Optimal contract under moral hazard with soft information

Guillaume Roger *

The University of New South Wales

October 31, 2012

Abstract

I study a model of moral hazard with soft information: the agent *alone* observes the stochastic outcome of her action; hence the principal faces a problem of *ex post* adverse selection. With limited instruments the principal cannot solve these two problems independently; the *ex post* incentive for misreporting interacts with the *ex ante* incentives for effort. This affects the shape and properties of the optimal contract, which fails to elicit truthful revelation in all states. In this set up audit and transfer become strategic *complements*; this is rooted in the non-separability of the problem.

Keywords: moral hazard, asymmetric information, soft information, contract, mechanism, audit. JEL Classification: D82.

*School of Economics, UNSW. Email: g.roger@unsw.edu.au. I owe special thanks to Andy Skrzypacz. Thank you also to an anonymous referee, Guido Friebel, Claudio Mezzetti, Mike Peters, Ralph Winter, my colleagues Richard Holden, Bill Schworm, Chris Bidner, Carlos Pimienta and Suraj Prasad for helpful comments. Financial support from the Australian School of Business at UNSW is gratefully acknowledged.

1 Introduction

The standard solution of a moral hazard problem requires the observation of some informative signal of the agent's action. It is then possible to design a second-best contract, which is conditioned on that information instead of the actual action. When performance is difficult to observe or noisy, the signal may be complemented. Sometimes however performance is not observed at all: an accounting report, for instance, is *not* a direct observation of the state of a firm.

This paper studies exactly this problem: the outcome realization is *not* observable by the principal, but must be reported by the agent. Then the principal is exposed to *ex ante* moral hazard and adverse selection *ex post*. The object of the paper is to characterize the optimal contract when these two problems *interact*. Bar for the issue of observability, the model mirrors that of a standard moral hazard problem. A risk-neutral principal delegates production to a risk-averse agent, who relies on a stochastic technology. The agent *alone* observes the outcome θ, which must therefore be elicited *ex post*. Because the principal otherwise observes nothing, the contract must include an audit and some (bounded) penalty.[1] The model attempts to be faithful to audit as a sampling process, which is imperfect.[2] Applications are broad-ranging. For example, after hiring the CEO, a board often asks of him (her) to report his (her) results while on the job; a regulated firm may be asked to reveal its production cost after investing in an uncertain technology. Kedia and Philippon (2006) also document the pervasiveness of "earnings management"–a euphemism for fraudulent accounting–which arises in the equilibrium of this model.

The paper's main contribution is to show that audit and transfer are strategic complements. More precisely, they respond to exogenous changes by moving in the same direction. This departs from standard results of the costly state verification literature, which shows they are substitutes. It also depart from the standard literature on moral hazard, where

[1] Doornik (2006) notes that penalties are always bounded for courts will not enforce any penalty in excess of the damages incurred.

[2] For example, financial audits are sampling processes. See the discussion for details.

monitoring and incentives are substitutes.[3] The more powerful the *ex ante* incentives for effort (i.e. the steeper the transfer function with respect to the outcome), the more attractive is the option to manipulate information *ex post*, especially when it is bad. Therefore the more accurate must the audit be. Yet the transfer function must remain steep enough to generate *ex ante* incentives in the first place. This is the fundamental conflict of this paper.

This work is closely connected to that of Mookherjee and Png (1989, now MP), who show that with enough instruments, the twin problem of moral hazard and *ex post* adverse selection can be treated separately. More precisely, separability allows for *ex post* truthful revelation without any consequence on the incentive device used to solve the *ex ante* moral hazard problem. Their connection becomes moot and the moral hazard problem can be solved in standard fashion, yielding standard results. The analysis of such interaction has received scant attention in economics, possibly because the Revelation Principle (applied by MP) is too powerful in some sense. Indeed the accounting literature roots misreporting of information in some failure of the Revelation Principle (e.g. Arya, Glover and Sunder, 1998 or Demski and Frimor, 1999). I suggest a different route.

The starting premise is that the real world does not accord with the results of MP; agents do mislead their principal. For example, Bally Total Fitness, a large chain of fitness clubs, fired its controller and treasurer, then its CFO, for misleading accounting in 2005. More recently, Howie Hubler, a "headstrong" trader at Morgan Stanley single-handedly lost the firm $9 Bn after covering up his trades, was terminated and yet paid out past boni.[4] Thus a model that systematically predicts truthful revelation has limited applicability. Second, the schemes suggested by MP are not observed in practice. Executive compensation contracts, for example, may specify a diverse array of contingent payments, but usually not a bonus for not misleading shareholders. A third objection is that the transfer (a reward) that is

[3] "Auditing" is understood to mean observing the output; to observe the agent;s action is to "monitor"; see for example Khalil and Lawarree (1995), who show this is not a trivial distinction. There is no monitoring in this model, only auditing.

[4] Bally's executive were not subject to legal proceedings of any kind, neither was Hubler. Sources: Motley Fool at motley.com. and "The Big Short", Michael Lewis, 2008.

necessary to induce information revelation may be arbitrarily large; it turns the principal into a source of money regardless of the value of the productive relationship with the agent.

