker is hired.

³Alternatively, the worker could report his own signal, which would then be used by the manager to update her belief. Given the nature of Bayesian updating, both reporting protocols yield the same posterior belief for the manager.

⁴As discussed below, the worker has no incentive to report her signal untruthfully.

At date 1, with probability β the manager receives her private signal, η_m , and then decides whether or not to divulge her information to the worker.

At date 2, the worker exerts effort e , which determines $h_w(e)$, and subsequently observes η_w .

At date 3, the worker reports his beliefs to the manager. The manager's posterior distribution (and her best estimate of η) is then determined, either by adopting the worker's report (if the manager was uninformed or if she shared her information with the worker), or by combining the worker's report with the manager's signal (if the manager did not share her information with the worker).

At date 4, the manager observes η , and the project either goes ahead (if $\bar{\eta} \in [\eta - q, \eta + q]$) or is abandoned (if $\bar{\eta} \notin [\eta - q, \eta + q]$). Conditional on the project going ahead, the worker receives his private benefit B .

3 The First-Best (Benchmark) Case

We start the analysis by examining a benchmark scenario in which both the worker's effort and the manager's information sharing decision are efficient, i.e., total surplus from the employment relationship is maximized. Because information sharing directly affects total surplus only through the worker's cost of effort, it should be clear that the first best requires that the manager always shares her information if she has observed a signal. Thus, we need to specify two efficient effort levels: Effort e_I , asked of the worker when the manager is informed, and effort e_U , asked of him when the manager is uninformed.

When the manager is uninformed, the total surplus is given by

$$prob(|\eta_N - \eta| < q) (\Pi + B) - C(e),$$

so that the first-best level of e_U , denoted e_U^{FB} , is given by the first-order condition

$$\frac{\partial prob(|\eta_N - \eta| < q)}{\partial h_N} \frac{\partial h_N}{\partial h_w} h'_w(e_U^{FB}) (\Pi + B) = C'(e_U^*). \quad (1)$$

When the manager has observed a signal, the total surplus is given by

$$prob(|\eta_S - \eta| < q) (\Pi + B) - tC(e),$$

which yields the first-best level of e_I , denoted e_I^{FB} , as the solution to the first-order condition

$$\frac{\partial \text{prob}(|\eta_S - \eta| < q)}{\partial h_S} \frac{\partial h_S}{\partial h_w} h'_w(e_I^{FB}) (\Pi + B) = tC'(e_I^*). \quad (2)$$

We summarize these observations in the following proposition:

Proposition 1 *In the first-best arrangement, the informed manager always divulges her information to the worker. The worker provides effort e_I^{FB} if the manager is informed and e_U^{FB} if she is uninformed, where e_U^{FB} and e_I^{FB} are given by expressions (1) and (2), respectively.*

4 Non-Verifiable Effort

4.1 The effects of information sharing on effort

In our baseline model, if $|\bar{\eta} - \eta| < q$ the worker receives his private benefit B , otherwise he receives zero. Observe that once the worker's effort is sunk, the worker shares with the principal the goal of making the principal's estimate as precise as possible. Consequently, he always has an incentive to report his posterior mean to the manager truthfully, as doing so maximizes the probability that the manager's posterior mean will fall within q of η . His report about the precision of his posterior belief has no impact on his payoff, so the worker reports this information truthfully as well (and in any case, the manager can infer $h_w(e)$ from the worker's equilibrium effort level). This leads to the following result, shown more formally in the appendix:

Lemma 1 *The worker always reports his posterior belief about the distribution of η truthfully.*

We now determine the equilibrium by backward induction: Suppose first that the manager has observed a signal and has shared it with the worker. Then at date 2 the worker exerts effort e_S^* such that:

$$e_S^* \in \arg \max_e \text{prob}(|\eta_S - \eta| < q) B - tC(e), \quad (3)$$

where

$$\eta_S = \frac{h_0\eta_0 + h_w(e)\eta_w + h_m\eta_m}{h_0 + h_w(e) + h_m} \quad (4)$$

is the mean of the worker's posterior belief regarding η .

Evidently, $\text{prob}(|\eta_S - \eta| < q)$ depends on the worker's effort: We show in the proof of Proposition 2 that

$$\text{prob}(|\eta_S - \eta| < q) = \frac{1}{\sqrt{2\pi}} \int_{-q\sqrt{h_S}}^{q\sqrt{h_S}} e^{-x^2/2} dx,$$

and that this probability is strictly increasing and concave in the precision h_S of the worker's posterior belief, which is given by

$$h_S = h_0 + h_w(e) + h_m. \quad (5)$$

As can be seen from (5), h_S is strictly increasing in the precision of the worker's personal observation, which in turn is positively affected by his effort. Thus, taking the first-order condition, e_S^* can be expressed as the solution to:

$$\frac{\partial \text{prob}(|\eta_S - \eta| < q)}{\partial h_S} \frac{\partial h_S}{\partial h_w} h'_w(e_S^*) B = tC'(e_S^*). \quad (6)$$

The strict concavity of $\text{prob}(|\eta_S - \eta| < q)$ and of $h(\cdot)$, together with the strict convexity of $tC(\cdot)$ imply a strictly concave program for the worker, and a unique solution e_S^* . Note that the strict concavity of $\text{prob}(|\eta_S - \eta| < q)$ with respect to h_S is not surprising: An increase in the precision of the worker's updated estimate increases his success probability, but since probabilities are by definition bounded at 1, the rate of increase in probability must fall as the precision rises.

Now suppose the manager has not communicated any information to the worker at date 1. This could be either because she has not observed any signal or because she has concealed her information from the worker. Lacking the manager's information, the mean and the precision of the worker's posterior belief regarding η are

$$\eta_N = \frac{h_0\eta_0 + h_w(e)\eta_w}{h_0 + h_w(e)} \quad (7)$$

and

$$h_N = h_0 + h_w(e). \quad (8)$$

The above also describes the manager's posterior belief when she was uninformed. When the manager is informed but strategically did not share her information with the worker, the

precision of her posterior belief is formed according to

$$h_P = h_0 + h_w(e) + h_m, \quad (9)$$

which is identical to (5), with the exception that $h_w(e)$ may be different due to the worker choosing a different level of effort.

Before choosing his effort, the worker makes a conjecture $\hat{\gamma}$ about the manager's strategy γ - where γ is the probability that the manager will divulge her information to the worker if she is informed (i.e. if she has received signal η_m). The worker then uses his conjecture $\hat{\gamma}$ to form a posterior belief $\hat{\beta}$ about the probability that the manager is informed given that she did not divulge information to him. Using Bayes' rule, we can express this belief as

$$\hat{\beta} = \frac{\beta(1 - \hat{\gamma})}{\beta(1 - \hat{\gamma}) + 1 - \beta}. \quad (10)$$

That is, the worker believes that once he reports his information to the manager, with probability $\hat{\beta}$ the manager will further aggregate it with her own signal to get a posterior distribution with precision h_P . With the residual probability $1 - \hat{\beta}$ the manager will have no private signal, so that the precision of her posterior will be h_N as given in 8.

