Working paper

HIDDEN ACTION OR HIDDEN INFORMATION? HOW INFORMATION GATHERING SHAPES CONTRACT DESIGN.[1]

April 17, 2013

Elisabetta Iossa[a] and David Martimort[b]

Abstract: A risk averse agent gathers information on productivity shocks and produces accordingly on behalf of his principal. Information gathering is imperfect so that the agent has either complete or no knowledge at all of those shocks. The model allows for moral hazard in information gathering, private information on productivity shocks and moral hazard on operating effort. Two polars scenarios of the agency literature with either pure hidden action (the agent exerts operating effort not knowing yet the realization of the shock) or pure hidden information (the agent knows that shock when exerting operating effort) arise with positive probability. An optimal menu of linear contracts mixes high-powered, productivity-dependent screening options following "good news " with a fixed low-powered option that solves a pure moral hazard problem otherwise.

Keywords: Incentive mechanisms, information gathering, hidden action, hidden information. .

1. INTRODUCTION

Since its early inception, the agency literature has put much emphasis on how uncertainty impacts on the design of contracts and the overall performances of principal-agent arrangements.[1] The two competing paradigms of hidden information and hidden actions are indeed both based on the idea that uncertainty somewhat blurs the relationship between the agent's local knowledge and his performances.[2]

Beyond sharing such broad common theme, those two alternative paradigms are nevertheless based on rather different assumptions and, as a result, provide also somewhat different insights. In the hidden action model, uncertainty is exogenous and affects the agent's performances after he had already chosen his (non-verifiable) action. Instead, in the hidden information paradigm, the agent chooses the level of his contractual activity for the principal after having privately learned some productivity shock that affects his performance. Quite surprisingly, the literature has so far failed to recognize that whether actions come before uncertainty is released or after is to a large extent endogenous and part of the organizational problem under scrutiny. Indeed, agents do invest in information gathering to better tailor their actions to realized shocks that may hit their performances. Henceforth, whether hidden action or hidden information is the most relevant paradigm to study is to a large extent a matter of the agent's choice himself (and not necessarily one freedom left to the modeler).

A more complete and probably best descriptive view of agency relationships should definitively take such endogeneity into account. Instead of drawing a stark line between the hidden action and hidden information paradigms, such theory should demonstrate

[1]For useful comments, we wish to thank Malin Arve, Antoine Bommier, Perrin Lefebvre, Manuel Marfan and Wilfried Sand-Zantman as well as participants to seminar in Universities of Pavia and Naples. All errors are ours.

[a]University of Rome Tor Vergata, CMPO, CEPR and EIEF, elisabetta.iossa@uniroma2.it
[b]Paris School of Economics-EHESS, david.martimort@parisschoolofeconomics.eu

[1]Laffont and Martimort (2002, Chapter 1) offers an historical survey on those issues.
[2]See Arrow (1986) and Hart and Holmström (1987) for early important overviews.

how those two alternative scenarios might endogenously arise as different facets of the same overall agency problem.

This paper provides such unifying framework. So doing, we unveil circumstances under which either hidden action or hidden information dominates as the driving force that shapes contractual design. In a nutshell, our main finding is that screening is valuable on the upper tail of the distribution of productivity shocks. Following such "good news," the agent communicates whatever information has been learned and operates under linear contracts stipulating bonuses that are positively correlated with productivity shocks. Incentives are rather high-powered. Instead, no such correlation arises when the agent reports "bad news" or claims being uninformed. In particular, an uninformed agent operates under a low-powered linear contract that entails a bonus independent on the mapping between operating effort and performances. This is akin to the solution of a pure hidden action problem.

Overview of the model. Consider a principal (referred to as "she" in the sequel) who hires a risk averse agent ("he") to gather information relevant for the productive task he exerts on her behalf. The return on that activity depends on the agent's operating effort but also on some uncertain productivity shock. Learning the value of that shock requires costly information gathering. The value of information is positive; ideally, the agent's effort at the operating stage should be tailored to the realized productivity shock.

The contractual relationship is plagued with three different agency problems. First, *ex ante moral hazard* arises because the agent's effort in gathering information is non-verifiable. Importantly, the information gathering technology is imperfect; the agent privately learns the mapping between his operating effort and performances only with some probability that depends on his information gathering effort. When information gathering succeeds, the agent gets *ex post private information.* The principal cannot infer whether a good performance results from a high operating effort or from a favorable productivity shock which has been privately learned by the agent. Finally, *ex post moral hazard* arises when the agent's effort at the operating stage is non-verifiable.

Once information gathering has succeeded, the agent benefits from rents associated to his private information. To reduce those rents and induce an operating effort tailored to the different realizations of the productivity shock, contracts must remain flexible enough and allow some screening. When information gathering has instead failed, the agent chooses his operating effort without knowing how productivity shocks affect performance. The contractual scenario then boils down to pure hidden action. In this context, hidden action and hidden information emerge endogenously within the same contractual framework but under different circumstances.[3]

Overview of the results. We now briefly review some of our main findings.

Productivity shocks are common knowledge. To isolate how the different dimensions of the agency problems interact, let us first consider the simple case where the principal shares with her agent whatever information has been gathered. When information gathering

[3]From a technical viewpoint, the agency model that we study below belongs to the class of so-called "mixed models"mixing elements of moral hazard and asymmetric information. That performances mixes the impact of effort and innate parameter is a well-known assumptions in the seminal model of taxation due to Mirrlees (1971) and in the model of regulation due to Laffont and Tirole (1986). Laffont and Martimort (2002, Chapter 7) offers a typology of those models.

succeeds, the principal and her agent both learn the exact value of the productivity shock so that perfectly aligning the agent's objectives with those of the principal at the operating stage is not an issue. The only incentive problem thus consists in designing correct incentives for information gathering. Had the agent been risk neutral, his objectives and those of the principal could still be perfectly aligned by means of a single fixed price contract. Such contract would make the agent residual claimant for his choice of efforts both ex ante in information gathering but also ex post at the operating stage. Of course, the principal would charge the agent a fixed fee that fully extracts his ex ante surplus.

Risk aversion. It is no surprise that such simple contractual solution fails with a risk averse agent. As a matter of fact, introducing risk aversion in an environment with both ex ante, ex post moral hazard and asymmetric information while still getting some tractable results is a task of notorious difficulty in the folklore of agency theory. In general, few results are available in such complex environments. To nevertheless investigate how those three different sorts of contractual impediments interact in our specific settings while still keeping the same tractability as under risk neutrality, we slightly depart from the familiar modeling of risk aversion by means of a concave Bernoulli utility function which has been widely adopted by most of the agency literature.[4] Instead, our modeling of risk behavior relies on the theory of rank-dependent utilities pushed forward by Quiggin (1982) and Yaari (1987). Risk attitude is then captured by having the risk averse party evaluate his expected payoffs through an increasing and convex transformation of the objective probabilities of the various events when ranked in terms of increasing payoffs.

To account for the fact that the sequentiality between the information gathering and operating stages renders the game intrinsically dymanic, we follow Epstein and Zin (1989) and assume that the agent's intertemporal utility is recursive. Roughly speaking, that modeling also allows us to vary risk attitude for different lotteries that may have been compounded. To maintain tractability of our model, we thus assume that the agent overweighs the probability that information gathering fails while he shares the same beliefs as his principal on the possible realizations of those shocks and is thus risk neutral with respect to those shocks. These assumptions when taken in tandem disentangle the impact on risk-aversion on incentives at the information gathering stage (where it does matter) from that on incentives at the operating stage (where it thus does not).

Uncertainty on whether information gathering succeeds thus imposes some (first-order) risk on the agent that is increasing in the probability of success. To diminish his own risk exposure, the agent reduces his information gathering effort. Moreover, since the principal ends up paying for a risk premium to ensure the agent's participation, she would like more effort reduction than what the agent undertakes by himself. To control for this vertical externality inducing over-effort, the principal still relies on fixed-price contracts but modifies fixed fees, asking for greater fees when information is gathered as a means to limit information gathering.

When information learned is common knowledge, contracts exhibit now some form of dichotomy: Bonuses are only used to solve the ex post moral hazard problem while fixed fees are used to deal with ex ante incentives.

Productivity shocks are private information. Suppose now that the agent gets private information on productivity shocks. Truthtelling constraints now bind together the profile

[4] See Mirrlees (1999), Holmström (1979) and Grossman and Hart (1983).

of the agent's information rents and his bonuses. As a result, that profile has an impact on the agent's incentives to gather information. The above dichotomy result no longer holds.

To illustrate, offering the same fixed-price contracts than when productivity shocks are common knowledge would generate two effects. First, it would induce the agent to claim that the productivity shock is less favorable than what it really is. This strategy reduces the fixed fee that the agent pays back to the principal so that he can now pocket some rent. The principal counters this *"worsening effect",* by making the contract targeted for lower levels of the shock less attractive. This force thereby reduces bonuses, informational rents, but also efforts at the operating stage for those lower levels of the shock.

Second, asking for a lower fixed fee when information is gathered, as a means to limit information gathering effort, would now give incentives to the agent to pretend that he is uninformed. The agent would then behave *de facto* as having observed a mean-value shock. This *"hiding effect"* is fought by making the contract targeted to an uninformed agent less attractive. Following "good news", the *"worsening"* and the *"hiding"* effects just add up. Following "bad news", those two effects countervail.[5]

As the result of those different incentives to manipulate information, the optimal menu of linear schemes is made of two different regimes. On the upper tail of the distribution of productivity shocks, i.e., when the agent has learned "good news" on productivity shocks, screening matters. The agent's bonus is positively correlated with the productivity shock.

On the lower tail of the distribution, the agent instead operates under the same linear contract whether he gets informed on productivity shocks or failed to be so and instead takes expectation over possible realizations of that shock. In other words, this regime is without much concerns for screening and very much akin to solving a pure hidden action problem. Yet, the low powered incentives that arise in that zone does not result from the standard trade-off between moral hazard and risk-aversion since the agent is risk-neutral at the operating stage. Instead, it results from the interaction between the countervailing worsening and hiding effects, which in turn arise from the desire to limit the informational rent and the over-provision of effort in information gathering.

Literature review. There is by now a sizable literature which aims at endogenizeing information structures in principal-agent models and social choice environments.[6] Earlier contributions have analyzed the principal's preferences over the agent's amount of private information but neglected his incentives to acquire information (Lewis and Sappington 1991, 1993, Sobel 1993). Subsequent research has instead pushed incentives on the forefront of the analysis. How much information is gathered at equilibrium, whether information is valuable per se or just a pure rent-seeking activity, and the shape of the optimal contract all depend on fine details of the modeling. This sensitivity makes definitive lessons on whether endogeneizing information structures in agency problems brings something new to our tools kit at best unsettled. The timing for information gathering and contracting (Cremer and Khalil 1992, Cremer, Khalil and Rochet 1998a, Compte

[5] See Lewis and Sappington (1989) for a seminal paper on countervailing incentives.

[6] Bergemann and Vällimäki (2006) provide an exhaustive survey. Beyond models involving a single agent which are close to ours, several authors have analyzed information gathering in multi-agent contexts (Bergemann and Vällimäki 2006, Cremer et al. 2009, Girardi and Yariv 2008, Gershkov and Szentes 2009, Shi 2012). A more tangential literature has studied incentives to gather information to prevent hierarchical collusion (Kofman and Lawarée 1993).

and Jehiel 2008 and Tiersiege 2012), the modeling of information as a continuous or a discrete variable (Cremer and Khalil 1994, Kessler 1998), the amount of competition between agents (Cremer and Khalil 1992, Compte and Jehiel 2008), and whether information gathering refers to the outcome of the agent's effort (Iossa and Legros 2004) are all ingredients that matter a great deal.

Two papers are closest to ours. Cremer, Khalil and Rochet (1998b) extend the seminal model of Baron and Myerson (1982) by allowing a risk neutral agent to gather information on his cost parameter before contract signing but after its mere offering by the principal. An informed agent then chooses a contract within a whole menu of screening options. Instead, the contract stipulates just a single output-payment pair if the agent remains uninformed so that the agent is left with no further choice if he chooses that option. There is no residual uncertainty in information gathering; either the agent learns perfectly his type or he does not, and each possibility may be optimal depending on parameters.