In this paper complete truth-telling can never be an equilibrium.[5] Furthermore, when truthful revelation is possible for at least some states, the agent misreports in the worse states, where the incentive is strongest and the cost is lowest. The reason is non-separability: a single transfer is used to solve the *ex ante* and the *ex post* problems. This is not enough to disentangle them and introduces a fundamental tension between *ex ante* effort provision, which requires a state-contingent compensation, and *ex post* information revelation, which is best addressed with a constant transfer. The interaction of these two problems implies that the optimal transfer function is "option-like". It must satisfy an implicit limited liability constraint (because of the bounded penalty), which creates systematic incentives for message inflation. This option shape accords well with many real-life contracts.

The papers closest to this one are MP and Kanodia (1985). Both consider a combination of moral hazard and *ex post* adverse selection with no observability. MP combine a Grossman-Hart (1983) model with an *ex post* revelation mechanism. The principal may use a transfer for each of the moral hazard and the adverse selection problems; the latter is a reward (by limited liability) that may be arbitrarily large. In the present paper, the principal can use only one payment, which also accords well with real-life contracts.[6] Border and Sobel (1987) construct an audit mechanism with endogenous penalties. The optimal probability of audit varies in the messages sent; truthful revelation obtains with arbitrarily large penalties (and ignoring the agent's participation decision, as pointed out by MP).[7] In all these papers, auditing is perfect but the principal controls the probability of running an audit. Here the audit is imperfect and closer to sampling, which is what real financial audits do, and has been modeled by Bushman and Kanodia (1996) or Demski and Dye (1999). Crocker and Morgan

[5]Whether any truthful revelation occurs is determined in equilibrium; it depends on the whole contract.

[6]MP's model yields a quirky byproduct: the agent strictly prefers being audited. This owes to the construction of the revelation constraint, which implicitly only allows *reward* to be offered for truth-telling.

[7]In Khalil (1997) truthful revelation obtains through a standard direct revelation mechanism. Auditing relaxes the agent's incentive constraint; the principal trades-off the audit cost with the information rent.

(1998) construct an optimal insurance contract in the presence of fraud: actual damages may be inflated *ex post*. In equilibrium there is always falsification, as here. It is necessary to induce separation, which is a condition of efficient insurance; here separation is a condition of effort. The payment scheme internalizes this fraud and is low-powered, as in this paper. Doornik (2006) considers the opposite problem: the principal only observe the outcome of the agent's effort and may renegotiate at the interim stage. The interim offer is informative; the equilibrium is Perfect Bayesian. A rejection triggers costly legal enforcement, which does occur with some probability in equilibrium. This is sufficient to be unable to implement the first best in spite of the agent risk neutrality. The contract is a one-step bonus: a (extreme) form of option. An "option-like" contract is derived by Jewitt, Kadan and Swinkels (2008, now JKS) when the agent must receive a minimum payment: it pays a constant wage below a threshold, and an increasing transfer beyond. In this paper bounded penalties imply a limited liability constraint; the contract takes a similar shape.

Close in spirit, Gromb and Martimort (2007) let (an) expert(s) search for some information by exerting some effort, who then has (have) to disclose it to the principal. To overcome the moral hazard problem, the expert's incentive contract must be made state-dependent. Like in this paper, this very fact introduces adverse selection. However, a contract can be conditioned on the final outcome, unlike here. For the purpose of this discussion, Krähmer and Strausz (2011) adopt a similar construct in the context of pre-project planning. Malcomson (2009) studies a problem where the agent acquires soft information that may be used by the agent to make a decision yielding a verifiable outcome. The principal may have incentives to distort the decision rule away from the first-best to foster information acquisition. Levitt and Snyder (1997) develop a contracting model in which the agent receives an early (soft) signal about the likely success of the project. With appropriate early information, the principal can decide whether to shut-down or continue. To obtain this information, he must commit to shut-down less frequently than the unconstrained solution prescribes. The eventual outcome is fully observed by the principal, hence contractible. In all these papers, information is still *exogenously* given although *ex ante* unknown to the agent. Here the private information emerges endogenously.

After introducing the model, Section 3 deals with the *ex post* information revelation problem. Next I characterize the optimal contract; Sections 5 explores some properties. Section 6 presents an extensive discussion. The proofs and some of the technical material are relegated to the Appendix.

2 Model

A principal delegates a task to an agent who undertakes an action $a \in \mathcal{A} \subset \mathbb{R}_+$. The action's cost $c(a)$ is increasing and convex, and yields a stochastic outcome $\theta \in [\underline{\theta}, \overline{\theta}] \equiv \Theta \subset \mathbb{R}$ with conditional distribution $F(\theta|a)$ and corresponding density $f(\theta|a) > 0$. The density $f(\theta|a)$ satisfies the MLRP: f_a/f is non-decreasing, concave in θ; therefore $F(\theta|a')$ stochastically dominates $F(\theta|a)$ in a first-order sense when $a' > a$. I make the additional assumption that $F_a(F^{-1}(\theta|a))$ is convex in (θ, a).[8] The agent *alone* observes the outcome θ and reports a message $\omega \in \Omega$ to the principal, whereupon she receives a transfer t. Her net utility is given by $u(t,a) = v(t) - c(a)$, where $v : \mathbb{R} \mapsto \mathbb{R}$ is a continuous, increasing, concave function with $v(0) = 0$. The principal receives a net payoff $S(t, \theta) = \theta - t$. If the true state θ were observable by the principal, the model would collapse to the textbook moral hazard problem. The principal can commit to the contract.