The worker's expected payoff when no information was shared is thus

$$\left[\text{prob}(|\eta_P - \eta| < q) \hat{\beta} + \text{prob}(|\bar{\eta}_N - \eta| < q) (1 - \hat{\beta}) \right] B - C(e),$$

and he chooses e_N^* such that:

$$\left[\frac{\partial \text{prob}(|\eta_P - \eta| < q)}{\partial h_P} \frac{\partial h_P}{\partial h_w} \hat{\beta} + \frac{\partial \text{prob}(|\bar{\eta}_N - \eta| < q)}{\partial h_N} \frac{\partial h_N}{\partial h_w} (1 - \hat{\beta}) \right] h'_w(e_N^*) B = C'(e_N^*). \quad (11)$$

Note from (11) that e_N^* is strictly increasing in $\hat{\gamma}$. If the worker's belief that an informed manager will divulge her information increases, then the probability that the manager is informed given that she did not divulge her information must go down. And the lower is the probability that the manager is informed, the greater is the expected marginal product of the worker's effort and hence also his equilibrium effort. On the other hand, e_S^* does not depend on $\hat{\gamma}$, as is apparent from (6). This differing effect of the worker's belief about the manager's strategy on his effort will play an important role in determining the manager's optimal infor-

mation sharing strategy. As a first step towards deriving this optimal strategy, a comparison of (11) with (6) yields the following comparative statics result.

Proposition 2 *For any given $\hat{\gamma} \in (0, 1)$, there exists a $t^*(\hat{\gamma}) \in (0, 1)$ such that if $t > t^*(\hat{\gamma})$ the worker exerts more effort when the manager's information is concealed from him, whereas if $t < t^*(\hat{\gamma})$ he exerts more effort when the manager's information is divulged to him.*

Proof: Let p_i , $i \in \{S, N\}$, be the worker's posterior belief after observing his signal. As in the proof of Lemma 1, rewrite the worker's probability of success, Q_i , in terms of the standard normal distribution as

$$\begin{aligned} Q_S &\equiv \text{prob}(|\eta_S - \eta| < q) \\ &= \frac{1}{\sqrt{2\pi}} \int_{-q\sqrt{h_S}}^{q\sqrt{h_S}} e^{-z^2/2} dz \end{aligned}$$

and

$$Q_N = \hat{\beta} \frac{1}{\sqrt{2\pi}} \int_{-q\sqrt{h_P}}^{q\sqrt{h_P}} e^{-z^2/2} dz + (1 - \hat{\beta}) \frac{1}{\sqrt{2\pi}} \int_{-q\sqrt{h_N}}^{q\sqrt{h_N}} e^{-z^2/2} dz.$$

Using Leibnitz' rule, we can show that the term $T_j \equiv \frac{1}{\sqrt{2\pi}} \int_{-q\sqrt{h_j}}^{q\sqrt{h_j}} e^{-z^2/2} dz$, $j \in \{S, N, P\}$, is a strictly increasing and strictly concave function of the precision of the worker's posterior distribution:

$$\begin{aligned} \frac{\partial T_j}{\partial h_j} &= \frac{q}{\sqrt{2\pi h_j}} \exp\left(-\frac{q^2}{2} h_j\right) > 0, \\ \frac{\partial^2 T_j}{\partial h_j^2} &= -\frac{q}{2\sqrt{2\pi h_j}} \left(\frac{3}{h_j} + q^2\right) \exp\left(-\frac{q^2}{2} h_j\right) < 0. \end{aligned}$$

When the manager shares information with the worker, we have $Q_S = T_S$. Using $\partial h_j / \partial h_w = 1$, the worker's first-order condition can therefore be written as

$$\frac{\partial T_S}{\partial h_S} h'_w(e_S^*) B = t C'(e_S^*). \quad (12)$$

When no information is shared, the worker's expected payoff is

$$\left[\hat{\beta} T_P + (1 - \hat{\beta}) T_N \right] B - C(e),$$

with the first-order condition

$$\left[\hat{\beta} \frac{\partial T_P}{\partial h_P} h'_w(e_N^*) + (1 - \hat{\beta}) \frac{\partial T_N}{\partial h_N} h'_w(e_N^*) \right] B = C'(e_N^*). \quad (13)$$

Given $h'_j(e_i) > 0$ and given the strict concavity of T_j in h_j , in both cases the worker's objective function is strictly concave in e_i , so the above first-order conditions are also sufficient.

Next, observe that for any given level of effort $\hat{e}$, the precision of the worker's estimate is higher when the manager has shared with him information than when she has not: $h_S > h_0 + h_w(\hat{e}) + h_m = h_0 + h_w(\hat{e}) = h_N$. This, together with the strict concavity of T_j with respect to h_j , implies $\frac{\partial T_S}{\partial h_S} = \frac{\partial T_P}{\partial h_P} < \frac{\partial T_N}{\partial h_N}$ when $e = \hat{e}$. At any effort level, the worker's marginal benefit of an effort increase is therefore higher under no information sharing than under information sharing.

Now, from (12), e_S^* decreases in t . For any given $\hat{\gamma}$, as $t \rightarrow 0$, e_S^* converges to ∞ , whereas for $t = 1$ it must be $e_S^* < e_N^*$ because, as argued above, the LHS of (13) exceeds the LHS of (12) for any given effort level. By monotonicity and continuity of e_S^* in t , there must exist a $t^*(\hat{\gamma}) \in (0, 1)$ such that $e_S^*(t) < e_N^*$ if $t \in (t^*(\hat{\gamma}), 1]$ and $e_S^*(t) > e_N^*$ if $t < t^*(\hat{\gamma})$. $\square$

Proposition 2 is driven by the following tradeoff: When the manager divulges her information to the worker, she reveals to him that she has received an informative signal about the project, which increases the project's success for any level of effort the worker exerts. But at higher success levels the marginal impact of effort on the success probability falls: The marginal impact of an increase in the precision of the worker's signal on the precision of the principal's posterior belief is always the same, regardless of the information sharing arrangement ($\frac{\partial h_S}{\partial h_w} = \frac{\partial h_P}{\partial h_w} = \frac{\partial h_N}{\partial h_w} = 1$), while for any given level of effort, the precision of the worker's signal is higher when the manager has shared her information than when she has not. This, together with the strict concavity of the probability function, implies that the marginal impact of an increase in the precision of the worker's signal on the project's success probability is lower with information sharing than without, because for any given e , we have $\frac{\partial \text{prob}(|\eta_S - \eta| < q)}{\partial h_S} = \frac{\partial \text{prob}(|\eta_P - \eta| < q)}{\partial h_P} < \frac{\partial \text{prob}(|\eta_N - \eta| < q)}{\partial h_N}$. A comparison of the left hand sides of (11) and (6) then reveals that the worker's marginal benefit of an effort increase is smaller when the manager shares information with him than when she does not. This effect tends to decrease the worker's effort, as the worker substitutes the information provided by the manager for his

personal effort.

The effect of information sharing on the worker's cost of effort, however, works in the opposite direction: Given that $t < 1$, the worker's effort is less costly if the manager shares with him her information, which tends to increase the worker's effort. This second effect prevails when avoiding duplication of effort and searching for information in the right direction is sufficiently important, as captured by small values of t .