We depart from this model along several lines. First, we introduce risk aversion as a prime source of the ex ante agency problem instead of the ex post participation constraints considered by Cremer, Khalil and Rochet (1998b). In our framework, contract signing takes place before information gathering, which is more in lines with standard models of hidden actions.[7]

Second, we introduce uncertainty in information gathering; there is always some non-zero probability that the agent remains uninformed on shocks before moving to its operating task. From an economic viewpoint, this assumption of course ensures that information gathering leads to a nontrivial moral hazard problem. From a technical viewpoint, this assumption also implies that contracts remain continuous while they exhibit strong discontinuities in Cremer, Khalil and Rochet (1998b).[8]

An important consequence of uncertainty on whether information gathering has succeeded or not is that the same level of performance may have been reached by the agent whether he got informed or not. When informed, the agent has perfect knowledge of the mapping between his operating effort and output. He chooses according to his observed shock within a menu of screening options a contract that determines his effort at the operating stage and the subsequent realized output. When uninformed, the agent still has the same freedom in choosing his operating effort, but he ignores how that effort will be mapped into realized outputs. This is in sharp contrast with Cremer, Khalil and Rochet (1998b). There, the uninformed agent has no choice by assumption and is forced to operate under a pooling contract. In our framework, it makes now no sense for the principal to commit to a single output-payment pair when the agent claims being uninformed since different realizations of that shocks command different output levels. This is precisely that possibility that transforms the contracting problem into a scenario with hidden action whenever the agent remains uninformed.

The second important paper we build on is Szalay (2009). This author provides a rich framework to model uncertain information gathering. Yet, the timing of the contractual game remains similar to Cremer, Khalil and Rochet (1998b). Szalay (2009) highlights that high-powered contracts make the agent's informational rents more risky which boosts

[7]In our view, such timing also better captures the contracting environments that motivated our study, especially in procurement contexts. For instance, Public-Private Partnerships are indeed long-term projects whose contours are designed much before any uncertainty on future demand and cost parameters is released and thus much before private contractors are able to learn relevant information on those parameters.

[8]See Szalay (2009) for an earlier critique of those discontinuities.

information acquisition,[9] while insensitive contracts are used to reduce incentives for information gathering. We differ from Szalay (2009) by restricting the class of information technologies (the agent gets either full information or not at all) but this restriction allows us to expand the analysis to introduce risk aversion and thus an agency cost of information gathering. This meaningful first stage agency cost has of course consequences on second stage incentives. This framework allows us to derive an optimal contract that features at the same time both high- and low-powered incentives depending on whether "good", "bad" or "no" news are gathered by the agent. Those modeling ingredients ensure that optimal contracts in our framework smoothly mix different features borrowed from the hidden information and hidden action literatures.

Finally, our concerns on how ex ante incentives in information gathering and ex post in truthtelling revealing such information is also shared by a burgeoning literature that analyzes the optimal organization of experts in specific contexts (Lambert 1986, Demski and Sappington 1987, Gromb and Martimort 2007, Malcomson 2009, Szalay, 2005, and Dai, Lewis and Loppomo 2006, among others). Experts must be incentivized to gather information and to truthful reveal their findings whenever such information is manipulable. Those papers investigate how those two agency problems interact even if those experts might not necessarily operate themselves the projects they recommend.

Organization of the paper. Section 2 presents the model. Section 3 discusses the benchmark where the principal gathers information on his own. Section 4 studies the case where productivity shocks can be verified so that there is no agency problem ex post, but there is an ex ante incentive issue on information gathering. It shows that there exists a dichotomy between providing incentives at the operating and information gathering stages. Section 5 describes the set of incentive feasible allocations when the agent gets also ex post private information on what he learns. Incentives at the operating and information gathering stages can no longer be separated. Section 6 characterizes the optimal contract and provides a number of comparative statics results. Section 7 briefly concludes. Proofs are relegated to an appendix.

2. THE MODEL

Consider a principal (thereafter "she") who hires an agent (or "he") to gather information on a productivity shock and then produce on her behalf according to the information he may have learned by adapting his operating effort. Of course, this abstract setting describes very well the kind of relationships that arise in some public procurement contexts (for instance in trendy Public-Private Partnerships). There, a public authority delegates to a firm the task gathering information on the cost or demand for a good or service. Such information is relevant to determine the level of service, its quality, its volume. Yet, our methodology and findings are sufficiently general to allow broader interpretations and would apply as well to other principal-agent relationships involving shareholders-CEOs, manufacturer-retailers, clients-lawyers, lender-borrowers, ...

[9]Also related is Baker and Jorgensen (2003). Those authors analyze a moral hazard model à la Holmström and Milgrom (1991) with an agent having a CARA utility function, but whose measured performances depend both on some additive noise and a multiplicative shock (coined as "volatility") that affects the returns on his effort and which is *a priori* known by the agent. Because of a lack of tractability of the CARA model to handle the ex post asymmetric information, the analysis is restricted to the case where a single linear contract is offered. Baker and Jorgensen (2003) demonstrate that high-powered incentives are preferred in more volatile settings.

Technology. The project yields a return y which depends both on the realization of a productivity shock θ and the agent's operating effort e. For simplicity, we assume the following multiplicative specification of the production function:

$$y = (1 + \theta)e.$$

A lower (resp. higher) shock θ reduces (resp. increases) the marginal return on effort.

Ex post asymmetric information. If the agent has succeeded in gathering information, he privately learns the value of the productivity shock θ. Such *ex post* asymmetric information is a first source of agency costs.

The productivity shock is drawn according to the common knowledge cumulative distribution $F(\theta)$ with an atomless and everywhere positive density $f(\theta)$ on the support $\Theta = [-\delta, \delta]$ (with the extra condition $\delta \leq 1$ to maintain positive outputs under all circumstances). The density function $f(\theta)$ is symmetric and centered around zero so that $E_\theta(\theta) = 0$ where $E_\theta(\cdot)$ denotes the expectation operator. Let $\sigma^2 = E_\theta(\theta^2)$ also denote the variance of the productivity shock.

To ensure monotonicity of optimal efforts in some of the scenarios considered below (a familiar condition inherited from incentive compatibility in screening models), let assume that $S(\theta) = \frac{F(\theta)}{f(\theta)(1+\theta)}$ and $R(\theta) = \frac{1-F(\theta)}{f(\theta)(1+\theta)}$ satisfy the following monotonicity properties:

Assumption 1

$$\dot{R}(\theta) \leq 0 \leq \dot{S}(\theta) \quad \forall \theta \in \Theta.$$

Ex post moral hazard. The operating effort is non-verifiable. The contract is thus subject to *ex post moral hazard.* Whether the agent gets informed or not, he still has some freedom in choosing effort at the operating stage. This second leg to the agency problem is important to keep comparable the two scenarios with either symmetric (but incomplete) or asymmetric information between the principal and the agent.

Let the agent's cost of exerting an operating effort e be $\psi(e) = \frac{e^2}{2\lambda}$ where $\lambda > 0$.[10] When the technology is more flexible, i.e., for greater values of λ, the agent's effort becomes more sensitive to what he may have learned on productivity shocks.

Ex ante moral hazard. The choice of how much to invest in information gathering is non-verifiable. Such ex ante moral hazard introduces a third leg to the agency problem.

At cost $\varphi(a)$, the agent gathers information on the productivity shock θ with probability a. With probability $1 - a$, the agent remains uninformed. We assume that $\varphi(a)$ is increasing, convex, has non-negative third-derivative and satisfies usual Inada conditions ($\varphi'(a) \geq 0$, $\varphi''(a) > 0$, $\varphi'''(a) \geq 0$, $\varphi'(0) = 0$ and $\varphi'(1) = +\infty$). These assumptions ensure interior solutions to all problems below.

Importantly, the fact that we have interior solutions in the first stage implies that both regimes where either the agent gets informed or not have positive probabilities at the optimal contract.

[10] This quadratic specification is used for tractability and simplifies the expression of some of our results below without any loss of economic insights.

Preferences. The principal pockets the project return $y = (1+\theta)e$ and pays back a monetary transfer t to the agent for his services. The principal's payoff is quasi-linear in (y, t):

$$W(y, t) = y - t.$$

The agent's (ex post) payoff net of his costs at the operating and the information gathering stages simply writes as:

$$U(t, e, a) = t - \psi(e) - \varphi(a).$$

As usual in the agency literature, assuming risk aversion is a key ingredient to introduce an agency cost of information gathering. Compounding that cost with those arising at the operating stage imposes specific modeling choices. In that respect, we assume that the agent exhibits risk aversion in the sense of the dual theory of choice pushed forward by Quiggin (1982) and Yaari (1987).

In a binary environment where the agent may learn or not the productivity parameter, this amounts to assuming that the agent, when computing his ex ante payoff puts a greater probability to not learning information. Instead of weighting with probability a the event where information is gathered is learned, the agent only gives a probability ρa to that event.[11] The parameter $1-\rho \in (0, 1)$ may be viewed as the degree of (first-degree) risk aversion for such preferences.[12]

From a technical viewpoint, that modeling of risk aversion allows us to maintain the linearity of payoffs in a principal-agent model under various sorts of frictions (ex ante and ex post moral hazard, asymmetric information) that would lose tractability had we assumed a concave Bernoulli utility function to capture risk behavior. The impact of linearity would become more explicit below when computing expected payoffs.

Contracts. Contracts are designed to give to the agent the right incentives to exert efforts at each node of the game, i.e., both at the information gathering and at the operating stages.

For tractability, we assume that contracts are linear in the realized outcome y, and of the form $t(y) = \alpha y - \beta$. Such contract stipulates a fixed fee paid by the agent to access the production process and a piece rate bonus that determines how the proceeds

[11] More precisely, in our framework, the agent is actually facing a compounded lottery with a first component of randomness coming from his learning or not of productivity shocks and a second component coming from the realization of the productivity shock itself. We follow Epstein and Zin (1989) in embedding the dual theory of risk into a recursive framework. Although the agent is risk neutral with respect to productivity shocks whether information gathering has succeeded or not, he overweighs the event that information gathering has failed. Guriev (2001) points out that this assumption can be justified as a reduced form for a model where a risk neutral agent faces financial constraints ex ante, inducing dual risk aversion at this stage.

[12] Although beliefs are not viewed as being subjective in the framework of dual risk theory, this modeling is also akin to assuming that the principal and the agent do not share the same prior beliefs; the agent is actually more pessimistic than his principal on the possibility of success in information gathering. Contracting problems between a principal and an (often more optimistic) agent with different risk perceptions have been studied in other contexts by De la Rosa (2007), Jeleva and Villeneuve (2004), Villeneuve (2005), Eliaz and Spiegler (2008), Spinnewijn (2008) and Grubb (2009).

of production are shared.[13,14]

Even when focusing on linear contracts, it remains true that the principal should, in full generality, offer a whole menu of contracts. Indeed, such menu helps screening the agent according to information he may have learned. This information is two-dimensional. First, the agent knows whether he gets information on the productivity shock or not. Second, if he does get information, he learns the exact realization of productivity shock. A direct revelation mechanism must induce revelation of both pieces of information. In other words, the agent can lie on the productivity shock when knowing it but he may also pretend having observed a given productivity shock even if he has not (and *vice versa*).[15]

From the Revelation Principle,[16] there is no loss of generality in having the principal offer the menu of linear schedules $\{(t(y,\hat{\theta}))_{\hat{\theta}\in\Theta}, t_u(y)\}$ to the agent where $\hat{\theta}$ is an announcement on the realized shock that the agent may have learned. When informed on the productivity shock θ, incentive compatibility implies that the agent optimally picks the scheme corresponding to the observed productivity shock $t(y,\theta) = \alpha(\theta)y - \beta(\theta)$. The agent then delivers an output $y(\theta)$ by adjusting his operating effort accordingly and he receives the corresponding payment $t(y(\theta),\theta)$. When instead uninformed, incentive compatibility implies that the agent optimally chooses the scheme $t_u(y) = \alpha_u y - \beta_u$, exerts an operating effort without knowing the productivity shock and gets a payment $t_u(y)$ when output y realizes.

3. PRELIMINARIES

In this section, we establish a couple of benchmarks that allow us to better understand contracting distortions in more complex scenarios later on. In passing, we also define a few notions that are useful in those more complex scenarios.