At the stage of information revelation, effort is sunk so all that matters is the utility $v(t)$, which can only be conditioned on the message ω. Given the monotonicity of $v(t)$, either all types pool to the same message if $t(\omega)$ is increasing, or have no effort incentive at all if it is constant. Auditing restores a measure of *ex post* observability; it has zero marginal cost. However it is imperfect and uncovers misreporting with probability $p(\omega - \theta; \alpha)$, where $p : \mathbb{R} \mapsto [0, 1]$ is a continuous, differentiable function in both arguments and $p(0; \alpha) = p(\cdot; 0) = 0$.[9] This breaks the monotonicity of $v(t)$. The technology $p(\cdot; \alpha)$ is costly to acquire; it is drawn from a family $\mathcal{P}$ of increasing, at least weakly convex functions parametrized by an investment

[8]This is sufficient for the Concave Local Informativeness condition of JKS.

[9]This is akin to a sampling process, as in Bushman and Kanodia (1996).

α at cost $k(\alpha)$, increasing and convex. The parameter α affects the slope of $p(\cdot;\alpha)$ at 0, that is, the precision of the audit. Auditing remains imperfect: $\forall\alpha,\ \partial p(z|\alpha)/\partial z|_{z=0} < \infty$ but there are no type-II errors. If discovered the agent receives nothing.[10] With this construction the expected utility function of an agent at the revelation stage is $U = v(t(\omega))\,[1 - p(\omega - \theta;\alpha)]$. Hence,

$$\frac{\partial U}{\partial t} = v'\,[1-p] \geq 0;\ \frac{\partial^2 U}{\partial t \partial \theta} = v'p' \geq 0 \tag{2.1}$$

is a sorting condition on the *ex post* expected utility of the agent, akin to the Spence-Mirrlees condition. The timing is almost standard:

1. The principal offers a contract $\mathcal{C} = \langle \Omega, t(\omega), p(\omega - \theta; \alpha)\rangle$ made of a message space, a transfer function and an audit technology;

2. The agent accepts or rejects the contract. If accepting, she also chooses an action a;

3. Action a generates an outcome $\theta \in \Theta$ observed by the agent only;

4. The agent reports a message $\omega \in \Omega$;

5. Audit occurs (because it has 0 marginal cost);

6. Transfers are implemented and payoffs are realized.[11]

3 Information Revelation

This Section focuses on information transmission. It takes advantage of some results contained in a companion paper (Roger 2012) that are briefly explained. Then it is shown that truthful revelation in any arbitrary state θ amounts to a condition relating the transfer function $t(\cdot)$ to the probability $p(\cdot|\alpha)$. This defines three regimes: complete, partial or no information revelation. The dependence on α is suppressed where convenient.

[10] See Section 6 for a discussion of these two assumptions.

[11] That payoffs are realized needs not imply that they are observed by the principal, as in the accounting example. Mathematically, not observing θ does not prevent maximizing $\mathbb{E}[S(t,\theta)]$ or any other monotone transformation $\mathbb{E}[S(t,g(\theta))]$. See also Grossman and Hart (1983), Remark 4.

3.1 Preliminaries

Consider a mechanism with message space Ω and suppose that the transfer function $t(\omega)$ is increasing and a.e. differentiable.[12] The agent sends a message ω such that $\max_{\omega\in\Omega} v(t(\omega))\,[1-p(\omega-\theta)]$, i.e. $\omega(\theta)$ solves:

$$v't'(\omega)[1-p(\omega-\theta)] - v(t(\omega))p'(\omega-\theta) = 0 \tag{3.1}$$

Condition (3.1) must bind for some ω because $p(\cdot)$ is monotonically increasing. For a mechanism to be truthful, $v(t(\theta)) \geq v(t(\omega))\,[1-p]$, or $v(t(\theta)) = \max_{\omega\in\Omega} v(t(\omega))\,[1-p(\omega-\theta)]$. Using (3.1), this is equivalent to there being some $\tilde{\theta}$ such that:

$$v't'(\tilde{\theta}) = v(t(\omega))|_{\omega=\tilde{\theta}} \cdot p'(0). \tag{3.2}$$

Roger (2012) establishes that *(i)* a direct mechanism where $\Omega = \Theta$ induces a measure of pooling and *(ii)* there is no loss of generality in restricting attention to a *separating* mechanism that uses an enlarged, but simple, message space labeled $\widehat{\mathcal{M}}$.[13] That is, choosing the appropriate message space becomes part of the design problem, unlike instances where attention may be restricted to direct mechanisms without loss. That sufficient message space is defined as follows: consider some set $\mathcal{M}$ such that $\Theta \subset \mathcal{M} \subset \mathbb{R}$. Let

$$\widehat{m}(\theta, t) = \arg\max_{m\in\mathcal{M}} v(t(m))\,[1-p(m-\theta)]\,,$$

then we have $\widehat{\mathcal{M}} \equiv \{\widehat{m}(\theta;t) \in \mathcal{M} | \widehat{m} \in \arg\max\ U\ \forall\theta\in\Theta\}$. The mapping $\widehat{m}$ is a function of the transfer t, which is fixed and committed to at the stage of information revelation. Already we see that information revelation and *ex ante* incentives interact. Lemma 8 (in the Appendix) shows there is no better message space than $\widehat{\mathcal{M}}$.[14]

Partial pooling may arise in a direct mechanism because *all* agents may have incentives to misreport upwards but the top type cannot report more than $\overline{\theta}$; this applies to a positive

[12]This is not a restriction: $p(\cdot;\alpha)$ is continuous, so must be $t(\cdot)$. See Roger (2012).