4.2 Optimal Sharing of Managerial Information

Moving back one period to date 1, we next determine the informed manager's optimal probability, γ^* , with which she will divulge her information to the worker when she is informed. If she shares her information, her expected payoff is $\text{prob}(|\eta_S - \eta| < q) \Pi$, while if she conceals it, her expected payoff is $\text{prob}(|\eta_P - \eta| < q) \Pi$. In both cases the manager's posterior belief $\bar{p} = N(\bar{\eta}, 1/\bar{h})$ is based on both the worker's and her own signals, but its precision can differ, depending on the effects of information sharing on the worker's effort, as is apparent from (5) and (9). This observation leads to our first main result.

Proposition 3 *There exist t_1 and t_2 , $0 < t_1 < t_2 < 1$, such that*

- (i) if $t \leq t_1$, the manager always shares her information with the worker, i.e., $\gamma^*(t) = 1$;*
- (ii) if $t \geq t_2$, the manager always conceals her information from the worker, i.e., $\gamma^*(t) = 0$;*
- (iii) if $t \in (t_1, t_2)$, the manager shares her information with the worker with conditional probability $\gamma^*(t) \in (0, 1)$. Moreover, $\gamma^*(t)$ strictly decreases in t .*

Proof: (i) Recall that $\hat{\gamma}$ denotes the worker's belief about the manager's strategy γ . As can be readily verified from (11), $e_N^*(\hat{\gamma})$ strictly increases in $\hat{\gamma}$, with $e_N^*(0) > 0$ and $e_N^*(1) < \infty$. On the other hand, $e_S^*(t)$ is independent of $\hat{\gamma}$ but strictly decreases in t , with $e_S^*(t) \rightarrow \infty$ as $t \rightarrow 0$, so that $e_S^*(t = 0) > e_N^*(\hat{\gamma} = 1)$. Moreover, as shown in the proof of Proposition 2, it must be $e_S^*(t = 1) < e_N^*(\hat{\gamma})$ for any $\hat{\gamma}$. Hence, there must exist a $t_1 > 0$ such that $e_S^*(t_1) = e_N^*(\hat{\gamma} = 1)$ and $e_S^*(t) > e_N^*(\hat{\gamma} = 1)$ for any $t < t_1$. That is, for these values of t , the worker provides a strictly higher level of effort if the manager shares with him her information than if she does not, making it optimal for the manager to always share her information (when she has received a signal). Thus, in this case $\gamma^* = 1$ and, in equilibrium, $\hat{\gamma} = \gamma^*$.

(ii) Because (a) $e_S^*(t = 1) < e_N^*(\hat{\gamma})$ for any $\hat{\gamma}$ (again by Proposition 2), (b) $e_S^*(t)$ strictly decreases in t , and (c) $e_N^*(\hat{\gamma})$ strictly increases in $\hat{\gamma}$, there must exist a $t_2 \in (t_1, 1)$ such that $e_S^*(t_2) = e_N^*(\hat{\gamma} = 0)$ and $e_S^*(t) < e_N^*(\hat{\gamma} = 0)$ for all $t > t_2$. In this case, divulging the manager's information to the worker decreases his effort, so the manager finds it optimal to conceal her information. That is, $\gamma^* = 0$ and, in equilibrium, $\hat{\gamma} = \gamma^*$.

(iii) For $t \in (t_1, t_2)$, we have $e_S^*(t) > e_N^*(\hat{\gamma} = 0)$ and $e_S^*(t) < e_N^*(\hat{\gamma} = 1)$. No equilibrium in pure strategies therefore exists: If the worker expects the manager to never share her information ($\hat{\gamma} = 0$), he provides less effort if the manager does not reveal any information to him than he would if the manager divulged her signal. Hence, the manager has an incentive to deviate by divulging her information. If the worker expects the manager to always share her information if she observes a signal ($\hat{\gamma} = 1$), he provides less effort when the manager reveals her information than he would if she concealed her signal. The manager therefore has an incentive to deviate by concealing her information. Thus, for these values of t the equilibrium requires that the worker believes that the manager plays a mixed strategy $\hat{\gamma}(t)$ such that $e_S^*(t) = e_N^*(\hat{\gamma})$. Such a $\hat{\gamma}(t)$ exists by continuity of $e_N^*(\hat{\gamma})$ in $\hat{\gamma}$. As always, equilibrium then requires that $\gamma^*(t) = \hat{\gamma}(t)$.

Finally, the result that $\gamma^*(t)$ strictly decreases in t for $t \in (t_1, t_2)$ is obtained from $e_S^*(t) = e_N^*(\gamma^*)$ and from the facts that $e_N^*(\hat{\gamma})$ strictly increases in $\hat{\gamma}$ and $e_S^*(t)$ strictly decreases in t . $\square$

Thus, for $t \leq t_1$ the manager chooses the first-best information sharing arrangement, but for $t > t_1$ she conceals her information from the worker (with positive probability) even though the first-best requires that the manager always shares her information. The logic behind this result is as explained earlier: By sharing her information with the worker, the principal decreases the worker's marginal benefit from effort. Given that this cannot be offset by strengthening the worker's formal incentives (or, as we show later, given that strengthening the worker's incentives through a monetary contract is costly), the manager prefers to conceal her information if this does not affect the worker's effort costs too much (i.e., if t is not too low).

4.3 A Time-inconsistency Problem

Even though the first-best arrangement calls for information sharing regardless of the value of t , the fact that the manager conceals her information with positive probability when $t > t_1$ does not imply that more information sharing would automatically improve efficiency. After

all, as we have seen, information sharing has detrimental effects on the worker's incentives, so perhaps the outcome described in Proposition 3 is the best that can be achieved in the presence of a moral hazard problem.

In this section we show that this is not the case; that is, more information sharing would improve efficiency even when the negative effect of information sharing on the worker's incentives is taken into account. The reason is that the manager suffers from a time inconsistency problem: By concealing her information, the informed managerial type imposes a negative externality on the uninformed type by sowing doubt in the worker's mind about whether a manager who does not share information is really uninformed. The informed type does not take this externality into account when making her decision to conceal her information from the worker, which leads to excessive information concealment, both from the point of view of efficiency and from the manager's *ex ante* point of view. This is demonstrated in the next proposition, which shows that if the manager could commit to an information sharing strategy at date 0, before she gets a chance to observe a signal, she would commit to more information sharing than the equilibrium level described in Proposition 3.

Proposition 4 *Suppose that at the time of contracting the principal can commit to a probability, $\gamma(t)$, with which she will share her signal with the worker. Then, holding B fixed, the manager commits to a $\gamma^{**}(t)$ such that $\gamma^{**}(t) \geq \gamma^*(t)$, with $\gamma^{**}(t) > \gamma^*(t)$ for all $t \in (t_1, t_2)$. This represents a Pareto improvement compared to the no commitment outcome of Proposition 3.*

Proof: Suppose first $t \leq t_1$. In this case, $\gamma^*(t) = 1$ by Proposition 3 and $e_S^*(t) > e_N^*(t)$ by the proof of Proposition 3. So, suppose, as a way of contradiction, that $\gamma^{**}(t) < \gamma^*(t)$. Then it must be that $e_N^*(\gamma^{**}) < e_N^*(\gamma^*)$ because $e_N^*(\gamma)$ strictly increases in γ . Furthermore, the worker exerts e_N^* more frequently (and $e_S^* > e_N^*$ less frequently), under γ^{**} than under γ^* . Hence, the manager's expected payoff is strictly lower under γ^{**} than it would be under γ^* . For these parameter values the manager therefore optimally commits to $\gamma^{**}(t) = \gamma^*(t) = 1$.