Suppose that the principal first gathers information by himself and then chooses accordingly the operating efforts. Once he has learned that the productivity shock is θ, the principal chooses an operating effort $e^{FB}(\theta)$ that trades off the marginal cost of effort

[13] The incentive properties of linear contracts have been studied, among others, by Laffont and Tirole (1986), Rogerson (1987, 2003) in the flied of regulation and Melumad, Mookherjee and Reichelstein (1992) and Mookherjee and Reichelstein (2001) when applied to the internal organization of the firm. In pure moral hazard environments, Holmström and Milgrom (1987, 1991) and Baker (1992) unveil conditions under which such schemes are optimal.

[14] Our focus on linear schemes is motivated not only by tractability and relevance in practice but also because of their robustness to the introduction of noise in the measurement of output. Our model can thus be viewed as a reduced form for a more complete contracting environment where such noise would be made explicit. To illustrate, suppose for instance that the verifiable measure of output $\bar{y}$ is only a garbled signal of real output $(1+\theta)e$:

$$\bar{y} = (1+\theta)e + \varepsilon$$

where ε is a random variable with zero mean and the real line as support. Thanks to the linearity of contracts and the agent's payoff (even with dual risk aversion), this extra noise is washed out when taking expectations under various scenarios below. In other words, this more complex contracting environment can be readily identified with our model where $y = (1+\theta)e$ is viewed as the realized output. From a technical viewpoint, the benefits of such alternative presentation in the background could be to ensure that there is no moving support in the distribution of verifiable outputs. Yet, even in such environments, complex discontinuous contracts à la Mirrless (1999) could be constructed to sometimes approximate the first best.

[15] Our analysis will illustrate that, in many respects, this bi-dimensionality of private information is much easier to handle than in the standard screening framework developed for multiproduct monopolists by Rochet and Choné (1998) for instance.

[16] Myerson (1982).

with its marginal return, namely:

$$\psi'(e^{FB}(\theta)) = 1 + \theta \Leftrightarrow e^{FB}(\theta) = \lambda(1+\theta).$$

This yields an output that depends positively on the productivity shock:

$$y^{FB}(\theta) = \lambda(1+\theta)^2.$$

A better productivity shock increases the marginal benefit of effort. The optimal levels of effort and output are thus increasing with θ. Also, a more flexible technology is associated with a lower marginal cost of effort and induces greater effort and output.

If the principal is unlucky and no information is gathered, he chooses the effort at the operating stage without knowing the exact realization of the productivity shock and its consequences on realized output. The corresponding (uninformed) level of effort e_u^{FB} is then:

$$\psi'(e_u^{FB}) = E_\theta(1+\theta) \Leftrightarrow e_u^{FB} = e^{FB}(0) = \lambda. \tag{3.1}$$

Because of our assumption that shocks have zero mean, this optimal effort is exactly the same as if the principal had learned a productivity shock with mean value.

We can now characterize the optimal first-stage effort in information gathering, say a^{FB}, as a solution to the following problem:

$$a^{FB} \equiv \arg\max_a aE_\theta\left((1+\theta)e^{FB}(\theta) - \psi(e^{FB}(\theta))\right) + (1-a)\left(e^{FB}(0) - \psi(e^{FB}(0))\right) - \varphi(a).$$

Necessary and sufficient conditions for optimality yield:

$$\varphi'(a^{FB}) = \frac{\lambda\sigma^2}{2}. \tag{3.2}$$

To understand this expression, it is useful to define the (positive) value of information in the absence of any agency problem as the difference in the (expected) operating surpluses when information is learned and when it is not:

$$\mathcal{V}^{FB} \equiv E_\theta\left((1+\theta)e^{FB}(\theta) - \psi(e^{FB}(\theta))\right) - \left(e^{FB}(0) - \psi(e^{FB}(0))\right) = \frac{\lambda\sigma^2}{2}.$$

The optimality condition (3.2) just means that the marginal cost of effort in information gathering is equal to the value of information. As a result, more uncertainty on productivity shocks (i.e., σ^2 larger) which increases the value of information makes it also more attractive to gather information.

4. INFORMATION GATHERING WITH VERIFIABLE PRODUCTIVITY SHOCKS

Suppose now that the productivity shock θ can be verified so that there is no agency problem ex post. Intuition suggests that the principal can certainly align her objectives with those of the agent at the operating stage under those circumstances by means of fixed-price contracts. Yet, the issue of providing ex ante incentives in information gathering remains. This benchmark thus allows us to focus on the sole consequences of moral hazard in information gathering.

Ex ante participation and incentive constraints. Because θ is verifiable, the agent operates under the scheme $t(y,\theta) = \alpha(\theta)y - \beta(\theta)$ when the productivity shock θ has been learned and $t_u(y) = \alpha_u y - \beta_u$ when it has not. Keeping our previous notation, we denote the agent's ex post payoffs under each of those scenarios respectively as:

$$U(\theta) \equiv \max_e \alpha(\theta)(1+\theta)e - \frac{e^2}{2\lambda} - \beta(\theta) = \frac{\lambda}{2}\alpha^2(\theta)(1+\theta)^2 - \beta(\theta) \tag{4.1}$$

and

$$U_u \equiv \max_e \alpha_u E_\theta((1+\theta)e) - \frac{e^2}{2\lambda} - \beta_u = \frac{\lambda}{2}\alpha_u^2 - \beta_u. \tag{4.2}$$

Remember that the agent exhibits dual risk aversion and as such overweights the probability of being uninformed. Anticipating his future streams of payoffs, the agent participates when his ex ante payoff remains non-negative:

$$\rho a E_\theta(U(\theta)) + (1-\rho a)U_u - \varphi(a) \geq 0. \tag{4.3}$$

The left-hand side above is strictly concave in a and the following incentive constraint characterizes the (interior) first-stage effort:

$$\frac{1}{\rho}\varphi'(a) = E_\theta(U(\theta)) - U_u. \tag{4.4}$$

The right-hand side of (4.4) is again the value of information for the agent. It is clearly non-negative because the principal could just offer $t(y,\theta) \equiv t_u(y)$ under all circumstances.

Because $\rho \leq 1$, the agent views information gathering as less likely than what the principal thinks. As a result, he does not have enough incentives to gather information.

The principal's problem. To understand the principal's problem let us decompose the agent's ex ante payoff as follows:

$$\rho a E_\theta\left(U\left(\theta\right)\right) + \left(1-\rho a\right)U_u - \varphi(a) = \mathcal{U} - \mathcal{R} - \varphi(a)$$

where

$$\mathcal{U} \equiv a E_\theta\left(U\left(\theta\right)\right) + \left(1-a\right)U_u - \varphi(a) \text{ and } \mathcal{R} \equiv a\left(1-\rho\right)\left(E_\theta\left(U\left(\theta\right)\right) - U_u\right). \tag{4.5}$$

Up to the fixed fee β, $\mathcal{U}$ is exactly the principal's expected payoff from gathering information by himself. It is also equal to the agent's ex ante payoff had the latter been risk neutral. With risk aversion, the agent's payoff accounts for the (first-order) risk premium R that he incurs when information gathering remains uncertain and the agent's payoffs are different depending on whether information acquisition is successful or not.

Expressions (4.4) and (4.5) when taken in tandem show that ex ante moral hazard at the information gathering stage combined with (dual) risk aversion force the principal to compensate the agent for the (first-order) risk premium $\mathcal{R}$. That is, to induce information gathering the principal must leave some risk on the agent (the difference $E_\theta\left(U\left(\theta\right)\right) - U_u$ must remain positive), but this comes at the cost of having to pay an extra risk premium $\mathcal{R}$ to induce his participation.

Inserting the expression of effort coming from the incentive constraint (4.4), this risk premium can also be expressed in terms of a only as:

$$\mathcal{R} = a(1-\rho)(E_\theta(U(\theta)) - U_u) = a\varphi'(a)\left(\frac{1}{\rho} - 1\right). \tag{4.6}$$

This risk premium represents an agency cost for the principal. In the Appendix, we show that the principal's contracting problem is exactly the same as if the effort in information gathering were observable except for this extra agency cost $\mathcal{R}$. Importantly, increasing the agent's information gathering effort raises the risk premium and increases that agency cost.

Observe that the risk premium $\mathcal{R}$ as defined in (4.6) does not depend on effort at the operating stage. This highlights a dichotomy between the moral hazard incentive problems at the operating and the information gathering stages. Making the agent residual claimant for effort provision at the operating stage is thus clearly optimal:

$$\hat{\alpha}_u = \hat{\alpha}(\theta) = 1 \quad \forall\theta. \tag{4.7}$$

Such fixed-price contracts align the principal's objectives with those of the agent and induce the efficient effort at the operating stage:

$$\hat{e}_u = e^{FB}(0) \text{ and } \hat{e}(\theta) = e^{FB}(\theta) \quad \forall\theta. \tag{4.8}$$

Because of this dichotomy, incentives to gather information are solely provided by means of fixed fees. Next proposition summarizes the design of the optimal contract.

PROPOSITION 1 *Assume that productivity shocks are verifiable but effort in information gathering is not. A menu of fixed-price contracts,* $\hat{t}(y,\theta) \equiv \hat{t}_i(y) = y - \hat{\beta}$ *and* $\hat{t}_u(y,\theta) = y - \hat{\beta}_u$ *with* $\hat{\beta}_u < \hat{\beta}$*, is optimal. Moreover, the effort in information gathering, although positive, is lower than at the first best,* $0 < \hat{a} < a^{FB}$ *with:*

$$\varphi'(\hat{a}) + (1-\rho)\hat{a}\varphi''(\hat{a}) = \frac{\rho\lambda\sigma^2}{2}. \tag{4.9}$$

Remember that the agent acts as if he were overly pessimistic about remaining uninformed. This is a source of a misalignment with the principal's objectives which is captured by the agency cost $\mathcal{R}$ in (4.6). This agency cost is reduced by charging a lower fixed fee when the agent is uninformed ($\hat{\beta}_u < \hat{\beta}$), so as to lower the agent's effort in information gathering below the first-best level. Of course, this distortion is more pronounced when the agent exhibits greater risk aversion (ρ smaller), as then the risk premium is higher.

It is interesting to observe that if the principal did not set a lower fee when the agent is uninformed, the agent would choose a level of effort still below the first best a^{FB} but nevertheless too high. Indeed, with $\hat{\alpha}_u = \hat{\alpha}(\theta) = 1$and $\hat{\beta}_u = \hat{\beta}(\theta)$ for all θ, the agent is made residual claimant both at the information gathering and at the operating stages. The expressions of the agent's rents whether he is informed or not would be respectively given by:

$$U(\theta) \equiv \max_e (1+\theta)e - \psi(e) - \beta = \frac{\lambda}{2}(1+\theta)^2 - \beta,$$

and

$$U_u \equiv \max_e E_\theta(1+\theta)e - \psi(e) - \beta = \frac{\lambda}{2} - \beta.$$

Therefore, the agent's ex ante payoff would be:

$$\rho a E_\theta(U(\theta)) + (1-\rho a)U_u - \varphi(a) = \frac{\lambda}{2}(1+\rho a\sigma^2) - \varphi(a) - \beta.$$

This leads to a level of effort a_0 that would satisfy:

$$\frac{1}{\rho}\varphi'(a_0) = \frac{\lambda\sigma^2}{2} = \varphi'(a^{FB}). \tag{4.10}$$

With the fixed-price contracts $\hat{\alpha}_u = \hat{\alpha}(\theta) = 1$and $\hat{\beta}_u = \hat{\beta}(\theta)$, the risk premium $\mathcal{R}$ would then be proportional to the value of information since:

$$E_\theta\left(U\left(\theta\right)\right) - U_u = \frac{\lambda\sigma^2}{2} > 0.$$

However, because $\rho \in (0,1)$, the agent puts an excessive weight on not learning information. His effort is then below the first best, $a_0 < a^{FB}$. More importantly, from (4.9) and (4.10), this effort would also be higher than the level of effort that the principal wishes to induce, $\hat{a} < a_0$.Intuitively, although both the principal and the agent wants to reduce effort a to decrease the risk premium , they want to do so with different intensities. The agent has less incentives to reduce the risk premium by himself because, in the end, the principal ends up paying for that premium to ensure the agent's participation. To induce such lower effort $\hat{a} < a_0$, the principal has then to ask for a higher fee when the agent gets informed, making information gathering less attractive.

As we shall see in the sequel, this fee differential plays a critical role on the shape of the optimal contract when shocks are non-verifiable.