[13]Kartik (2009) derives similar results in a model of almost cheap talk, for essentially the same reasons. See the Discussion.

[14]This observation stems from working through an anonymous referee's comment, whom I must thank.

measure of agents. This partial pooling dampens *ex ante* incentives because the contract does not sufficiently discriminate between outcomes: the compensation scheme is flat for a range of outcomes. A completely separating mechanism always dominates. Furthermore, removing some messages from the set $\widehat{\mathcal{M}}$ also does not help. It always renders the contract more expensive to the principal. If $\widehat{\mathcal{M}}$ is truncated from below, the agent must overstate her optimal message for some realisation (which is risky and therefore costly). If some interior messages are prohibited, the *ex post* expected utility U is no longer monotonic in θ. But any non-monotonic scheme is dominated by a monotone one (Carlier and Dana, 2005).

3.2 Degrees of information revelation

Equation (3.2) embodies a requirement on the precision of the audit at 0; that is, it defines a subset $\mathcal{P}_0(t) \subseteq \mathcal{P}$ of audit functions that can elicit truthful revelation for at least some values of θ, given the transfer t. Condition (3.2) is necessary and sufficient for truthful revelation at $\tilde{\theta}$, which does not mean it holds for *all* values. There may be three cases of interest; which of these the principal faces is determined in equilibrium.

Case 1: Partial truthful revelation. This corresponds to condition $v't'(\tilde{\theta}) = v(t(\tilde{\theta}))p'(0)$ for some value $\tilde{\theta} \in \left(\underline{\theta}, \overline{\theta}\right)$. If $v(t(\cdot))$ is concave, then $v't'|_{\theta \geq \tilde{\theta}} \leq v(t(\tilde{\theta}))p'(0)$ and truth-telling obtains above $\tilde{\theta}$. Similarly, $v't'|_{\theta < \tilde{\theta}} > v(t(\tilde{\theta}))p'(0)$ and truth-telling is out of reach below $\tilde{\theta}$ (so $\widehat{m}(\theta) > \theta$). The converse is true for $v(t(\cdot))$ convex. Figure 1 (left panel) depicts an interior example of $\tilde{\theta}$ when $v(t(\cdot))$ is a concave function.

The next two cases are special instances of the first one.

Case 2: Truthful revelation. Condition (3.2) is satisfied for all values of the private information θ; more precisely, $\forall \theta,\ v't'(\theta) \leq v(t(\theta))p'(0)$. Jointly with the transfer, the audit technology $p(\cdot;\alpha)$ is sufficiently precise so $\forall \theta,\ \widehat{m}(\theta) = \theta$.

Case 3: No truthful revelation. Condition (3.2) fails to hold anywhere on the range Θ, i.e. $\forall \theta \in \Theta,\ v't'(\theta) > v(t(\theta))p'(0)$. This is shown on the right panel of Figure 1. This

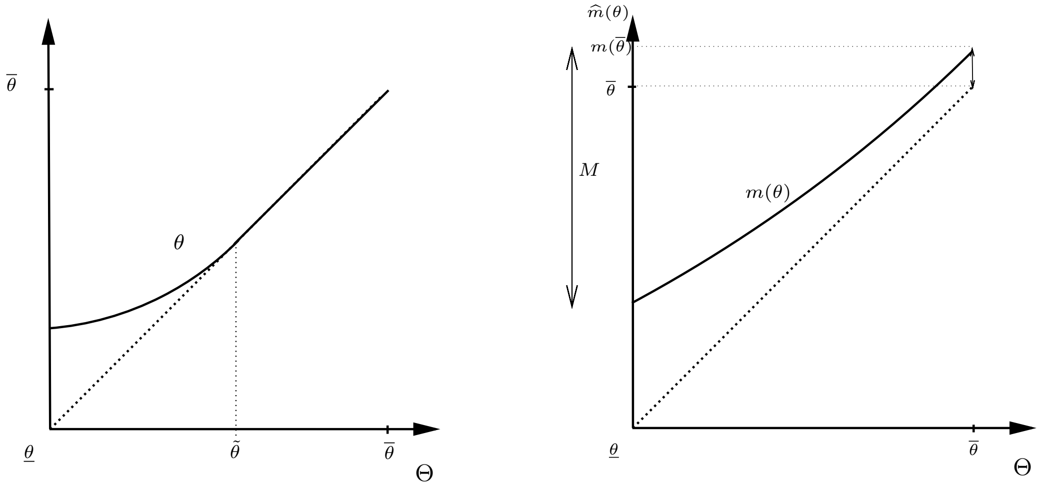

Figure1: Optimal messages above and below $\tilde{\theta}$ (left); with extended message space (right)

problem is the reason for constructing a separating mechanism.

This rich array of outcomes obtains because of non-separability of the twin problem of *ex ante* moral hazard and *ex post* adverse selection. That non-separability stems from the combination of the imperfect audit technology and the limited number of instruments. It implies a fundamental tension between *ex ante* effort incentives, which require a state-contingent transfer, and *ex post* information revelation that is best addressed with state-independent transfers. The consequence is that in Cases 2 and 3, an agent who is induced to exert any effort necessarily misreports her private information with positive probability. Indeed, re-arrange the truth-telling condition as $p(\omega - \theta) \geq 1 - v(t(\theta))/v(t(\omega))$: for a given α, this inequality is more difficult to satisfy when $t(\cdot)$ is steep.