Next assume $t \in (t_1, t_2)$. The proof of Proposition 3 shows that $\gamma^*(t)$ is in this case such that $e_S^*(t) = e_N^*(\gamma^*)$. The manager's expected payoff is therefore the same as if the worker always provided effort $e_S^*(t)$. Because $e_N^*(\gamma)$ increases in γ , if the manager committed to a $\gamma^{**}(t) < \gamma^*(t)$, the effort e_N^* would decrease to $e_N^*(\gamma^{**}) < e_N^*(\gamma^*)$ while e_S^* would remain unchanged.

Hence, the manager's expected payoff would decrease. By the same token, a commitment to a $\gamma^{**}(t) > \gamma^*(t)$ would increase the effort e_N^* to $e_N^*(\gamma^{**}) > e_N^*(\gamma^*)$, which would increase the manager's expected payoff. Hence, it must be $\gamma^{**}(t) > \gamma^*(t)$ for these values of t .

Finally, when $t \geq t_2$, we have $\gamma^*(t) = 0$ by Proposition 3. Thus, $\gamma^{**}(t) < \gamma^*(t)$ is not feasible, so it must be $\gamma^{**}(t) \geq \gamma^*(t)$.

Now, to see that $\gamma^{**}(t)$ Pareto improves upon $\gamma^*(t)$, notice first that the manager has to be at least as well off under γ^{**} as she would be under γ^* because γ^* was a feasible commitment. Thus, it remains to prove that the worker is also better off under γ^{**} . Given that $\gamma^{**}(t) \geq \gamma^*(t)$ for all t , we will show this by showing that the worker's expected utility strictly increases in γ .

The worker's ex ante utility as a function of γ is

$$EU(\gamma) = \beta\gamma [\text{prob}(|\eta_S - \eta| < q) |_{e=e_S^*} B - tC(e_S^*)] + (1-\beta\gamma) \left[\text{prob}(|\eta_P - \eta| < q) |_{e=e_N^*} \hat{\beta} + \text{prob}(|\bar{\eta}_N - \eta| < q) \right]$$

Using 10 and the fact that in equilibrium $\hat{\gamma} = \gamma^*$, this can be rearranged as

$$EU(\gamma) = \beta\gamma [\text{prob}(|\eta_S - \eta| < q) |_{e=e_S^*} B - tC(e_S^*)] + [\text{prob}(|\eta_P - \eta| < q) |_{e=e_N^*} \beta(1-\gamma) + \text{prob}(|\bar{\eta}_N - \eta| < q)]$$

The Envelope Theorem then yields

$$EU'(\gamma) = \beta [\text{prob}(|\eta_S - \eta| < q) |_{e=e_S^*} B - tC(e_S^*) - [\text{prob}(|\eta_P - \eta| < q) |_{e=e_N^*} B - C(e_N^*)]] .$$

We thus get that $EU'(\gamma) > 0$ iff

$$\text{prob}(|\eta_S - \eta| < q) |_{e=e_S^*} B - tC(e_S^*) > \text{prob}(|\eta_P - \eta| < q) |_{e=e_N^*} B - C(e_N^*) .$$

But this has to be true because $t < 1$ implies

$$\begin{aligned} \text{prob}(|\eta_P - \eta| < q) |_{e=e_N^*} B - C(e_N^*) &< \text{prob}(|\eta_P - \eta| < q) |_{e=e_N^*} B - tC(e_N^*) \\ &\leq \text{prob}(|\eta_S - \eta| < q) |_{e=e_S^*} B - tC(e_S^*) , \end{aligned}$$

where the last inequality follows because e_S^* solves 3. $\square$

Proposition 4 tells us that the inefficiently low level of information sharing that plagues the firm is at least partly due to commitment problems. If the manager could *ex ante* publicly commit to the frequency with which she shares her information with her subordinate, then for

any given contract she would never choose less information sharing than she does in the absence of commitment, and for a range of parameter values she would choose strictly more information sharing. Thus, commitment brings the firm closer to the first-best arrangement, which requires complete information sharing.

As explained earlier, the intuition for this result is that it stems from an externality imposed by the informed managerial type on the uninformed type. To see the mechanics of this externality in greater detail, consider a disclosure probability γ . Under full commitment, if the principal fails to share with the worker any information, the worker understands that the probability the principal is actually informed is $\frac{\beta(1-\gamma)}{1-\beta+\beta(1-\gamma)}$. Thus, in this case, any change from the information disclosure probability γ – say, a decrease to $\gamma' < \gamma$ – would get fully reflected in the worker's updating process: he would now believe that in the absence of information sharing the principal is informed with probability $\frac{\beta(1-\gamma')}{1-\beta+\beta(1-\gamma')} > \frac{\beta(1-\gamma)}{1-\beta+\beta(1-\gamma)}$. The increase in the worker's belief that the principal is informed decreases his marginal benefit of effort, which induces him to choose less effort than he would under γ .

In contrast, in the absence of commitment the worker cannot directly observe the principal's choice of γ ; instead, he forms a belief about her strategy. But a deviation by the principal from this strategy cannot affect the worker's belief about the likelihood the principal is informed, which means that the above negative effect on the worker's effort is absent in the case of no commitment. This makes it more tempting for the principal to choose a γ smaller than γ^{**} .

Commitment removes this externality problem and makes both the manager and the worker at least weakly better off. It is immediate that the manager cannot be worse off under commitment than under no commitment. To see that the worker is also better off, notice that by not sharing her information, the informed manager effectively "tricks" the worker into putting in more effort than he would if he knew that the manager is informed. Similarly, when the manager is uninformed, the worker puts in less effort than he would if he were positive that the manager is indeed uninformed. By committing to more information sharing, the manager decreases the worker's incentives to under-provide effort when the manager is uninformed and to over-provide effort when the manager is informed, which makes the worker better off.

The above discussion raises the natural question whether the manager would want to commit to the first-best level of information sharing, that is, whether $\gamma^{**}(t) = 1$ even if $\gamma^*(t) < 1$. Proposition 5 below shows that this is not always the case: When the probability that the

manager receives a signal is relatively large, she prefers a commitment according to which she conceals her information from the worker with positive probability.