Quadratic specification. Specializing our analysis to the case of a quadratic disutility of effort, namely $\varphi(a) = \frac{\varphi a^2}{2}$, and assuming that φ is large enough to ensure that the optimal first-stage efforts below remain in $[0,1]$ (i.e., $a^{FB} = \frac{\lambda\sigma^2}{2\varphi} < 1$) those efforts can be expressed in closed forms as:

$$\hat{a} = \frac{\rho\lambda\sigma^2}{2\varphi(2-\rho)} < a_0 = \frac{\rho\lambda\sigma^2}{2\varphi} < a^{FB} = \frac{\lambda\sigma^2}{2\varphi}. \tag{4.11}$$

■

5. INFORMATION GATHERING WITH NON-VERIFIABLE PRODUCTIVITY SHOCKS: INCENTIVE FEASIBLE ALLOCATIONS

We now turn to the more complex environment where information, when gathered, can still be manipulated by the agent. Clearly, the solution found in Proposition 1 is no longer feasible. Charging a lower fee $\hat{\beta}_u < \hat{\beta}(\theta)$ when information has not been gathered, while keeping the same constant bonus $\hat{\alpha}_u = \hat{\alpha}(\theta) = 1$, would violate incentive compatibility: The agent would always report being uninformed so as to minimize his payment, and this would destroy his incentive to gather information. Raising the fee when the agent is uninformed, so as to have $\hat{\beta}_u = \hat{\beta}(\theta)$, would induce truthful revelation but then, as we have seen, the level of effort chosen by the agent a_0 would be excessive.

In such context, we expect that the principal can obtain a higher payoff by modifying not only fees but also bonuses in order to induce information gathering and revelation.

This rough intuition already suggests that the dichotomy between ex post and ex ante incentives will now disappear.

We first describe the set of incentive feasible allocations when all incentive problems taking place both ex ante (moral hazard on a) and ex post (moral hazard on e and private information on θ) have to be considered.

Ex post private information. The agent must be induced to reveal that he is informed in which case he must also report the value of the productivity shock he has learned. Turning to this second aspect of the revelation problem, let again denote by $U(\theta)$ the information rent of an agent who has private information on the shock θ. Incentive compatibility for this screening side of the contracting problem amounts to:

$$U(\theta) \equiv \max_{(\hat{\theta},e)} \alpha(\hat{\theta})(1+\theta)e - \psi(e) - \beta(\hat{\theta}). \tag{5.1}$$

This compact expression encompasses already all incentive constraints preventing an informed agent to lie on the productivity shock he has observed. The agent picks then within the schedule $t(y,\theta)$ his most preferred output $y(\theta) = (1+\theta)e(\theta)$ where the operating effort reflects actual bonus:

$$e(\theta) = \lambda(1+\theta)\alpha(\theta) \quad \forall\theta \in \Theta. \tag{5.2}$$

We immediately get the standard characterization of the rent profile $U(\theta)$.

LEMMA 1 *$U(\theta)$ is absolutely continuous, non-decreasing and convex with*

$$U(\theta) = U(0) + \lambda \int_0^\theta (1+x)\alpha^2(x)dx \quad \forall\theta \in \Theta. \tag{5.3}$$

The bonus $\alpha(\theta)$ is non-decreasing in θ.

To prevent an informed agent from pretending being uninformed and choosing an output y along the scheme $t_u(\cdot)$, the following incentive compatibility constraint must also hold:

$$U(\theta) \geq V(\theta) \equiv \max_e \alpha_u(1+\theta)e - \beta_u - \psi(e) \quad \forall\theta \in \Theta. \tag{5.4}$$

Note that the reservation payoff $V(\theta)$ is also an implementable profile corresponding to a fixed bonus α_u. As such, it satisfies Lemma 1 and thus:

$$V(\theta) = V(0) + \frac{\lambda\alpha_u^2}{2}\theta(2+\theta) \quad \forall\theta \in \Theta. \tag{5.5}$$

Instead, when uninformed, the agent must prefer to pick the scheme $t_u(\cdot)$ and choose his effort not knowing the productivity shock θ rather than pretending being informed on a shock $\hat{\theta}$. This yields the following incentive constraint:

(5.6)
$$U_u = \max_e \alpha_u e E_\theta(1+\theta) - \beta_u - \psi(e) \equiv V(0) \geq \max_{(e,\hat{\theta})} \alpha(\hat{\theta})eE_\theta(1+\theta) - \beta(\hat{\theta}) - \psi(e) \equiv U(0).$$

The right-hand side above simply comes from observing that the mean productivity shock θ is zero. The uninformed agent can always behave as being hit by such mean-value shock.

Simplifying the incentive feasible set. Reciprocally, the agent who has learned that the shock is just the mean of the distribution could behave as being uninformed. When uninformed, the agent thus gets the same payoff as when he knows that the productivity shock is the mean of the distribution. Next Lemma immediately follows.

LEMMA 2 *$U(\theta)$ is more convex than $V(\theta)$*[17] *with:*

$$U(0) = V(0) = U_u. \tag{5.7}$$

Since $U(\theta)$ is absolutely continuous and convex, it admits at least right- and left-derivatives at 0 where its subgradient is $\partial_\theta U(0) = [\lambda\alpha(0^-), \lambda\alpha(0^+)]$. Because $V(\theta)$ is also convex, minorizes $U(\theta)$ at $\theta = 0$ and has derivative α_u at that point, it must be that:

$$\alpha(0^-) \leq \alpha_u \leq \alpha(0^+). \tag{5.8}$$

If $\alpha(\cdot)$ is continuous at 0, condition (5.8) is akin to a "smooth-pasting" requirement, i.e., the rent profile $U(\theta)$ tangentially touches $V(\theta)$ at 0. Smooth-pasting holds at the optimal contract, as we will see below.

Turning now to the information gathering stage, using the characterization of the rent profile $U(\theta)$ obtained from Lemma 1 and integrating by parts, we get the following expression of the ex ante moral hazard incentive constraint.

LEMMA 3 *The information gathering incentive constraint (4.4) can be rewritten as:*

$$\frac{1}{\rho}\varphi'(a) = \lambda E_\theta \left(\frac{1_{\theta>0} - F(\theta)}{f(\theta)} (1+\theta)\alpha^2(\theta) \right) \tag{5.9}$$

where $1_{\theta>0} = \begin{cases} 1 & \textit{if } \theta > 0 \\ 0 & \textit{otherwise} \end{cases}$.

The incentive constraint (5.9) has important implications for the shape of the optimal contract. When information remains private to the agent, his incentives to gather information are driven by his information rents obtained at the operating stage. Equation (5.9) can be interpreted as earlier; the right-hand side is again the value of information from the agent's viewpoint. This quantity depends now on the whole profile of bonuses at the operating stage. It implicitly encapsulates the fact that the profile of ex post rents also satisfies incentive compatibility conditions. The left-hand side is again the agent's marginal cost of effort conveniently pondered by his risk-aversion parameter.

An important implication of (5.9) is that first- and second-stage efforts are now linked. There is no longer any dichotomy between ex post and ex ante incentives when the agent produces non-verifiable information. Bonuses at the operating stage must be modified to boost information gathering.

[17] Formally, our definition of being "more convex" means that the convex epigraph of U is a subset of the convex epigraph of V: $epi \quad U = \{y| \quad y \geq U(\theta), \theta \in \Theta\} \subset epi \quad V = \{y| \quad y \geq U(\theta), \theta \in \Theta\}$with (5.7) ensuring that those epigraphs have a common point. This definition is consistent with other notions of relative convexity like in Palmer (2003).

6. INFORMATION GATHERING WITH NON-VERIFIABLE PRODUCTIVITY SHOCKS: OPTIMAL CONTRACTS

We are now ready to write the principal's optimization problem when the agent has ex post private information on productivity shocks. Ex post private information puts quite a bit of structure on this problem. First, it implies that $\alpha(\theta)$ must remain non-decreasing and second, that incentives at the information gathering stage depend also on the power of incentives at the operating stage as shown through (5.9).

Taking into account the expression of operating efforts in terms of bonuses (see the corresponding expressions (A.1) in the Appendix), we may write this optimization problem as:

$$(\mathcal{P}): \max_{(\alpha(\theta),\alpha_u,a)} a\lambda E_\theta\left((1+\theta)^2\left(\alpha(\theta)-\frac{1}{2}\alpha^2(\theta)\right)\right)+(1-a)\lambda\left(\alpha_u-\frac{1}{2}\alpha_u^2\right)-\varphi(a)$$

$$\underbrace{-a\varphi'(a)\left(\frac{1}{\rho}-1\right)}_{\text{Risk-premium}}$$

subject to $\alpha(\theta)$ being non-decreasing and satisfying (5.8); and (5.9).

From Lemmas 2 and 3 above, these constraints are sufficient to fully describe the constrained set. We now proceed in two steps. First, we replace (5.8) and the requirement of monotonicity of $\alpha(\cdot)$ by the weaker constraints:

$$\alpha(\theta)\le\alpha_u \quad \forall\theta\le 0;\ \text{and}\ \alpha_u\le\alpha(\theta) \quad \forall\theta\ge 0. \tag{6.1}$$

Second, we optimize the new problem $(\mathcal{P}^*)$ so obtained with (5.9) and (6.1) as constraints. Lastly, we will show that the solution to this relaxed problem is such that $\alpha(\cdot)$ is indeed monotonically increasing.

6.1. *Linear Contracts*

Before proceeding to a full-fledged optimization, and for the sake of building preliminary intuition for our more general findings, consider the simple case where the principal offers a single "pooling" scheme which applies whether the agent gets informed or not, $t_i(y,\theta)=t_u(y)\equiv\alpha y-\beta$ for some given pair $(\tilde{\alpha},\tilde{\beta})$. Since it is independent on whether information has been gathered or not and on the announcement on the productivity shock, that scheme obviously solves the screening problem. The value of this preliminary analysis is thus to highlight the consequences of the agency conflict in information gathering without perturbating this analysis with any screening consideration. Second, we will see below that a fixed linear contract is indeed *part* of the optimal linear contract, at least over some range of possible realizations of the productivity shocks, even when more general schemes are allowed. Finally, this analysis is also useful because linear contracts have been the workhorse of a huge applied literature based on moral hazard considerations and it may be of interest of knowing how such contracts behave in slightly different informational settings.

Integrating by parts on $[-\delta,0]$ and $[0,\delta]$ respectively yields $E_\theta\left(\frac{1_{\theta>0}-F(\theta)}{f(\theta)}\right)=\frac{\sigma^2}{2}$. Using this expression, we may thus rewrite the incentive constraint (5.9) in the case of a linear contract with slope α as:

$$\frac{1}{\rho}\varphi'(a)=\frac{\lambda}{2}\sigma^2\alpha^2. \tag{6.2}$$

A lower (fixed) bonus reduces the agent's value of information and his incentives to gather information. It thus better aligns the agent's and the principal's objectives.

PROPOSITION 2 *Assume that the principal offers the same linear scheme whether the agent is informed or not.*

1. *The agent is made only partially claimant:*

$$\tilde{\alpha} < 1. \tag{6.3}$$

2. *The optimal effort in information gathering $\tilde{a}$ is lower than what the agent would choose on his own:*

$$\tilde{a} < a_0. \tag{6.4}$$

Remember from Section 4 that, when operating under a fixed-price contract and being residual claimant for his operating effort, the agent has too much incentives to invest in information gathering compared to what the principal would like to induce. To curb those incentives, the principal makes the agent's payoff at the operating stage less sensitive to his information. When restricted to use contracts with constant bonuses, this is obtained by making the agent only partially residual claimant for his effort at the operating stage so that $\tilde{\alpha} < 1$.

Quadratic specification (continued). To get more explicit expressions for the optimal bonus, we come back to the case of a quadratic disutility of effort, namely $\varphi(a) = \frac{\varphi a^2}{2}$ where we again impose the condition $a^{FB} = \frac{\lambda\sigma^2}{2\varphi} < 1$ to ensure that optimal efforts in information gathering always lie in $[0, 1]$. In the Appendix, we show that the first-order condition for optimality of $\tilde{\alpha}$ becomes:

$$1 - \tilde{\alpha} = \frac{\frac{\rho\lambda\sigma^4\tilde{\alpha}^2}{2\varphi}}{1 + \frac{\rho\lambda\sigma^4\tilde{\alpha}^2}{2\varphi}}\left((3-\rho)\,\tilde{\alpha} - 2\right). \tag{6.5}$$

The right-hand side is a function of $\tilde{\alpha}$ on $(0, 1)$, which starts negative but becomes positive and increasing for $\tilde{\alpha} \geq \frac{2}{3-\rho}$. The left-hand side is instead strictly decreasing in $\tilde{\alpha}$ and zero at $\tilde{\alpha} = 1$. The solution to (6.5) is thus unique in the interval $\left(\frac{2}{3-\rho}, 1\right)$.