One last remark is in order. There may exist many contracts satisfying $t' \geq 0$: some may include jumps, there may be intervals on which $t' = 0$ and so on, with implications for the message $\widehat{m}$. It is not obvious that $\widehat{m}$ must be continuous, as it is depicted in Figure 1. To see why, consider a scheme $t(\cdot)$ that is flat on some range, say, on $\Theta_f \equiv [\theta_1, \theta_2]$. If $\tilde{\theta} \geq \theta_2$ the agent misreports her information on Θ_f as anywhere else below $\tilde{\theta}$. If $\tilde{\theta} \leq \theta_1$, she may face the conditions $v't'(\theta_1) \leq v(t(\theta_1))p'(0;\alpha)$ but $v't'(\theta_2) \geq v(t(\theta_2))p'(0;\alpha)$, i.e. $t(\cdot)$ may be steeper at θ_2 than at θ_1 and (3.2) is reversed. Then one moves from truthful revelation above $\tilde{\theta}$ and below θ_1 to misreporting from θ_2 on, i.e. there is a jump in the optimal message (because

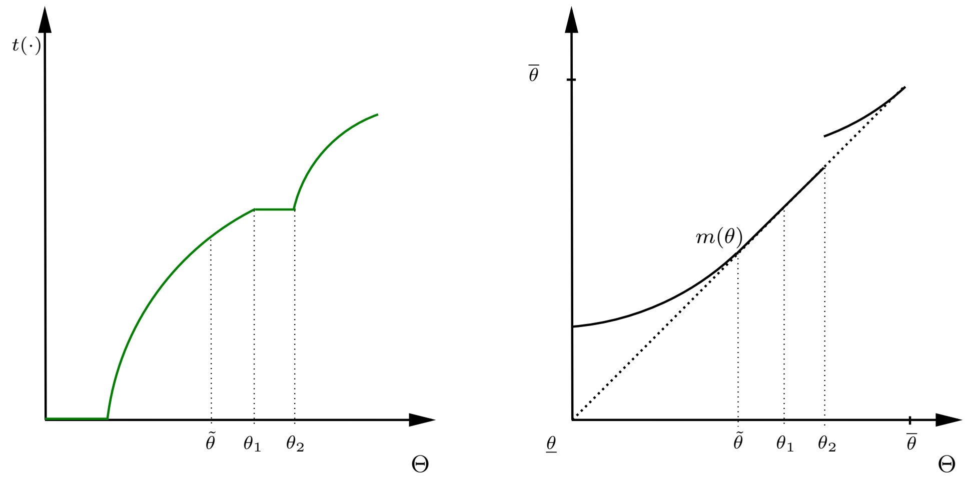

Figure 2: contract (left) may induce jump(s) in the optimal message (right)

$v(t(\tilde{\theta})) \geq (1-p)v(t(m(\tilde{\theta})))$ at $\tilde{\theta}$ but $v(t(\theta_2)) < (1-p)v(t(m(\theta_2)))$– see Figure 2.

4 Characterization

To proceed, I first seek to understand the behavior of the contract for some fixed audit technology $p(\cdot;\alpha)$. Then I endogenize α, to which all other endogenous variables also respond, and optimize fully over the whole set of instruments t, a, α. I use the first-order approach.[15]

The ensuing analysis may be problematic in that the agent's utility

$$U = \begin{cases} v(t(\theta)), & \theta \geq \tilde{\theta}; \\ (1-p(m-\theta))v(t(m)), & \theta < \tilde{\theta}. \end{cases}$$

may not be smooth, nor even continuous, at $\tilde{\theta}$. It turns out that it must be both; the formal statement takes the form of Lemma 6 in the Appendix. From this it follows that the optimal message is also a smooth function of θ at $\tilde{\theta}$ by the Theorem of the Maximum (see Figure 1), so the regime change at $\tilde{\theta}$ is "smooth".[16] Defining $t : \widehat{\mathcal{M}} \mapsto \mathbb{R}$, the principal's program is

[15]See Jewitt (1988) or Conlon (2009) for validations; Jewitt (1988) specifically for sufficient conditions.

[16]The other potential source of discomfort is that highlighted in Figure 2, i.e. a jump away from truth-telling above $\tilde{\theta}$; this is addressed later.

Problem 1

$$\max_{\alpha,t,a} \int_{\underline{\theta}}^{\tilde{\theta}} [x - (1 - p(\widehat{m}(x) - x;\alpha))t(m(x))]\, dF(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} [x - t(x)]\, dF(x|a) - k(\alpha)$$

s.t.

$$\widehat{m}(\theta) = \arg\max_{m\in\mathcal{M}} (1 - p(m - \theta))v(t(m)) \tag{4.1}$$

$$\int_{\underline{\theta}}^{\tilde{\theta}} v(t(\widehat{m}(x)))[1 - p(\widehat{m}(x) - x)]dF(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t(x))dF(x|a) - c(a) \geq 0 \tag{4.2}$$

$$\int_{\underline{\theta}}^{\tilde{\theta}} v(t(\widehat{m}(x)))[1 - p(\widehat{m}(x) - x)]dF_a(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t(x))dF_a(x|a) = c'(a) \tag{4.3}$$

where $\tilde{\theta} \equiv \tilde{\theta}(p(\cdot;\alpha), t, a)$. The *ex post* message may be entirely truthful (only drawn from Θ), not at all (and only drawn from $\widehat{\mathcal{M}}$) or some of both depending on where $\tilde{\theta}$ lies.[17] From an *ex ante* standpoint the principal must account for any of these possibilities, which the objective function and the constraints reflect. Condition (4.1) is the agent's information revelation constraint – the novelty in this paper. Let λ be the Lagrange multiplier of constraint (4.3), μ that of (4.2) and t^O denote the solution of the following conditions.