Proposition 5 *For any $t > t_1$, there exists a $\beta^+ < 1$ (possibly depending on t) such that $\gamma^{**}(t) < 1$ for all $\beta \geq \beta^+$.*

Proof: To see that we don't always get $\gamma^{**}(t) = 1$, we need to examine the manager's expected payoff as a function of γ^{**} , $E\pi(\gamma^{**})$:

$$E\pi(\gamma^{**}) = [\text{prob}(|\eta_S - \eta| < q) \gamma^{**} \beta + \text{prob}(|\eta_P - \eta| < q) \beta (1 - \gamma^{**}) + \text{prob}(|\eta_N - \eta| < q) (1 - \beta)] \Pi.$$

Differentiating with respect to γ^{**} yields

$$\begin{aligned} \frac{\partial E\pi(\gamma^{**})}{\partial \gamma^{**}} = & \left[\beta [\text{prob}(|\eta_S - \eta| < q) - \text{prob}(|\eta_P - \eta| < q)] + \beta (1 - \gamma^{**}) \frac{\partial \text{prob}(|\eta_P - \eta| < q)}{\partial e_N^*} \frac{\partial e_N^*}{\partial \gamma^{**}} \right. \\ & \left. + (1 - \beta) \frac{\partial \text{prob}(|\eta_N - \eta| < q)}{\partial e_N^*} \frac{\partial e_N^*}{\partial \gamma^{**}} \right] \Pi. \end{aligned}$$

Suppose $t > t_1$, so that $\gamma^* < 1$ (by Proposition 3) and assume $\gamma^{**} = 1$, so that $\hat{\beta} = 0$ for any $\beta < 1$. Parts (ii) and (iii) in the proof of Proposition 3 imply $e_N^*(\gamma^{**} = 1) > e_S^*(t)$ for all $t > t_1$, from which $\text{prob}(|\eta_S - \eta| < q) - \text{prob}(|\eta_P - \eta| < q) < 0$. Observe that $e_N^*(\gamma^{**} = 1)$ and $e_S^*(t)$ do not directly depend on β ; hence, this inequality holds for any $\beta < 1$. Since the second term in the brackets disappears when $\gamma^{**} = 1$, and the third term converges to zero when $\beta \rightarrow 1$, we get that for β close to one it must be $\frac{\partial E\pi(\gamma^{**})}{\partial \gamma^{**}}|_{\gamma^{**}=1} < 0$. This implies that for large values of β , the optimal $\gamma^{**}(t)$ is strictly less than one. $\square$

To sum up, full commitment does not always lead to fully efficient information sharing, although it tends to induce the principal to share her information more efficiently than in the absence of commitment. Of course, full commitment to an arbitrary probability for sharing information may not be a realistic arrangement. More realistically, the manager may be able to commit to the simpler, and easier to enforce, arrangement, in which she always shares her information with the worker. An example of this could be the "open book" management approach, whereby a company explains to its employees its goals and commits to share with them all of its operating information. This management strategy has been adopted by a number of firms, including Southwest Airlines, Home Depot, and Whole Foods (Dixon et al, 2004).

As shown in the next proposition, whenever the manager's optimal information sharing strategy in the no commitment case is to randomize, the limited commitment option induces her to instead commit to full information sharing.

Proposition 6 *Suppose that at the time of contracting, the only possible commitment option available to the principal is to always share her signal with the worker. Then she adopts this commitment (i.e., sets $\gamma^{**}(t) = 1$) for all $t < t_3$, where $t_3 \geq t_2$. If $t_3 < 1$, she never shares her information with the worker when $t \geq t_3$.*

Proof: The proof of Proposition 3 already shows the result for $t \leq t_1$; thus, let $t \in (t_1, t_2)$. The same argument as the one in the proof of Proposition 4 shows that the manager's expected payoff is higher under any $\gamma^{**}(t) > \gamma^*(t)$ than under $\gamma^*(t)$. Hence, the manager prefers $\gamma^{**}(t) = 1$ to $\gamma^*(t)$.

Now suppose $t \geq t_2$. We have that a commitment to $\gamma^{**} = 1$ is better than no commitment (in which case the principal chooses $\gamma^* = 0$) iff

$$\begin{aligned} E\pi(\gamma^{**} = 1) &= [\text{prob}(|\eta_S - \eta| < q) |_{e=e_S^*(t)}\beta + \text{prob}(|\eta_N - \eta| < q) |_{e=e_N^*(1)}(1 - \beta)] \Pi \\ &> [\text{prob}(|\eta_P - \eta| < q) |_{e=e_N^*(0)}\beta + \text{prob}(|\eta_N - \eta| < q) |_{e=e_N^*(0)}(1 - \beta)] \Pi = E\pi(\gamma^* = 0). \end{aligned}$$

Now, as is apparent from the above inequality, $E\pi(\gamma^* = 0)$ does not depend on t , whereas $E\pi(\gamma^{**} = 1)$ decreases in t (because $e_S^*(t)$ decreases in t). Hence, there must exist a cutoff level $t_3 \in [t_2, 1]$ such that the manager commits to $\gamma^{**} = 1$ if $t \leq t_3$ and does not commit otherwise. $\square$

It should be pointed out here that Proposition 6 does not imply that limited commitment is better from the efficiency point of view than full commitment. While limited commitment leads to more efficient information sharing than full commitment, it can have adverse effects on the worker's equilibrium effort.

5 Discussions and Extensions

This section has three goals. The first is to examine how the manager's information sharing strategy depends on her ability. Second, we wish to explore the link between information

sharing and authority in organizations. Finally, we want to demonstrate that our insights carry over into a setting in which contingent contracts are feasible.

5.1 Managerial Ability

In this subsection, we examine how the optimal information sharing arrangement depends on the manager's ability. In particular, should we expect better informed managers to share their information with subordinates more frequently? In the context of our model, one way to capture how well the manager is informed is through the precision h_m of her signal. At first blush, it might appear as though an increase in the accuracy of the manager's information should increase the opportunity cost of concealing that information from the worker, which should favor more information sharing. In fact, the opposite is true in our model:

Proposition 7 *Let t_1 and t_2 be as in Proposition 3. An increase in h_m decreases both t_1 and t_2 , and leads to a decrease in $\gamma^*(t)$ for each $t \in (t_1, t_2)$. That is, managers with more precise information are less likely to share their information with their subordinates.*

Proof: From the proof of Proposition 2, we have

$$\frac{\partial T_j}{\partial h_j} = \frac{q}{\sqrt{2\pi h_j}} \exp\left(-\frac{q^2}{2} h_j\right) > 0, \quad j \in \{S, N, P\}.$$

Differentiating with respect to h_m and using $\partial h_j / \partial h_m = 1$ for $j = S, P$ (from (5) and (9)) and $\partial h_N / \partial h_m = 0$ (from (8)), we get

$$\begin{aligned} \frac{\partial}{\partial h_m} \left[\frac{\partial T_j}{\partial h_j} \right] &= \frac{\partial^2 T_j}{\partial h_j^2} = -\frac{q}{2\sqrt{2\pi h_j}} \left(\frac{3}{h_j} + q^2 \right) \exp\left(-\frac{q^2}{2} h_j\right) < 0, \quad j = S, P \\ \frac{\partial}{\partial h_m} \left[\frac{\partial T_N}{\partial h_N} \right] &= 0 \end{aligned}$$

Implicit differentiation of (12) with respect to h_m then yields

$$\frac{\partial e_S^*}{\partial h_m} = \frac{\partial^2 T_S}{\partial h_S^2} \frac{h'_w(e_S^*) B}{t C''(e_S^*) - \frac{\partial T_S}{\partial h_S} h''_w(e_S^*) B} < 0.$$

Similarly, implicit differentiation of (13) yields

$$\frac{\partial e_N^*}{\partial h_m} = \frac{\hat{\beta} \frac{\partial^2 T_P}{\partial h_P^2} h'_w(e_N^*) B}{C''(e_N^*) - \left[\hat{\beta} \frac{\partial T_P}{\partial h_P} h''_w(e_N^*) + (1 - \hat{\beta}) \frac{\partial T_N}{\partial h_N} h''_w(e_N^*) \right] B}.$$