Equation (6.5) is highly nonlinear which makes it difficult to get general comparative statics, especially when comparing $\hat{a}$ and $\tilde{a}$. To nevertheless better understand how risk aversion modifies first-stage effort, let suppose that $\rho = 1 - \epsilon$ for ϵ close enough to zero, which means that the agent is almost risk neutral and should almost made residual claimant at both stages. Taylor expansions help us to uncover variations in optimal bonuses in that neighborhood.

COROLLARY 1 *Assume the disutility of effort is quadratic, namely $\varphi(a) = \frac{\varphi a^2}{2}$ with $a^{FB} = \frac{\lambda\sigma^2}{2\varphi} < 1$ and that the principal offers linear contracts. For $\rho = 1 - \epsilon$ for ϵ close enough to zero, the agent exerts more effort in information gathering when he has private information on productivity shocks than when these shocks are verifiable:*

$$\hat{a} < \tilde{a}. \tag{6.6}$$

There are two effects at play in determining the optimal first stage effort under a linear contract. Each of those effects can be unveiled by a careful look at the marginal benefits and costs of increasing the first-stage effort. On the benefits side, the value of information for the principal must now be evaluated with the surplus induced by a fixed bonus at the operating stage, namely:

$$\lambda E_\theta\left(\left((1+\theta)^2-1\right)\left(\alpha-\frac{1}{2}\alpha^2\right)\right)=\frac{\lambda\sigma^2}{2}\left(\alpha-\frac{1}{2}\alpha^2\right).$$

Because the principal's value of information is lower when the agent only receives a fraction of the surplus (i.e., when $\alpha < 1$) and reduces accordingly his efforts at the operating stage below the first-best level, this benefit is less than when the agent operates under a fixed-price contract ($\alpha = 1$). This first effect reduces information gathering effort below $\hat{a}$.

On the cost side, counting both the disutility of effort for the agent and the risk premium, the marginal cost of effort writes as

$$\varphi'(a)+\left(\frac{1}{\rho}-1\right)\frac{d}{da}(a\varphi'(a))=\frac{\varphi'(a)}{\rho}+\frac{1-\rho}{\rho}a\varphi''(a)$$

where of course a is positively linked to the bonus through the incentive constraint (6.2). From the analysis of the benchmark case where productivity shocks verifiable, we already know that the principal wants to reduce the agent's effort in information gathering to save on the risk premium. With a single linear contract, a differential on fixed fees can no longer be used and (6.2) show that this effort reduction can only come from also reducing the bonus below 1. But because such move also reduces the benefits as we have just seen, the principal finds this strategy less attractive which explains (6.6).

6.2. *General Schemes*

Let us now turn to the more general analysis where the principal is no longer restricted in the menus of screening options she can offer. The next proposition characterizes optimal bonuses and effort at the optimal contract under that scenario. To help presentation, let us define $\theta^* = \Omega(a)$ as the implicit solution in $(0, \delta)$ to the following equation:[18]

$$R(\theta^*)=\frac{a\int_{-\delta}^{\theta^*}(1_{\theta>0}-F(\theta))(1+\theta)d\theta}{1-a+a\int_{-\delta}^{\theta^*}(1+\theta)^2f(\theta)d\theta}. \tag{6.7}$$

This variable helps to distinguish two regimes in the optimal contract where either pure asymmetric information or pure moral hazard are the driven forces.

In the sequel, and for the sake of simplicity, we will assume quasi-concavity of the principal's problem.[19]

[18] We prove in the Appendix that this solution is indeed unique.

[19] Although much has been written on the validity of the first-order approach for the agent's problem in a moral hazard context, little is known on the (quasi-)concavity of the principal's problem itself. For a similar remark, see Szalay (2009). Nevertheless, a weaker but still true statement of Theorem 1 below is available if we dispense with the assumption of quasi-concavity and just say that the contract exhibited in Theorem 1 improves on a fixed-price contract, i.e., choosing $\alpha_u^* = 1 - \eta^*$ for η^* small enough and a menu of screening bonuses on the upper tail of the distribution strictly improves the principal's payoff.

Theorem 1 *The optimal contract and effort in information gathering can be characterized as follows.*

1. *For $\theta \in [-\delta, \Omega(a^*)]$, the agent is made only partially claimant. The incentive bonus is independent of θ and identical whether the agent gets such "bad news" or remains uninformed:*

$$\alpha^*(\theta) = \alpha_u^* < 1.$$

2. *For $\theta \in [\Omega(a^*), \delta]$, the incentive bonus $\alpha^*(\theta)$ is increasing with the productivity shock and greater than when the agent remains uninformed:*

$$\alpha_u^* \leq \alpha^*(\theta) = \frac{\alpha_u^*}{\alpha_u^* + (1-\alpha_u^*)\frac{R(\theta)}{R(\Omega(a^*))}} \leq 1 \quad (6.8)$$

with an equality on the r.h.s at $\bar{\theta}$ only and on the l.h.s at $\Omega(a^*)$ only.

3. *The optimal information gathering effort a^* is less than what the agent would choose on his own:*

$$a^* < a_0.$$

Screening for "good news." Using a menu of linear contracts is of course *always* optimal and strictly dominates the restricted option consisting in offering a single linear scheme. Indeed, a menu always helps the principal to better screen the agent according to the productivity shock he has learned. Of course, this does not come as a surprise. What is more surprising is precisely when and how those extra possibilities are used. Quite strikingly, the optimal scheme entails much bunching on the lower tail $[-\delta, \theta^*]$ of the distribution while screening only arises for sufficiently "good news", i.e., on an interval $[\theta^*, \delta]$. On the bunching area, the agent gets the same bonus as when uninformed.

Lemma 2 provides the key intuition to understand these distortions. From there, we already know that a more convex profile of rents at the operating stage boosts the agent's incentives to gather information. To increase convexity, the principal would like to give more of the surplus to the agent as he announces having observed a rather good shock and insulate him from productivity shocks when he announces instead a bad shock. At the extreme, leaving all revenues to the agent for "good news" (i.e., the bonus is $\alpha(\theta) \equiv 1$ for $\theta > 0$) and paying only a fixed fee for "bad news" (i.e., the bonus is $\alpha(\theta) \equiv 0$ for $\theta < 0$) would give a very convex profile of rents which certainly boosts incentives to gather information. Of course, this benefit on ex ante incentives must be traded off with the cost that such profile brings in terms of insufficient rent extraction.

As far as screening is concerned, this menu actually leaves rent to an informed agent much beyond what he gets by remaining uninformed. Of course, the agent could also operate under a single linear scheme (as described in Proposition 2). Such contract trivially solves the screening problem but the rent profile it implements is "too flat" to provide enough incentives to gather information. The optimal mechanism is somewhere in between the "convex" two-item menu and the "everywhere flat" option we just described.

To figure out how the optimal contract is constructed, we need to come back again on Lemma 2 which also imposes another incentive compatibility requirement: The agent who has failed to get an informative signal should get the same payoff as if he was informed on a mean-value shock. This condition acts as an anchor which, together with incentive compatibility, fully defines the whole profile of rents of an informed agent at the operating stage. On both sides of that anchor, the agent faces nevertheless different incentives to manipulate information. Those incentives in turn requires different kinds of distortions.

A careful analysis of those incentives to manipulate information relies on the specificities of the informational problem under scrutiny, and especially its bi-dimensionality. On the one, the informed agent may want to claim that returns will be lower than they are really: a *"worsening" effect*. So doing, he pays back a lower fee to the principal and earns some rent. On the other hand, the informed agent may also want to claim he is uninformed: a *"hiding" effect*. This strategy is particularly attractive when the principal asks for a lower fee if the agent claims being uninformed, a feature of the optimal contract that we found optimal when shocks are verifiable. This *"hiding" effect* and the *"worsening" effect* reinforce each other when shocks are "good news" while they are conflicting with "bad news."

To illustrate, if "good news" have been observed, the agent would like to pretend that those shocks are actually closer to the mean value than what they really are. To limit those incentives to manipulate information downwards, bonuses must be reduced (i.e., setting $\alpha(\theta) \leq 1$ for such "good news"). This distortion is captured on (6.8). On the upper tail of the distribution, bonuses are positively correlated with shocks as shown by this latter formula. This monotonicity preserves enough convexity of the rent profile so as to boost information gathering. In other words, the screening distortion then also helps to ease the ex ante moral hazard problem.

If instead "bad news" have been observed, reporting a mean value amounts to lying upwards while the *"worsening" effect* still calls for lying downwards. Incentives to lie upwards arise because the principal wishes to reduce effort in information gathering by increasing the fixed fee differential between an informed and an uninformed agent, as in the case of verifiable productivity shocks. This incentive to report ignorance would be curbed by having increasing bonuses on the lower tail, since this would increase the cost for a low-productivity agent of reporting ignorance (and thus the mean value of productivity) but such distortions are not compatible with those needed to control the *"worsening" effect* that calls for lower bonuses. To reconcile those countervailing requirements, the principal just makes information revelation irrelevant on the lower tail by offering the same scheme whether the agent is uninformed or has learned sufficiently "bad news".

The resulting profile of rents is thus "convex" on the upper tail and flatter on the lower tail. The agent is thus somewhat protected against the risk of learning bad news. He acts upon the value of productivity shock he has learned on the upper tail and as if he had not learned anything on the lower tail.

Bonuses and fixed fees. On the screening area $[\Omega(a^*), \delta]$, bonuses are strictly increasing with, in the limit, $\alpha(\delta) = 1$; a *"no-distortion at the top"* result which is familiar from the screening literature. The agent who announces having observed the most productive shock is thus made residual claimant for his performances. Less attractive reports are followed by lower bonuses.

As far as fixed fees are concerned, the first-order necessary (and here sufficient) conditions for optimality at the revelation stage imply that the fixed fee $\beta^*(\theta)$ solves the following differential equation:

$$\dot{\beta}^*(\theta) = \lambda(1+\theta)^2 \dot{\alpha}^*(\theta).$$

From Proposition 1, it follows that $\beta^*(\theta)$ is increasing on the upper tail $[\Omega(a^*), \delta]$ and constant on the lower tail $[-\delta, \Omega(a^*)]$, $\beta^*(\theta) = \beta^*_u$.

On the other hand, the agent's participation constraint being binding at the optimum, we get:

$$\rho a^* E_\theta(U^*(\theta)) + (1 - \rho a^*)U^*(0) = \varphi(a^*).$$

Rearranging this expression, using (4.4) and (5.7), we obtain the expression of β_u^* :

$$U^*(0) = \varphi(a^*) - a^*\varphi'(a^*) = \frac{\lambda}{2}\alpha_u^{*2} - \beta_u^* < 0$$

where the right-hand side equality follows from direct computations of the agent's payoff when he remains uninformed and the inequality from $\varphi(\cdot)$ convexity.

This condition together with the continuity and the monotonicity of the optimal profile $U^*(\cdot)$ implies that $U^*(\theta) \leq 0$ on a whole interval $[-\delta, \theta_1]$ for some $\theta_1 > 0$. Only if the agent reports sufficiently good news will he make a positive rent. The optimal contract is *most of the time* biased towards negative payoff and the agent's ex ante participation constraint only holds because his payoff increases sufficiently quickly with θ.

Value of communication. Laffont and Tirole (1986) have shown how a menu of linear schemes can be used, under some conditions, to implement the optimal nonlinear contract when the information structure is exogenous. We may ask here whether the reverse finding holds when the information structure is endogeneized. In fact, the principal would achieve a very different payoff by offering the (nonlinear) upper envelope of the linear schemes that we described above. In other words, there is a positive value of communication between the principal and the agent and using a menu of linear options strictly dominates the single offer of its upper envelope.[20]

To see why, let us consider the upper envelope of the optimal menu of linear schemes found above which is defined as:

$$T^*(y) = \max_{\theta \geq \Omega(a^*)} \alpha^*(\theta)y - \beta^*(\theta).$$

As a maximum of linear functions which are increasing, $T^*(y)$ is itself increasing, convex with maximal slope 1 at the highest possible output (again, a consequence of the familiar *"no distortion at the top"* property) and minimal slope α_u^* on the upper tail.