Lemma 1 *Fix a and α. The first-order conditions of Problem 1 are given by:-*

$$\frac{1}{v'(t(\widehat{m}(\theta)))} = \mu + \lambda\frac{f_a}{f}; \tag{4.4}$$

for $\theta < \tilde{\theta}$; and

$$\frac{1}{v'(t(\theta))} = \mu + \lambda\frac{f_a}{f}; \tag{4.5}$$

for $\theta \geq \tilde{\theta}$, where $\widehat{m}(\theta)$ is determined by (4.1) and $\mu, \lambda \geq 0$.

The case of complete information revelation is obtained by extending $\tilde{\theta}$ to $\underline{\theta}$. Then the first-order condition is standard; (4.5) holds over Θ. Case 3 corresponds to $\tilde{\theta} \geq \overline{\theta}$.

[17] Note that although the problem does not specify a distribution over the message space $\mathcal{M}$, $F(\theta|a)$ is still the relevant distribution because $\widehat{m}(\theta)$ is injective. For details, see Roger (2012). More comprehensively the program allows for jumps as described in Section 3; the principal's objective is then $\int_{\underline{\theta}}^{\tilde{\theta}} [x - (1 - p(m(x) - x;\alpha))t(m(x))]\, dF(x|a) + \int_{\tilde{\theta}}^{\theta_2} [x - t(x)]\, dF(x|a) + \int_{\theta_2}^{\hat{\theta}} [x - (1 - p(m(x) - x;\alpha))t(m(x))]\, dF(x|a) + \int_{\hat{\theta}}^{\overline{\theta}} [x - t(x)]\, dF(x|a) - k(\alpha)$, with a jump at θ_2 and two thresholds $\tilde{\theta}, \hat{\theta}$–and the agent's utility is similarly modified. The analysis extends immediately.

4.1 Incentives under soft information

The agent's incentives, and the principal's response, are altered in two ways that are extensively explored in Roger (2012). This section presents a brief summary. First, the potential for misrepresentation affects the agent's incentives in Problem 1. Denote by $\hat{a}$ the action solving $\int_\Theta v(t(x))dF_a(x|a) = c'(a)$–the standard moral hazard constraint. Formally,

Lemma 2 *Fix the transfer function t across models. Whenever $\tilde{\theta} \geq \underline{\theta}$, the agent selects an action a^* solving (4.3) below the action $\hat{a}$.*

For a given transfer schedule the agent's moral hazard constraint is hardened because her expected payoff is higher in any state thanks to message inflation–so the marginal benefit of effort is lower. Very low outcomes are tempered by the option to exaggerate them; they do not provide strong incentives.

Second, the *ex post* penalties act like a default payment that interacts with the *ex ante* incentives. Indeed, the agent can always do better than accepting a negative transfer: she can simply take the lottery $\{p, 1-p\}$ over 0 and some positive $v(t(m))$ by exaggerating her message. In the words of JKS, this penalty becomes "payment binding"; that is, it becomes a limited liability constraint. The transfer function is modified in consequence, as in JKS. Because the ratio f_a/f is monotonic and $\mathbb{E}_\Theta[f_a/f] = 0$, for some action a there exists some θ_a such that $f_a(\theta_a|a)/f(\theta_a|a) = 0$.

Lemma 3 *Fix a. The optimal transfer t^O takes the form*

$$\frac{1}{v'(t^{SB})} = \begin{cases} \kappa, & \forall\ \widehat{m}(\theta) \leq \theta_a; \\ \kappa + \lambda\frac{f_a}{f}, & \forall\ \widehat{m}(\theta) > \theta_a. \end{cases}$$

where $\kappa \geq 0$, $\kappa \neq \mu$ and $\widehat{m}(\theta)$ solves (4.1). Furthermore, the multiplier λ of the moral hazard constraint (4.3) is strictly positive.

JKS call this kind of scheme option contracts. The constant κ corresponds to the minimum payment the agent must receive *ex post*. It is evident that the option contract generates weaker incentives for the agent because failure does not carry great consequences. For the

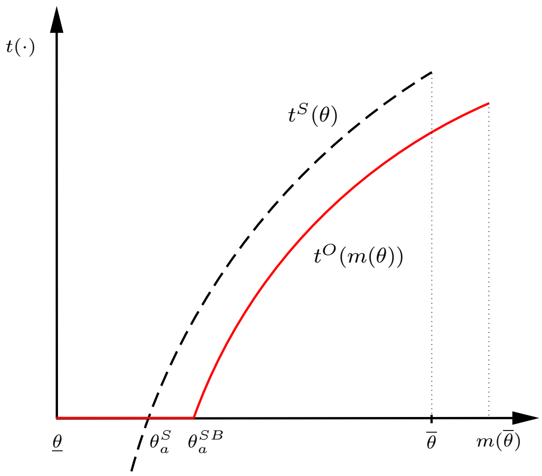

Figure 3: Transfers functions–standard and distorted.