Now, at any $t \in [t_1, t_2]$, we have $e_S^* = e_N^*$. Denote this effort as e^* (suppressing the fact that e_S^* depends on t), so that we can write

$$\left. \frac{\partial e_S^*}{\partial h_m} \right|_{e_S^*=e^*} = \frac{\frac{\partial^2 T_S}{\partial h_S^2} h'_w(e^*) B}{t C''(e^*) - \frac{\partial T_S}{\partial h_S} h''_w(e^*) B}$$

$$\left. \frac{\partial e_N^*}{\partial h_m} \right|_{e_N^*=e^*} = \frac{\hat{\beta} \frac{\partial^2 T_P}{\partial h_P^2} h'_w(e^*) B}{C''(e^*) - \left[\hat{\beta} \frac{\partial T_P}{\partial h_P} + \left(1 - \hat{\beta}\right) \frac{\partial T_N}{\partial h_N} \right] h''_w(e^*) B}.$$

As we have argued in the proof of Proposition 2, for any given effort level it must be $0 < \frac{\partial T_S}{\partial h_S} = \frac{\partial T_P}{\partial h_P} < \frac{\partial T_N}{\partial h_N}$. Furthermore, $h_S = h_P$ implies $\left| \frac{\partial^2 T_S}{\partial h_S^2} \right| = \left| \frac{\partial^2 T_P}{\partial h_P^2} \right|$. Hence, given that (i) $t < 1$, (ii) $\hat{\beta} \leq 1$, and (iii) $h'_w > 0$ and $h''_w < 0$, we have $\left| \frac{\partial e_S^*}{\partial h_m} \right| > \left| \frac{\partial e_N^*}{\partial h_m} \right|$. Thus, for any $t \in [t_1, t_2]$, e_S^* declines more than e_N^* if $\hat{\gamma}$ (and hence also $\hat{\beta}$) is held fixed. Therefore, t_1 and t_2 must both decrease when h_m increases. Moreover, $\hat{\gamma}$ (which in equilibrium is equal to γ^*) must also decrease, in order to increase $\hat{\beta}$ and restore the equality $e_S^* = e_N^*$ for $t \in (t_1, t_2)$. $\square$

As explained earlier, the marginal impact of an increase in the precision of the worker's signal on the project's success probability is lower with information sharing than without, and this effect tends to decrease the worker's effort. Moreover, the bigger is the precision of the signal that the manager shares with the worker, the stronger is this effect. This is what leads to the results in Proposition 7. A caveat is in order here, however: We are holding t , which measures the impact of the manager's information on the worker's marginal cost, constant. More realistically, t might be decreasing in the precision of the shared information. This effect would work in favor of more information sharing and, if the decrease in t were sufficiently pronounced, could overturn the above result.

An alternative way to measure how well the manager is informed would be through the probability β with which she observes a signal. As demonstrated in the next proposition, an increase in the probability that the manager is informed has very different implication than an increase in the quality of the manager's signal.

Proposition 8 *Let t_1 and t_2 be as in Proposition 3. An increase in β leaves t_1 unchanged, increases t_2 , and leads to an increase in $\gamma^*(t)$ for each $t \in (t_1, t_2)$. That is, a manager who is more likely to be informed is also more likely to share her information with the worker, conditional on being informed.*

Proof: As shown in the proof of Proposition 2, the cutoff level t_1 is given by the condition $e_S^*(t_1) = e_N^*(\hat{\gamma} = 1)$. Since $\hat{\gamma} = 1$ implies $\hat{\beta} = 0$, both of these efforts are independent of β . Consequently, t_1 is also independent of β .

The cutoff level t_2 is given by the condition $e_S^*(t_2) = e_N^*(\hat{\gamma} = 0)$. Effort $e_S^*(t_2)$ does not depend on β . On the other hand, given that $\hat{\beta} = \beta$ when $\hat{\gamma} = 0$, $e_N^*(\hat{\gamma} = 0)$ decreases in β , as can be seen from (11). Thus, an increase in β implies that $e_S^*(t_2) > e_N^*(\hat{\gamma} = 0)$ at the initial level of t_2 , which in turn implies an increase in t_2 , because e_N^* does not depend on t and e_S^* decreases in t .

Finally, for $t \in (t_1, t_2)$, the manager plays a mixed strategy, sharing her information with the worker with (conditional) probability $\hat{\gamma} \in (0, 1)$, which is determined by the requirement that $\hat{\beta}$ is such that $e_S^*(t) = e_N^*(\hat{\gamma})$. Given that $\hat{\beta} = \frac{\beta(1-\hat{\gamma})}{\beta(1-\hat{\gamma})+1-\beta}$ increases in β and decreases in $\hat{\gamma}$, an increase in β requires that $\hat{\gamma}$ increases, in order to keep $\hat{\beta}$ unchanged at the level that equalizes $e_S^*(t)$ and $e_N^*(\hat{\gamma})$. In equilibrium, it must be $\gamma^* = \hat{\gamma}$, that is, γ^* must increase as well. $\square$

Propositions 7 and 8 contain potentially testable predictions. A possible interpretation of our setup is that sharing information with a subordinate takes the form of advising or mentoring him. The results of this section then suggest that (i) managers with higher quality information (say, due to their longer experience on the job or due to their higher educational attainment) should be less likely to advise and mentor their subordinates, but (ii) managers who are more likely to be informed about a particular project (say, because they oversee fewer projects), should be more likely to advise and mentor their subordinates. Moreover, (ii) is true not just because managers who are more likely to be informed are automatically more often in a position to be able to offer useful advice, but also because they are more inclined to advise the worker *conditional* on being informed.

5.2 Information Sharing and Authority

The foregoing analysis points to a natural connection between information and authority, and suggests the following question: Does information confer authority in the organization? Prior research, especially Aghion and Tirole (1997), suggests that the answer is Yes and at a first glance our model appears to support this answer. To reach this conclusion, one might reason as

follows: When the manager divulges her information to the worker, she subsequently “rubberstamps” the worker’s posterior distribution $p_S = N(\eta_S, 1/h_S)$, using it as her own posterior distribution. In contrast, when the manager conceals her information from the worker, the ex post rubberstamping disappears: The manager combines the worker’s report $p_N = N(\eta_N, 1/h_N)$ with her own signal to arrive at the posterior distribution $p_P = N(\eta_P, 1/h_P)$. Thus, in the language of Aghion and Tirole (1997), the manager in our model appears to delegate more “real” authority to the worker when divulging her information, because it is in these cases that the rubberstamping occurs. (She always retains “formal” authority whether or not she shares her information with the worker, because she has control over the decision regarding the project.)

In our framework, however, rubberstamping may not be the most meaningful measure of the degree of influence the worker has on the final decision: If his report relies much more on the information divulged by the manager than on his own information, then the worker in fact has little influence on the final decision, even if the manager systematically rubberstamps his report. Conversely, the worker may have little real authority, but a lot of influence on the project choice decision if the manager’s final estimate places a lot of weight on the information contributed by the worker. In line with these observations, a more appropriate measure of the worker’s authority in our model is the weight placed on the worker’s opinion in the determination of the manager’s posterior belief $\bar{p} = N(\bar{\eta}, 1/\bar{h})$ - a measure we refer to as the worker’s *implicit* authority.