To see how the principal may lose from using this convex envelope, consider an uninformed agent facing that nonlinear scheme $T^*(y)$ instead of the linear payment rule $\alpha_u^* y - \beta_u^*$ that he should be taking within the menu of linear options. Incentives at the operating stage are rather different with those two contracts. Using a (necessary and sufficient first-order) condition for optimality and taking into account that $T^{*\prime}(y) \geq \alpha_u^*$, the uninformed agent now chooses an effort e_u at the operating stage that satisfies:

$$\psi'(e_u) = E_\theta((1+\theta)T^{*\prime}((1+\theta)e_u)) \geq \alpha_u E_\theta(1+\theta) = \alpha_u^* \Rightarrow e_u > e_u^*$$

where e_u^* is the operating effort under the optimal linear scheme $\alpha_u^* y - \beta_u^*$. Hence, an uninformed agent would have too much incentives at the operating stage with the nonlinear scheme. Averaging over possible realizations of the productivity shocks, the marginal value of increasing effort is too "high" compared with those arising with linear contracts.

[20]This finding is reminiscent of the work of Melumad and Reichelstein (1989), although these authors address this issue in the context of a risk-neutral agent who has exogenous private information on a technological parameter before exerting a non verifiable effort whose return is random.

A contrario, when informed on the realization of the shock θ, the agent exerts the same effort when facing $T^*(y)$ or when taking within the menu of linear contracts his most preferred option $\alpha^*(\theta)y - \beta^*(\theta)$ (for $\theta \geq \Omega(a^*)$). Hence, we may conclude that the value of communication through a menu of contracts comes precisely from the possibility that the agent have to (truthfully) claim that he is uninformed and operates thereby on a less powered incentive scheme.

7. CONCLUSION

We have analyzed a model where an agent may invest to gather private information on a productivity shock that will affect the marginal benefit of his operating effort. Both polar models of hidden action and hidden information arise endogenously with positive probabilities as a result of the agent's investment in (imperfect) information gathering. The optimal contract merges well-known features from both paradigms but under quite specific circumstances. Screening arises on the upper tail of the distribution while, as a result of countervailing incentives in information gathering and screening, incentives to operate do not depend on realized shock on the lower tail of the distribution.

This finding suggests that any empirical analysis of contracts that would aim at disentangling whether hidden action or hidden information is the true driver of contract design should be particularly cautious in dealing with an endogenity problem. Hidden information-like contracts are expected in favorable environments and should come with higher average productivity, and high-powered incentives. Hidden action-like contracts are more likely with low powered incentives and bad average performances. Although this insight strikes us as being particularly important from an applied perspective, we are unaware of any research on that front. Probably, this is explained by the scarcity of analysis where information structures are endogenized.

On a more theoretical front, a few important extensions of our analysis would be worth to pursue. A first one would be to also introduce (dual) risk aversion at the operating stage so that expected rents at this stage are evaluated with extra weights on "bad news". This effect would further reduce the agent's expected payoffs when informed and thus his incentives to gather information in the first place. Formally, it is thus akin to a decrease in the parameter ρ in our model which assumes risk-neutrality at the operating stage. The impact of such risk aversion is thus to amplify the main features of the optimal contract we exhibited above.

Second, it would be interesting to develop our results further by analyzing the implications for organization design of the spillover effect that the transfer of operating risk generates on ex ante incentives for information acquisition. On the one hand, bundling the tasks of information gathering and operating brings some benefits because, on the upper tail of the distribution, screening convexifies payoffs and facilitates information gathering. On the other hand, those two tasks are also conflicting on the lower tail. There, splitting tasks is beneficial.[21] Averaging those two effects, the trade-off between bundling or not those tasks remain unclear at this stage.

We have focused on a case where the mapping between effort and performance is exogenously determined. In practice, when the design of a project is delegated to the agent, his choice may increase the likelihood of high productivity shocks or reduce uncertainty surrounding the implementation of the project. It would be interesting to extend our

[21]Specific organizational issues raised when planning and implementation are either merged or split between two separate agents have been addressed in Lewis and Sappington (1997), Khalil, Kim and Shin (2006), Krähmer and Strausz (2009) and Hoppe and Schmitz (2012).

analysis in this direction, studying the shape of the optimal contract when the agent must be given incentives to innovate in project design and to exert information gathering effort.

Finally, we have restricted our analysis to the case of menus of linear contracts, only briefly touching at the possible benefits that expanded communication through menus may yield. It remains to elaborate further and assess whether more complex menus of nonlinear contracts could help as well.

REFERENCES

Arrow, K., 1986, "Agency and Markets," in *Handbook of Mathematical Economics,* (K. Arrow and M. Intriligator, Eds.) Vol. 3, North-Holland, Amsterdam.

Baker, G., 1992, "Incentive Contracts and Performance Measurement," *Journal of Political Economy*, 100: 598-614.

Baker, G. and B. Jorgensen, 2003, "Volatility, Noise and Incentives," Mimeo Harvard University.

Baron, D. and R. Myerson, 1982, "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50: 911-930.

Bergemann, D. and J. Välimäki, 2002, "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70, 1007-1033.

Bergemann, D. and J. Välimäki, 2006, "Information in Mechanism Design," in *Proceedings of the 9th World Congress of the Econometric Society* (R. Blundell, W. Newey, and T. Persson, Eds.), Cambridge University Press, New York, 186-221.

Compte, O. and P. Jehiel, 2008, "Gathering Information Before Signing a Contract: A Screening Perspective", *International Journal of Industrial Organization*, 26: 206-212.

Cremer, J. and F. Khalil, 1992, "Gathering Information Before Signing a Contract," *American Economic Review*, 82: 566-578.

Cremer, J. and F. Khalil, 1994, "Gathering Information Before a Contract is Offered," *European Economic Review*, 38: 675-682.

Cremer, J., F. Khalil and J.C. Rochet, 1998a, "Strategic Information Gathering before a Contract is Offered," *Journal of Economic Theory*, 81: 163-200.

Cremer, J., F. Khalil and J.C. Rochet, 1998b, "Contracts and Productive Information Gathering," *Games and Economic Behavior*, 25: 174-193.

Cremer, J., Y. Spiegel and C. Zheng, 2009, "Auctions with Costly Information Acquisition," *Economic Theory*, 38: 41-72.

Dai, C., T. Lewis and G. Loppomo, 2006, "Delegating Management to Experts," *The RAND Journal of Economics*, 37: 503-520.

De La Rosa, L., 2011, "Overconfidence and Moral Hazard," *Games and Economic Behavior,* 73: 429-451.

Demski, J. and D. Sappington, 1987, "Delegated Expertise," *Journal of Accounting Research*, 25: 68-89.

Eliaz, K. and R. Spiegler, 2008, "Consumer Optimism and Price Discrimination," *Theoretical Economics,* 3: 459-497.

Epstein, L. and S. Zin, 1990, "First-Order Risk Aversion and the Equity Premium Puzzle," *Journal of Monetary Economics*, 26: 387-407.

Gerardi, D. and L. Yariv, 2008, "Information Acquisition in Committees," *Games and Economic Behavior,* 62: 436-459.

Gershkov, A. and B. Szentes, 2009, "Optimal Voting Schemes with Costly Information Acquisition," *Journal of Economic Theory,* 144: 36-68.

Gromb, D. and D. Martimort, 2007, "Collusion and the Organization of Delegated Expertise," *Journal of Economic Theory*, 137: 271-299.

Grossman, S. and O. Hart, 1983, "An Analysis of the Principal-Agent Problem," *Econometrica*, 51: 7-45.

Grubb, M., 2009, "Selling to Overconfident Consumers," *American Economic Review*, 99: 1770-1807.

Guriev, S., 2001, "On Microfoundations of the Dual Theory of Choice," *The Geneva Papers on Risk and Insurance Theory*, 26: 117-113.

Hart, O. and B. Holmström, 1987, "The Theory of Contracts," in *Advances in Economic Theory, Proceedings of the 1985 World Congress of the Econometric Society*, T. Bewley ed., Cambridge, Cambridge University Press

Holmström, B. and P. Milgrom, 1987, "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica* 55: 303-328.

Holmström, B. and P. Milgrom, 1991, "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, and Organization,* 7: 24-52.

Hoppe, E. and P. Schmitz, 2012, "Public-Private Partnerships versus Traditional Procurement: Innovation Incentives and Information Gathering," forthcoming *The RAND Journal of Economics.*

Iossa, E. and P. Legros, 2004, "Auditing and Property Rights", *The RAND Journal of Economics*, 35: 356-372.

Jeleva, M. and Villeneuve, B., 2004, "Insurance Contracts with Imprecise Probabilities and Adverse Selection," *Economic Theory*, 23: 777-794.

Kessler, A., 1998, "The Value of Ignorance", *The RAND Journal of Economics*, 29: 339-354.

Khalil, F., D. Kim and D. Shin, 2006, "Optimal Task Design: To Integrate or Separate Planning and Implementation?" *Journal of Economics and Management Strategy*, 15: 457-478.

Kofman, F. and J. Lawarée, 1993, "Collusion and Hierarchical Agency", *Econometrica*, 61: 629-656.

Krähmer, D. and R. Strausz, 2011, "Optimal Procurement Contracts with Pre-Project Planning," *Review of Economic Studies*, 78: 1015-1041.

Laffont, J.J. and D. Martimort, 2002, *The Theory of Incentives: The Principal-Agent Model*, Princeton University Press.

Laffont, J.J. and J. Tirole, 1986, "Using Costs to Regulate Firms," *Journal of Political Economy*, 94: 614-641.

Lambert, R., 1986, "Executive Effort and Selection of Risky Projects," *The RAND Journal of Economics,* 17: 77-88.

Lewis, T. and D. Sappington, 1989, "Countervailing Incentives in Agency Problems," *Journal of Economic Theory* , 49: 294-313.

Lewis, T. and D. Sappington, 1991, "All or Nothing Information Control," *Economics Letters*, 37: 111-113.

Lewis, T. and D. Sappington, 1993, "Ignorance in Agency Problems," *Journal of Economic Theory*, 61: 169-183.

Lewis, T. and D. Sappington, 1997, "Information Management in Incentive Problems," *Journal of Political Economy*, 105: 796-821.

Malcomson, J., 2009, "Principal and Expert Agent," *The B.E. Journal of Theoretical Economics,* 9 (Online) 1935-1704, DOI: 10.2202/1935-1704.1528.

Melumad, N., D. Mookherjee and S. Reichelstein, 1992, "A Theory of Responsibility Centers ," *Journal of Accounting and Economics,* 15: 445-484.

Melumad, N. and S. Reichelstein, 1989, "Value of Communication in Agency," *Journal of Economic Theory,* 47: 334-368.

Mirrlees, J., 1971, "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies,* 38: 175-208.

Mirrlees, J., 1999, "The Theory of Moral Hazard and Unobservable Behaviour: Part I," *Review of Economic Studies,* 66: 3-21.

Mookherjee, D. and S. Reichelstein, 2001, "Incentives and Coordination in Hierarchies," *Advances in Theoretical Economics*, vol1-iss1-art4.

Myerson, R., 1982, "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics,* 10: 67-81.

Palmer, J., 2003, "Relative Convexity," mimeo.

Quiggin, J., 1982, "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3: 225-243.

Rochet, J.C. and P. Choné, 1998, "Ironing, Sweeping, and Multidimensional Screening," *Econometrica,* 66: 783-826.

Rogerson, W., 1987, "On the Optimality of Linear Contracts," Mimeo Northwestern University.

Rogerson, W., 2003, "Simple Menus of Contracts in Cost-Based Procurement and Regulation," *American Economic Review*, 93: 919-926.

Shi, X., 2012, "Optimal Auctions with Information Acquisition," *Games and Economic Behavior,* 74: 666-686.

Sobel, J., 1993, "Information Control in the Principal-Agent Problem," *International Economic Review*, 34: 259-269.

Spinnewijn, J., 2008, "Insurance and Perceptions: How to Screen Optimists and Pessimists," Working Paper MIT.