principal that means that a high action becomes more expensive to implement. To complete the description of the transfer function,

Proposition 1 *The optimal transfer function t^O solving Problem 1 is continuous, non-decreasing over Θ and described by Lemma 3; in particular, it:-*

- *is continuous but with a kink at θ_a;*
- *is non-decreasing concave for all θ above θ_a; and*
- *pays zero below θ_a.*

Figure 3 depicts the transfer function. The zero payment below θ_a stems from the zero *ex post* penalty. Notice that when $\tilde{\theta}$ is interior, t^O is still continuous at $\tilde{\theta}$. The reason is that $\widehat{m}(\theta)$ smoothly converges to θ at $\tilde{\theta}$ because the function U is smooth. (See the left panel of Figure 1). An immediate consequence of Proposition 1 is

Corollary 1 *Completely truthful revelation (Case 2) can never occur in equilibrium.*

Because the optimal contract pays zero on the range $[\underline{\theta}, \theta_a]$, the agent is strictly better off taking the lottery $\{p, 1-p\}$ over zero (if detected) and a positive payoff obtained by reporting $\widehat{m}(\theta) > \theta$. Furthermore, because the optimal transfer function is concave, misreporting always occurs "at the bottom" (see Case 2). Indeed, the agent with the strongest incentives to misreport is the one in the worse states. It is also the agent whose cost of misreporting is the lowest. From the collection of the previous results, it is also true that:

Corollary 2 *The optimal message $\widehat{m}(\theta)$ is everywhere continuous on Θ; i.e. there are no jumps.*

This follows from the fact that the optimal transfer function t^O is monotone strictly concave from θ_a on. There can be no pair $\theta_1 < \theta_2$ such that $v't'(\theta_2) > v't'(\theta_1)$; thus Condition (3.2) cannot be simultaneously holding at θ_1 but reversed at θ_2. Furthermore, there can be only at most one threshold $\tilde{\theta}$, and the three simple regimes described in Section 3 are exhaustive.

4.2 Optimal contract

As part of the optimal contract the principal selects his audit technology $p(\cdot;\alpha) \in \mathcal{P}$ by choice of α. This may have two effects. First, fixing $t(\cdot)$ and a, it may alter the degree of information revelation, i.e. the cutoff $\tilde{\theta}$ (Cases 1 to 3). Second, $t(\cdot)$ and a are endogenous variables, so they too adjust to a change in α. The optimal contract balances all these effects.

Proposition 2 *The optimal contract is characterised by:-*

1. *a continuous transfer scheme* $t^O = \begin{cases} t^O(\widehat{m}(\theta)), & \theta < \tilde{\theta}; \\ t^O(\theta), & \theta \geq \tilde{\theta}. \end{cases}$ *determined by Proposition 1, and Conditions (4.4) and (4.5) on the relevant ranges;*

2. *an action a^O solving the first-order condition*

$$\begin{aligned} &\int_{\underline{\theta}}^{\tilde{\theta}} [x - t(\widehat{m}(x))(1-p)]dF_a + \int_{\tilde{\theta}}^{\overline{\theta}} [x - t(x)]dF_a \\ &+\lambda \left[\int_{\underline{\theta}}^{\tilde{\theta}} v(t(\widehat{m}(x)))(1-p)dF_{aa} + \int_{\tilde{\theta}}^{\overline{\theta}} v(t(x))dF_{aa} - c''(a) \right] = 0 \end{aligned} \quad (4.6)$$

3. *and an audit investment $\alpha^O = \alpha_1^O + \alpha_2^O$, where α_1^O solves*

$$v't'(\underline{\theta}) = v(t(\underline{\theta}))p'(0;\alpha_1^O) \quad (4.7)$$

and $\alpha^O \geq \alpha_1^O$ solves

$$\int_{\underline{\theta}}^{\tilde{\theta}} t(\widehat{m}) p_\alpha dF(x|a) + \lambda \int_{\underline{\theta}}^{\tilde{\theta}} v(t(\widehat{m})) p_\alpha dF_a(x|a) = k'(\alpha) \tag{4.8}$$

The cut-off $\tilde{\theta} \in [\underline{\theta}, \overline{\theta}]$ is determined by (3.2) given t^O, a^O, α^O.

The threshold $\tilde{\theta}$ is free to lie at either boundary or to be interior; it is endogenous to the contract and so is the regime one operates under. The first two items of Proposition 2 resemble standard ones. The last one determines the level of investment in the audit technology. It allows for α_2^O to be zero, that is, $\tilde{\theta} = \underline{\theta}$. If so, the technology is sufficiently inexpensive (or equivalently, precise) for Condition (3.2) to hold at $\underline{\theta}$. Condition (4.7) thus pins down the smallest investment necessary for truthful revelation. In that case, the transfer is determined by (4.5) and (4.6) collapses to the standard expression; the *pair* t^O, α^O, together with the zero penalty, are such that they compel truthful revelation. If α_1^O is not sufficient, the investment may be increased from α_1^O to α^O (i.e. by α_2^O), and this entails a trade-off given by (4.8); that is, further distortions arise. The total marginal benefit (LHS) includes saving on undue transfers, as well as *relaxing* the moral hazard constraint. When truthful revelation is impossible, the transfer is determined solely by (4.4) and (4.6) is modified by extending $\tilde{\theta}$ to $\overline{\theta}$. Truth-telling cannot be guaranteed (unlike in Mookherjee and Png [26]), because t^O, α^O are *jointly* determined. Whether truthful revelation obtains does not just depends on the audit procedure because the problems of moral hazard (*ex ante*) and adverse selection (*ex post*) are meshed.