Consider the mean of the manager’s posterior belief:

$$\bar{\eta} = \frac{h_0\eta_0 + h_w(e)\eta_w + h_m\eta_m}{h_0 + h_w(e) + h_m}. \quad (14)$$

Posterior mean $\bar{\eta}$ is a weighted average of the prior mean, and of the worker’s and the manager’s respective signals. The worker’s implicit authority is therefore captured by the weight, α_w , that the manager places on the worker’s signal:

$$\alpha_w = \frac{h_w(e)}{h_0 + h_w(e) + h_m}. \quad (15)$$

Moreover, α_w can also be viewed as measuring the fraction of the manager’s posterior precision (given by $\bar{h} = h_0 + h_w(e) + h_m$) that is due to the information contributed by the worker. This reinforces the notion that α_w is a natural measure of the worker’s authority within our

framework.

It is straightforward that α_w is a strictly increasing function of $h_w(e)$. The impact of information sharing on the weight placed on the worker's opinion then depends on how divulging information affects the worker's effort. Building on Proposition 3, we obtain the following result:

Proposition 9 *Suppose the manager is informed. The effect of information sharing on the worker's implicit authority α_w is described by (i) - (iii) below:*

- (i) When $t < t_1$, the worker's implicit authority is larger when the manager divulges to him her information than it would be if she concealed it.*
- (ii) When $t > t_2$, the worker's implicit authority is larger when the manager conceals her information than it would be if she divulged it to the worker.*
- (iii) When $t \in [t_1, t_2]$, the worker's implicit authority is the same when the manager divulges to him her information as when she conceals it.*

Proof: All three parts follow from Proposition 3. When $t < t_1$, the worker's effort is larger when information is shared than when it is concealed from him. Hence, so is α_w . When $t > t_2$, the worker's effort (and therefore also α_w) is larger when the manager conceals her information. Finally, when $t \in [t_1, t_2]$ the manager plays a mixed strategy and the worker's effort level is the same whether the manager shares her information with him or not. $\square$

Of course, in equilibrium, both the probability of information sharing and the worker's effort are endogenous and are jointly determined by the cost parameter t . In particular, a smaller t leads to both more information sharing and more effort, and therefore also to a greater implicit authority for the worker, α_w . Thus, if we want to talk about the effect of information sharing on the worker's implicit authority in a meaningful way, we need to imagine that the manager makes an unexpected decision regarding information sharing, say, in response to an exogenous and unexpected event. Proposition 9 characterizes the effects of information sharing in such a thought experiment.

5.3 Monetary Incentives

In this section, we demonstrate that our core arguments continue to hold when monetary incentives are feasible, as long as contracting is imperfect. Specifically, assume that it is possible

to write contracts in which the worker's pay is contingent on whether the project is undertaken. Assume also that the worker is protected by limited liability, and has zero initial wealth, zero reservation utility, and no private benefits: $B = 0$. Let b denote the bonus the agent gets if the project goes ahead and note that, as is standard in limited liability models, an optimal contract will pay the agent zero if the project does not go ahead.

In this setting, the worker's efforts e_S and e_N are given by first-order conditions analogous to (6) and (11), except that the private benefit B is replaced by the bonus b . The manager thus chooses b , e_S , e_N , and γ to solve the following optimization program:

$$\max_{b, e_S, e_N, \gamma} [\text{prob}(|\eta_S - \eta| < q) \gamma \beta + \text{prob}(|\eta_P - \eta| < q) \beta (1 - \gamma) + \text{prob}(|\eta_N - \eta| < q) (1 - \beta)] (\Pi - b) \quad (16)$$

subject to

$$\frac{1}{t} \frac{\partial \text{prob}(|\eta_S - \eta| < q)}{\partial h_S} h'_w(e_S) b = C'(e_S); \quad (17)$$

$$\left[\frac{\partial \text{prob}(|\eta_P - \eta| < q)}{\partial h_P} \hat{\beta} + \frac{\partial \text{prob}(|\eta_N - \eta| < q)}{\partial h_N} (1 - \hat{\beta}) \right] h'_w(e_N) b = C'(e_N); \text{ and} \quad (18)$$

$$\gamma = \arg \max_{\gamma'} [\text{prob}(|\eta_S - \eta| < q) \gamma' + \text{prob}(|\eta_P - \eta| < q) (1 - \gamma')] (\Pi - b). \quad (19)$$

Here, (17) and (18) are the worker's incentive compatibility constraints for efforts e_S and e_N respectively, and (19) is the manager's incentive compatibility constraint that ensures that if informed, the manager is willing to divulge her signal to the worker with probability γ . The worker's participation constraint is implied by (17) and (18) and we omit it in the statement of the problem.

As in our baseline model, the trade-off due to information sharing is reflected in the difference between the worker's two IC constraints (17) and (18): By concealing her information, the manager increases the worker's marginal benefit of effort, *ceteris paribus*, but also his marginal cost. What is different here is that, unlike the private benefit, B , in the baseline model, the bonus, b , is chosen by the manager and will in general depend on the cost parameter t . Nevertheless, Proposition 10 below shows that this difference is immaterial for the manager's optimal information sharing strategy.

Proposition 10 *When monetary incentives are feasible, the manager's optimal information sharing strategy continues to be characterized by the same two cutoff levels t_1 and t_2 as described in Proposition 3.*

Proof: For any given contract b , the manager's choice of γ is driven by a comparison of e_S and e_N : If, for a given $\hat{\beta}$, it is true that $e_S > e_N$ then the manager sets $\gamma = 1$, if $e_S < e_N$ then she sets $\gamma = 0$, and if $e_S = e_N$ she employs a mixed strategy, with $\gamma \in [0, 1]$. The comparison between e_S and e_N depends solely on the left hand sides (LHS) of the two first-order conditions (17) and (18), as follows: Let $e_S^{**} = e_S^{**}(b)$ be the solution to the first-order condition (17). Then by concavity of the worker's problem

$$\begin{aligned} LHS(17)|_{e_S=e_S^{**}} &> LHS(18)|_{e_N=e_S^{**}} &\implies e_S^{**} > e_N^{**}; \\ LHS(17)|_{e_S=e_S^{**}} &= LHS(18)|_{e_N=e_S^{**}} &\implies e_S^{**} = e_N^{**}; \\ LHS(17)|_{e_S=e_S^{**}} &< LHS(18)|_{e_N=e_S^{**}} &\implies e_S^{**} < e_N^{**}. \end{aligned}$$

But b cancels out in the above comparisons.

Now consider the manager's choice of her information sharing strategy γ . It is apparent from (19) that the optimal γ is given solely by comparing $\text{prob}(|\eta_S - \eta| < q)$ with $\text{prob}(|\eta_P - \eta| < q)$. Furthermore, $\text{prob}(|\eta_S - \eta| < q) > \text{prob}(|\eta_P - \eta| < q)$ if $e_S^{**} > e_N^{**}$ and the reverse is true when $e_S^{**} < e_N^{**}$. Therefore, the manager's optimal γ does not depend on b (as long as $b > 0$, which can be easily verified to hold for the optimal contract). Moreover, the same logic also applies to our baseline model (where the worker is motivated by the private benefit B), which means that the cutoff levels t_1 and t_2 in Proposition 3 do not depend on B . Since the only difference between the first-order conditions that characterize e_S^* and e_N^* in the baseline model versus e_S^{**} and e_N^{**} here is that B in the baseline model was replaced by b here, the same cutoff levels t_1 and t_2 that described the manager's optimal information sharing strategy in Proposition 3 must also apply here. $\square$

Proposition 10 tells us that the presence of monetary incentives does not alter our baseline model's qualitative conclusions about the effects that information sharing has on incentives. The manager's optimization problem is somewhat more complex here because she needs to choose an optimal bonus for the worker, but this additional consideration does not affect the fundamental trade-offs that we identified and studied. The reason is that at the time the manager chooses whether to share her information with the worker, she treats the bonus b as given. Her goal at this point is therefore to simply maximize the probability of success, which is the same problem she faced in the baseline model.