Szalay, D. , 2005, "The Economics of Clear Advice and Extreme Options," *Review of Economic Studies,* 72: 1173-1198.

Szalay, D. , 2009, "Contracts with Endogenous Information," *Games and Economic Behavior,* 65: 586-625.

Tersiege, S., 2012, "Endogenous Information and Stochastic Contracts," *Games and Economic Behavior,* 76: 535-567.

Villeneuve, B., 2005, "Competition between Insurers with Superior Information," *European Economic Review,* 49: 321-340.

Yaari, M., 1987, "The Dual Theory of Risk Under Uncertainty," *Econometrica*, 55: 95-115.

APPENDIX

PROOF OF PROPOSITION 1. We suppose here that the agent's effort at the information gathering stage cannot be verified. Whether information has been gathered or not and the realized productivity shock (in case information has been successfully gathered) can be verified. Observing that optimal efforts are given by,

(A.1) $e(\theta) = \lambda(1+\theta)\alpha(\theta)$ and $e_u = \lambda\alpha_u$, the principal's problem can be written as

$$\max_{(\alpha(\theta),\alpha_u,U(\theta),U_u,a)} aE_\theta\left(\lambda(1+\theta)^2\left(\alpha(\theta)-\frac{1}{2}\alpha^2(\theta)\right)-U(\theta)\right)+(1-a)\left(\lambda\left(\alpha_u-\frac{1}{2}\alpha_u^2\right)-U_u\right)$$

subject to (4.3) and (4.4).

Inserting (4.4) into (4.3), we may rewrite the agent's participation constraint as:

$$U_u \geq -a\varphi'(a)+\varphi(a). \tag{A.2}$$

Observing that (A.2) is necessarily binding, we may rewrite the optimization problem as:

$$(\mathcal{P}^c): \max_{(\alpha(\theta),\alpha_u,a)} a\lambda E_\theta\left((1+\theta)^2\left(\alpha(\theta)-\frac{1}{2}\alpha^2(\theta)\right)\right)+(1-a)\lambda\left(\alpha_u-\frac{1}{2}\alpha_u^2\right)$$

$$-\varphi(a)-\underbrace{a\varphi'(a)\left(\frac{1}{\rho}-1\right)}_{\text{agency cost}=\mathcal{R}}.$$

This objective is exactly the same as if the effort at the information gathering stage were observable except for the extra agency cost of information gathering which is equal to the risk premium $\mathcal{R}$ (from (4.4) and (4.6)).

Optimizing immediately yields the expressions of the optimal bonuses given by (4.7). The corresponding efforts at the operating stage are then given by (4.8). The optimality condition with respect to a yields (4.9)

Taking into account (4.7) and inserting it into (4.1) and (4.2), we obtain the following expressions of the rents depending on whether the agent gets informed or not:

$$\hat{U}(\theta)=\frac{\lambda}{2}(1+\theta)^2-\beta(\theta) \text{ and } \hat{U}_u=\frac{\lambda}{2}-\beta_u. \tag{A.3}$$

With those expressions, the agent's incentive constraint at the information gathering stage (4.4) can be rewritten as:

$$\frac{1}{\rho}\varphi'(\hat{a})=E_\theta\left(\frac{\lambda}{2}(1+\theta)^2-\beta(\theta)\right)-\left(\frac{\lambda}{2}-\beta_u\right). \tag{A.4}$$

Since this incentive constraint at the information gathering stage depends only on the expected fee $E_\theta(\beta(\theta))$, there is actually no loss of generality in setting a constant fee:

$$\beta(\theta)\equiv\hat{\beta} \quad \forall\theta.$$

From (A.4) and (4.9), it then follows that:

$$\hat{\beta}_u-\hat{\beta}=\frac{1}{\rho}\varphi'(\hat{a})-\frac{\lambda\sigma^2}{2}=-\left(\frac{1}{\rho}-1\right)\hat{a}\varphi''(\hat{a})<0$$

where the negative sign comes from the fact that (4.9) also implies $\hat{a}>0$. *Q.E.D.*

Proof of Lemma 1. When the agent knows the realization of the shock θ and picks up the linear contract $t(y,\theta)=\alpha(\theta)y-\beta(\theta)$, he chooses an effort level such that:

$$e(\theta)=\arg\max_e \alpha(\hat{\theta})(1+\theta)e-\beta(\hat{\theta})-\psi(e). \tag{A.5}$$

The first-order condition for this strictly concave problem yields (5.2).

Taking into account the expression of the effort level, we can rewrite

$$U(\theta) = \max_{\hat{\theta}} \frac{\lambda}{2}\alpha^2(\hat{\theta})(1+\theta)^2 - \beta(\hat{\theta}). \tag{A.6}$$

From this, we immediately obtain that $U(\theta)$ is the maximum of convex functions and as such it is absolutely continuous, convex and so admits a sub-differential. It is even almost everywhere differentiable and, at any such point of differentiability, we have:

$$\dot{U}(\theta) = \alpha(\theta)e(\theta) = \lambda(1+\theta)\alpha^2(\theta).$$

Integrating yields (5.3).

Using simple revealed arguments and (A.6), we get for $\theta \geq \hat{\theta}$

$$\frac{\lambda}{2}\alpha^2(\theta)(1+\theta)^2 - \beta(\theta) \geq \frac{\lambda}{2}\alpha^2(\hat{\theta})(1+\theta)^2 - \beta(\hat{\theta}).$$

Writing the same incentive constraint but now between $\hat{\theta}$ and θ, summing on both sides and simplifying yields immediately that $\alpha(\theta) \geq \alpha(\hat{\theta})$ for $\theta \geq \hat{\theta}$. $\alpha(\cdot)$ is monotonically increasing and thus a.e. differentiable with at any point of differentiability:

$$\dot{\alpha}(\theta) \geq 0.$$

Q.E.D.

Proof of Lemma 2. From (5.4), we get $U(0) \geq V(0)$. But taking expectations in the left-hand side of (5.6), we get $U_u = V(0) \geq U(0)$.

From (5.7) and (5.4), we get for $\theta \geq 0$

$$\frac{1}{\theta}(U(\theta) - U(0)) = \frac{\lambda}{\theta}\int_0^\theta (1+x)\alpha^2(x)dx \geq \frac{\lambda\alpha_u^2}{2}(2+\theta)$$

Taking limits on both sides when $\theta \mapsto 0^+$ yields $\alpha^2(0^+) \geq \alpha_u^2$. Proceeding similarly for $\theta \mapsto 0^-$, we get $\alpha_u^2 \geq \alpha^2(0^-)$. Finally, we obtain (5.8).

Lemma 1 and condition (5.8) altogether imply that $\alpha(\theta) \geq \alpha(0^+) \geq \alpha_u$ for $\theta \geq 0$. From this, it follows that

$$\lambda\int_0^\theta (1+x)\alpha^2(x)dx \geq \frac{\lambda\alpha_u^2}{2}\theta(2+\theta) \quad \forall\theta \geq 0$$

and thus $U(\theta) \geq V(\theta)$ for all $\theta \geq 0$. The proof is similar for $\theta \leq 0$. Condition (5.4) holds thus everywhere. *Q.E.D.*

Proof of of Lemma 3. Using (5.3), we write:

$$E_\theta(U(\theta)) = U(0) + \int_0^\delta f(\theta)\left(\int_0^\theta (1+x)\alpha^2(x)dx\right)d\theta + \int_{-\delta}^0 f(\theta)\left(\int_0^\theta (1+x)\alpha(x)dx\right)d\theta.$$

Integrating by parts each of those integrals on the right-hand side leads to

$$E_\theta(U(\theta)) - U(0) = E_\theta\left(\lambda\frac{1_{\theta>0} - F(\theta)}{f(\theta)}(1+\theta)\alpha^2(\theta)\right).$$

Taking into account Lemma 2 yields then the result. *Q.E.D.*

Proof of of Proposition 2. With linear contracts, the principal's problem can be rewritten as

$$(\tilde{\mathcal{P}}): \quad \max_{(\alpha,a)} \lambda\left(\alpha - \frac{1}{2}\alpha^2\right)(1 + a\sigma^2) - \varphi(a) - a\varphi'(a)\left(\frac{1}{\rho} - 1\right) \text{ subject to (6.2).}$$

From (6.2), we may define a as a function of α, say $\Gamma(\alpha)$, with the derivative

$$\Gamma'(\alpha) = \frac{\rho\lambda\sigma^2\alpha}{\varphi''(\Gamma(\alpha))} \geq 0.$$

Inserting this into the above objective, the maximand in $(\tilde{\mathcal{P}})$ then becomes a function of α only, say $V_l(\alpha)$. We first compute the derivative of $V_l(\alpha)$ as:

$$\dot{V}_l(\alpha) = \lambda(1-\alpha)(1+\sigma^2\Gamma(\alpha))$$
$$+\left(\lambda\sigma^2\left(\alpha - \frac{1}{2}\alpha^2\right) - \frac{\varphi'(\Gamma(\alpha))}{\rho} - \Gamma(\alpha)\varphi''(\Gamma(\alpha))\left(\frac{1}{\rho} - 1\right)\right)\Gamma'(\alpha).$$

Manipulating $\dot{V}_l(\alpha) = 0$ we obtain (6.5). Observe that $a_0 = \Gamma(1)$. Hence, we get:

$$\dot{V}_l(1) = -a_0\varphi''(a_0)\left(\frac{1}{\rho} - 1\right)\Gamma'(1) < 0.$$

From this, it follows that:

$$1 > \tilde{\alpha}$$

and thus

$$a_0 > \tilde{a}.$$

Similarly, we get $0 = \Gamma(0)$ and

$$\dot{V}_l(0) = \lambda > 0.$$

Hence, we have:

$$0 < \tilde{\alpha}$$

and thus

$$0 < \tilde{a}.$$

Q.E.D.

Proof of Corollary 1. A first-order Taylor expansion of the first equality in (4.11) already yields the following approximation for $\hat{a}$:

$$\hat{a} = a^{FB}(1 - 2\epsilon). \tag{A.7}$$

Turning to (6.5) and denoting $\tilde{\alpha} = 1 - \tilde{\eta}$, another first-order Taylor expansion gives us:

$$\tilde{\eta} = \frac{\frac{\lambda\sigma^4}{2\varphi}}{1 + \frac{\lambda\sigma^4}{2\varphi}}(\epsilon - 2\tilde{\eta}) \geq 0. \tag{A.8}$$

A last Taylor expansion using (6.2) finally yields:

$$\tilde{a} = a^{FB}(1 - \epsilon - 2\tilde{\eta}). \tag{A.9}$$

From (A.7), (A.8) and (A.9), (6.6) immediately follows. *Q.E.D.*

PROOF OF THEOREM 1. We first observe that the fact that $\alpha(\theta)$ is non-decreasing and satisfy (5.8) altogether with Lemma 2 imply that the incentive constraint (5.4) binds only on a single connected interval $[\theta_*, \theta^*]$ (with $\theta^* \geq \theta_*$) including $\theta = 0$ (but possibly reduced to a single point at $\theta = 0$).

Suppose indeed it is binding on two such disconnected intervals $[\theta_1, \theta_2]$ (with possibly $\theta_1 = \theta_2$) and $[\theta_*, \theta^*]$ with $-\delta \leq \theta_1 \leq \theta_2 < \theta_*$. On (θ_2, θ_*) we have $U(\theta) > V(\theta)$ which implies that necessarily there exists a subinterval with non-empty interior $(\theta_2 + \epsilon, \theta_2 + 2\epsilon)$ (with ϵ small enough) such that $\dot{U}(\theta) > \dot{V}(\theta)$, i.e., $\alpha(\theta) > \alpha_u$, on that subinterval. This would contradict the monotonicity of $\alpha(\cdot)$ and in particular the fact that $\alpha(0^-) \leq \alpha_u$.

Consider now the interval $[\theta_*, \theta^*]$ and suppose it has a non-empty interior. Since $U(\theta) = V(\theta)$ on that interval, differentiating with respect to θ yields $\dot{U}(\theta) = \dot{V}(\theta)$ and thus $\alpha(\theta) = \alpha_u$. Of course, on $[\underline{\theta}, \theta_*]$, the monotonicity of $\alpha(\cdot)$ implies that we must have $\alpha(\theta) \leq \alpha_u$. This property is used below to prove that indeed, $\theta_* = \underline{\theta}$.