6 Concluding Remarks

In this paper, we examine optimal information flows between a manager and a worker who is in charge of evaluating a project. The manager may possess information about the project that can make the worker's effort more productive. However, if the manager's information is shared with the worker, the worker's incentives to collect additional information are dampened, which in the end may decrease the total effort the worker is willing to provide. Thus, when effort is difficult to measure and control, it may be optimal for the manager to conceal her information, even though in the first-best outcome she would always share her information with the worker.

We demonstrate that the inefficiently low level of information sharing is exacerbated by a time-inconsistency problem that the manager faces due to not being able to commit *ex ante* to an information sharing policy. We also show that managers with more precise information have a stronger incentive to withhold their information from the worker and that information sharing decreases the worker's implicit authority, which we define as the weight that the manager puts on the worker's opinion when deciding whether to undertake the project.

In order to provide clear and intuitive results, we have abstracted from some considerations that may be important in practice. For example, we have assumed that both the precision of the manager's signal and the probability that she receives a signal are exogenous. In reality, managers may be able to influence both of these parameters by expending costly effort. Our framework could readily incorporate such an extension and could then be used to address additional questions of interest, such as whether the manager might have an incentive to publicly commit to limit the amount of effort she expends on information gathering, in order not to dilute the worker's incentives.

A related consideration that is absent from our model but potentially important in reality is that most employment relationships involve multiple workers. A multi-agent setting would bring to light a whole set of new issues. For example, as in Gromb and Martimort (2007), one could ask whether the manager would be better off hiring one or two agents to collect signals. And if two agents are hired, should they collect their respective signals simultaneously or sequentially? A multi-agent extension of our framework could be used to shed light on these interesting problems.

A Appendix

A.1 Proof of Lemma 1

Suppose the mean of the worker's posterior belief after observing his signal is p_i , where $i = S$ if the manager has shared information with the worker and $i = N$ if she has not. The worker can misrepresent his information by reporting mean $\eta_i + x$ and he will choose x that maximizes the probability he receives his private benefit, Q_i . When $i = S$, this probability is given by

$$Q_S = \text{prob}(|\eta_S + x - \eta| < q) = \text{prob}(\eta_S + x - q < \eta < \eta_S + x + q), \text{ where } \eta \sim N(\eta_S, 1/h_S).$$

Rewriting this in terms of the standard normal distribution, we get

$$Q_S = \text{prob}\left((x - q)\sqrt{h_S} < z < (x + q)\sqrt{h_S}\right), \text{ where } z \sim N(0, 1),$$

which can be expressed as

$$Q_S = \frac{1}{\sqrt{2\pi}} \int_{(x-q)\sqrt{h_S}}^{(x+q)\sqrt{h_S}} e^{-z^2/2} dz.$$

When $i = N$, the worker believes that with probability $\hat{\beta}$ the manager has observed a signal, but has not divulged it. In such a case, the precision of the manager's posterior belief after receiving the worker's report is h_P , obtained through the same formula as h_S , except that the signals are aggregated in a different order. But the order of signal aggregation is irrelevant for the final belief. Thus, using again the standard normal distribution, we can write the worker's belief about the probability of success in this case as

$$Q_N = \hat{\beta} \frac{1}{\sqrt{2\pi}} \int_{(x-q)\sqrt{h_P}}^{(x+q)\sqrt{h_P}} e^{-z^2/2} dz + (1 - \hat{\beta}) \frac{1}{\sqrt{2\pi}} \int_{(x-q)\sqrt{h_N}}^{(x+q)\sqrt{h_N}} e^{-z^2/2} dz.$$

Maximizing with respect to x using Leibnitz' rule, one can easily show that $\partial Q_i / \partial x = 0$ if and only if $x = 0$; and that $\partial^2 Q_i / \partial x^2 < 0$ at $x = 0$, $i = S, N$. By continuity of Q_i , $x = 0$ therefore maximizes the probability of receiving the private benefit B : Truthfully reporting the mean of the posterior distribution is optimal for the worker.

His report of h_i has no impact on his payoff, and he reports it truthfully also. $\square$

References

- [1] Aghion, Philippe and Jean Tirole (1997), "Formal and Real Authority in Organizations," *Journal of Political Economy* **105**(1), 1-29.
- [2] Alonso, Ricardo, Wouter Dessein, Niko Matouschek (2007), "When Does Coordination Require Centralization," *American Economic Review* **98**(1), 145-179.
- [3] Alonso, Ricardo and Niko Matouschek (2007), "Relational Delegation," *Rand Journal of Economics* **38**(4), 1070-1089.
- [4] Baker, George, Robert Gibbons, and Kevin. J. Murphy (1999), "Informal Authority in Organizations," *Journal of Law, Economics, and Organization* **15**(1), 56-73.
- [5] Core, John and Jun Qian (2002), "Project Selection, Production, Uncertainty, and Incentives," Working Paper, Boston College.
- [6] Crawford, Vincent and Joel Sobel (1982), "Strategic Information Transformation," *Econometrica* **50**(6), 1431-1451.
- [7] Davis, Tim R.V. (1997), "Open-book Management: Its promise and Pitfalls," *Organizational Dynamics* **25**(3), 7-20.
- [8] Demski, Joel and David Sappington (1986), "On the Timing of Information Release," *Information Economics and Policy* **2**(4), 307-316.
- [9] Demski, Joel and David Sappington (1987), "Delegated Expertise," *Journal of Accounting Research* **25**(1), 68-89.
- [10] Dessein, Wouter (2002), "Authority and Communication in Organizations," *Review of Economic Studies* **69**(4), 811-838.
- [11] Dixon, Mark R., Linda J. Hayes, and Jack Stack (2004), "Changing Conceptions of Employee Compensation," *Journal of Organizational Behavior Management* **23**(2-3), 95-116.
- [12] Gromb, Denis and David Martimort (2007), "Collusion and the Organization of Delegated Expertise," *Journal of Economic Theory* **137**(1), 271-299.
- [13] Lambert, Richard (1986), "Executive Effort and Selection of Risky Projects," *Rand Journal of Economics* **17**(1), 77-88.
- [14] Levitt, Steven D. and Christopher M. Snyder (1997), "Is no News Bad News? Information Transmission and the Role of "Early Warning" in the Principal-Agent Model," *RAND Journal of Economics* **28**(4), 641-661.
- [15] Malcomson, James (2009), "Principal and Expert Agent," *The B.E. Journal of Theoretical Economics (Contributions)* **9**(1), Article 17.
- [16] Prendergast, Canice (1993), "A Theory of 'Yes Men'," *American Economic Review* **83**(4), 757-770.
- [17] Zabojnik, Jan (2002), "Centralized and Decentralized Decision Making in Organizations," *Journal of Labor Economics* **20**(1), 1-22.