With those earlier findings in mind, we can rewrite the information gathering incentive constraint (5.7) as:

$$\lambda\left(\int_{\theta^*}^{\delta}(1-F(\theta))(1+\theta)\alpha^2(\theta)d\theta - \int_{-\delta}^{\theta_*}F(\theta)(1+\theta)\alpha^2(\theta)d\theta + \int_{\theta_*}^{\theta^*}(1_{\theta>0} - F(\theta))(1+\theta)\alpha_u^2 d\theta\right)$$

$$= \frac{1}{\rho}\varphi'(a) \tag{A.10}$$

First, we omit constraint (6.1) for a while. Second, we write the Lagrangean of the so-relaxed problem $(\mathcal{P}^*)$ as:

$$L(\alpha(\cdot), \alpha_u, a, \mu) = a\lambda\left(\int_{\theta^*}^{\delta}(1+\theta)^2\left(\alpha(\theta) - \frac{1}{2}\alpha^2(\theta)\right)f(\theta)d\theta + \int_{-\delta}^{\theta_*}(1+\theta)^2\left(\alpha(\theta) - \frac{1}{2}\alpha^2(\theta)\right)f(\theta)d\theta\right)$$

$$+\lambda\left(1 - a + a\int_{\theta_*}^{\theta^*}(1+\theta)^2 f(\theta)d\theta\right)\left(\alpha_u - \frac{1}{2}\alpha_u^2\right) - a\varphi'(a)\left(\frac{1}{\rho} - 1\right) - \varphi(a)$$

$$+\mu\lambda\left(\int_{\theta^*}^{\delta}(1-F(\theta))(1+\theta)\alpha^2(\theta)d\theta - \int_{-\delta}^{\theta_*}F(\theta)(1+\theta)\alpha^2(\theta)d\theta + \alpha_u^2\int_{\theta_*}^{\theta^*}(1_{\theta>0} - F(\theta))(1+\theta)d\theta\right)$$

$$-\frac{\mu}{\rho}\varphi'(a)$$

where μ is the multiplier of (A.10), shown below to be negative. From now on, we assume that this Lagrangean is quasi-concave and derive the optimality conditions by means of first-order conditions.

- We optimize w.r.t. α_u and pointwise w.r.t. $\alpha(\theta)$.

1. On $[\theta_*, \theta^*]$, we get:

$$\left(1 - a + a\int_{\theta_*}^{\theta^*}(1+\theta)^2 f(\theta)d\theta\right)(1-\alpha_u) + 2\alpha_u\mu\int_{\theta_*}^{\theta^*}(1_{\theta>0} - F(\theta))(1+\theta)d\theta = 0. \tag{A.11}$$

2. On $[\theta^*, \delta]$, we get:

$$a(1+\theta)^2(1-\alpha(\theta))f(\theta) + 2\mu(1-F(\theta))(1+\theta)\alpha(\theta) = 0$$

$$\text{(A.12)} \quad \Leftrightarrow \alpha(\theta) = \frac{1}{1 - \frac{2\mu R(\theta)}{a}}.$$

Note that, under Assumption 1, $\alpha(\theta)$ is monotonically increasing on $[\theta^*, \delta]$ when $\mu < 0$ and is worth 1 at $\theta = \delta$ if $\theta^* < \delta$.

3. On $[-\delta, \theta_*]$, and if $\alpha(\theta)$ is not constrained by (6.1), we get:

$$a(1+\theta)^2 (1 - \alpha(\theta)) f(\theta) - 2\mu F(\theta)(1+\theta)\alpha(\theta) = 0$$

$$\text{(A.13)} \quad \Leftrightarrow \alpha(\theta) = \frac{1}{1 + \frac{2\mu S(\theta)}{a}}.$$

Observe that, under Assumption 1, $\alpha(\theta)$ so defined is monotonically increasing on $[-\delta, \theta_*]$ when $\mu < 0$ and $\alpha(-\delta) = 1$ if $\theta_* > -\delta$, so that $\alpha(\theta) \geq 1$ for all $\theta \in [-\delta, \theta_*]$. In particular, this unconstrained solution always violates (6.1) since the monotonicity condition on $\alpha(\cdot)$ implies that $\alpha_u \leq \alpha(\theta^*) \leq \alpha(\bar{\theta}) = 1$. Necessarily, the constraint (6.1) thus binds on the whole interval $[-\delta, \theta_*] \bigcup [\theta_*, \theta^*]$. We conclude that

$$\text{(A.14)} \quad \alpha(\theta) = \alpha_u \quad \forall \theta \in [-\delta, \theta^*].$$

We may then also rewrite (A.10) as:

$$\text{(A.15)} \quad \lambda \left(\int_{\theta^*}^{\delta} (1 - F(\theta))(1+\theta)\alpha^2(\theta) d\theta + \int_{\underline{\theta}}^{\theta^*} (1_{\theta > 0} - F(\theta))(1+\theta)\alpha_u^2 d\theta \right) = \frac{1}{\rho}\varphi'(a).$$

• Inserting the condition (A.14) into (A.11) and taking into account that $\theta_* = -\delta$ yields:

$$\text{(A.16)} \quad \alpha_u = \frac{1}{1 - \frac{2\mu \int_{-\delta}^{\theta^*} (1_{\theta > 0} - F(\theta))(1+\theta) d\theta}{1 - a + a\int_{-\delta}^{\theta^*} (1+\theta)^2 f(\theta) d\theta}}.$$

Observe that $\alpha_u < 1$ when $\mu < 0$.

• Optimizing the Lagrangean w.r.t. θ^* yields:

(A.17)

$$a(1+\theta^*)^2 f(\theta^*) \left(\alpha(\theta^{*+}) - \frac{1}{2}\alpha^2(\theta^{*+}) - \alpha_u + \frac{1}{2}\alpha_u^2 \right) + \mu(1 - F(\theta^*))(1+\theta^*) \left(\alpha_u^2 - \alpha^2(\theta^{*+}) \right) = 0.$$

An obvious solution to (A.17) consists in having $\alpha(\theta)$ continuous at θ^*:

$$\text{(A.18)} \quad \alpha(\theta^{*+}) = \frac{1}{1 - \frac{2\mu R(\theta^*)}{a}} = \alpha_u$$

• Finally, optimizing the Lagrangean w.r.t. a yields:

$$\varphi'(a) + ((1-\rho)a + \mu)\varphi''(a) = \rho\lambda \left(\int_{\theta^*}^{\delta} (1+\theta)^2 \left(\alpha(\theta) - \frac{1}{2}\alpha^2(\theta) \right) f(\theta) d\theta \right)$$

$$-\rho\lambda\left(\alpha_u-\frac{1}{2}\alpha_u^2\right)\left(1-\int_{-\delta}^{\theta^*}(1+\theta)^2f(\theta)d\theta\right). \tag{A.19}$$

• **Proof that $\theta^*>0$.** Taking together (A.12), (A.14) and (A.18) and assuming that $\mu\neq 0$ (an assertion to be checked below), we obtain that θ^* must solve (6.7). From Assumption 1, the left-hand side is a decreasing function of θ^* which is equal to $\frac{1}{2f(0)}$ for $\theta^*=0$ and zero for $\theta^*=\delta$. The right-hand side, say $\chi(\theta)$ is negative at $\theta=0$ and positive for $\theta=\delta$ (since after integrating by parts, one finds that the numerator is then worth $\int_{-\delta}^{\delta}(1_{\theta>0}-F(\theta))(1+\theta)d\theta=\frac{1}{2}E_\theta((1+\theta)^2-1)=\frac{\sigma^2}{2}>0$). It is zero for θ_0 such that $\frac{\sigma^2}{2}=\int_{\theta_0}^{\delta}(1-F(\theta)(1+\theta)d\theta$. Moreover, it can be checked that $\dot\chi(\theta^*)=0$ and $\dot\chi(\theta)<0$ in the neighborhood of $\theta=\delta$. Hence, there exists a unique $\theta^*>0$ that solves (6.7) and $\chi(\theta)$ has a maximum there. Let denote this solution as $\theta^*=\Omega(a)$. It can be readily seen that $\Omega(a)$ decreases with a.

• **Proof that $\alpha_u<1$ and $\mu<0$.** Eliminating μ/a from (A.13) and (A.16), we get:

$$\alpha(\theta)=\Lambda(\alpha_u,a,\theta)\equiv\frac{\alpha_u}{\alpha_u+(1-\alpha_u)\frac{R(\theta)}{R(\Omega(a))}}\quad\forall\theta\geq\Omega(a). \tag{A.20}$$

This, together with the fact that $\theta^*=\Omega(a)>0$, allows us to rewrite (A.15) as:

$$\lambda\left(\int_{\Omega(a)}^{\delta}(1-F(\theta))(1+\theta)\left[\alpha^2\right]_{\alpha_u}^{\Lambda(\alpha_u,a,\theta)}d\theta+\frac{\sigma^2}{2}\alpha_u^2\right)=\frac{1}{\rho}\varphi'(a).\tag{A.21}$$ [22]

The principal's problem can be rewritten as a simple optimization over (α_u,a) as:

$$(\mathcal{P}):\quad\max_{(\alpha_u,a)}\lambda\left(\alpha_u-\frac{1}{2}\alpha_u^2\right)(1+a\sigma^2)+\lambda a\int_{\Omega(a)}^{\delta}(1+\theta)^2f(\theta)\left[\alpha-\frac{1}{2}\alpha^2\right]_{\alpha_u}^{\Lambda(\alpha_u,a,\theta)}d\theta$$

$$-\varphi(a)-a\varphi'(a)\left(\frac{1}{\rho}-1\right)\text{ subject to (A.21).}$$

From (A.21), we may define a as a function of α_u, say $\Xi(\alpha_u)$. Using this expression, the maximand of $(\mathcal{P})$ then becomes a function of α_u only, say $V(\alpha_u)$. From now on, we assume that $V(\alpha_u)$ is quasi-concave so that that the solution to $(\mathcal{P})$ is actually characterized by means of a first-order condition. We compute the derivative of $V(\alpha_u)$ as:

$$V'(\alpha_u)=\lambda(1-\alpha_u)(1+\sigma^2\Xi(\alpha_u))$$

$$+\left(\lambda\sigma^2\left(\alpha_u-\frac{1}{2}\alpha_u^2\right)-\frac{\varphi'(\Xi(\alpha_u))}{\rho}-\Xi(\alpha_u)\varphi''(\Xi(\alpha_u))\left(\frac{1}{\rho}-1\right)\right)\dot\Xi(\alpha_u)$$

$$+\lambda\left(\int_{\Omega(\Xi(\alpha_u))}^{\delta}(1+\theta)^2f(\theta)\left[\alpha-\frac{1}{2}\alpha^2\right]_{\alpha_u}^{\Lambda(\alpha_u,\Xi(\alpha_u),\theta)}d\theta\right)\dot\Xi(\alpha_u)$$

$$+\lambda\Xi(\alpha_u)\int_{\Omega(\Xi(\alpha_u))}^{\delta}(1+\theta)^2f(\theta)$$

$$\times\left[(1-\Lambda(\alpha_u,\Xi(\alpha_u),\theta))\left(\frac{\partial\Lambda}{\partial\alpha_u}(\alpha_u,\Xi(\alpha_u),\theta)+\frac{\partial\Lambda}{\partial a}(\alpha_u,\Xi(\alpha_u),\theta)\dot\Xi(\alpha_u)\right)-(1-\alpha_u)\right]d\theta. \tag{A.22}$$

[22] We use the compact notation $[f]_y^x=f(x)-f(y)$.

Observe that $\Lambda(1, a, \theta) = 1$. Inserting into (A.21) yields $\Xi(1) = a_0$. That $\Lambda(1, a, \theta) = 1$ also implies that the last three terms in (A.22) are zero. Finally, we get:

$$V'(1) = \left(\frac{\lambda\sigma^2}{2} - \frac{\varphi'(a_0)}{\rho} - a_0\varphi''(a_0)\left(\frac{1}{\rho} - 1\right)\right)\dot{\Xi}(1) = -a_0\varphi''(a_0)\left(\frac{1}{\rho} - 1\right)\dot{\Xi}(1) < 0.$$

From this, it follows that:

$$\alpha_u < 1.$$

Inserting into (A.16), it follows that $\mu < 0$ as supposed. *Q.E.D.*