5 The relationship between audit and transfers

Because both the transfer t and the audit investment α are costly to the principal, a question of practical importance is to understand how they relate. Indeed, the costly state verification literature (as Mookherjee and Png, 1989 or Border and Sobel, 1987 among others) establishes and exploits the fact that transfer and audit are substitutes. Similarly we know that monitoring and transfers are substitutes in standard moral hazard problems. On the other hand, we know that the expected cost of effort defined as $T(a) \equiv \int_\Theta t(x)dF(x|a)$ is increasing, concave in a (see Conlon, 2008). So too in this model:

$$\begin{aligned} T^O(a) &\equiv \int_\Theta t^O(m(x))[1 - p(m(x) - x)]dF(x|a) \\ &= \int_{\underline{\theta}}^{\theta_a} t^O(m(x))[1 - p(m(x) - x)]dF(x|a) + \int_{\theta_a}^{\overline{\theta}} t^O(m(x))[1 - p(m(x) - x)]dF(x|a) \\ &= 0 + \int_{\theta_a}^{\overline{\theta}} t^O(m(x))[1 - p(m(x) - x)]dF(x|a) \end{aligned}$$

follows the same properties (for $p(\cdot)$ is weakly convex and the product of two concave functions t^O and $[1 - p]$ is concave). So at face value it is not clear how t^O and α^O relate. Indeed, in this model, α is used to elicit information *ex post*: the better the audit, the less over-reporting and the lower the transfer in a given state. But we know this relaxes the moral hazard constraint, which induces a higher action. This allows for a higher transfer in any state (as the principal re-optimizes). Because all these variables are jointly determined, to make a statement about their behaviour I run a comparative statics exercise on the primitives of the problem.

Proposition 3 *The transfer t^O and the audit investment α^O both:-*

1. *decrease in the dispersion of the distribution F (in the sense of SOSD);*
2. *decrease in the agent's risk-aversion;*
3. *decrease as the cost of effort $(c(a))$ increases;*
4. *increase in the principal's payoffs.*

Thus high-power contracts are necessarily accompanied with a large enough investment in the audit technology. Conversely, it is because the audit is sufficiently precise that the contract can be high-powered. Increasing t in isolation in response the the moral hazard problem is destructive; it requires a simultaneous increase in audit.

This claim may be counterintuitive; it stems from the non-separability of the problem and is explained as follows. In a costly state verification problem the transfer's only purpose is to provide incentives for information revelation (either as a penalty or a reward).

Because transfer and audit enter the incentive constraint multiplicatively they naturally are substitutes. In fact, the Maximum Punishment Principle of Baron and Besanko (1984) tells us that only one transfer is necessary (the most extreme one). Here the transfer's primary purpose is to induce effort, so it must be upward sloping. This is the very source of adverse selection: it is this responsiveness of the transfer to the state that generates the incentives to manipulate information. So the stronger are these incentives, the more beneficial is the audit.[18]

From a practical standpoint, Proposition 3 together with Condition (3.2) suggest it may not be the lack of audit that is the culprit in corporate embezzlement and earnings manipulation. There is little doubt that firms of that nature are subject to audit. Rather the audit may not have been sufficient given the incentives offered.

6 Discussion

Limited liability

In this model bounded penalties act like a limited liability constraint, which drives the shape of the optimal contract. Introducing a "proper" limited liability constraint on transfers would not change the substance of the paper. Consider such a constraint $t \geq \underline{t}$ and some penalty $-l \leq 0$. The relevant constraint for the agent facing some bad state θ is $\max\{\underline{t}, -l \cdot p\}$: only one constraint really matters. In this paper I effectively let the constraint on penalties be the relevant one. This may not be to an entirely trivial effect in that even a wealthy agent may be shielded by bounded penalties, but it does fit the examples of the introduction.

Other penalties

The paper purposefully bounds penalties; left unconstrained they necessarily lead to truthful revelation (unless they conflict with a limited liability constraint, which is essentially

[18] A friend of mine sits on a few board of very large, publicly listed corporations. To her this statement is equivalent to an economist hearing that agents respond to incentives.

equivalent to the present model). Here I discuss two potential modifications in this respect.

Harsher penalties. The model could allow for penalties $-l < 0$. Then the information revelation condition (3.1) would become $v't'(1-p) - p'(m-\theta)\left[v(t(m)) - v(-l)\right] = 0$ and clearly *(i)* there would be less exaggeration and *(ii)* for some l large enough, $\widehat{m}(\theta) = \theta \ \forall\theta$ (no misreporting). That is, one would revert to model closer to that of MP.[19] If l were not too large, the problem would remain as here, albeit muted. The only significant difference is that the threshold θ_a would be such that f_a/f would be negative.

Penalties conditioned on offense. The Maximal Punishment Principle (see Baron and Besanko, 1984, now MPP) asserts that the penalty should be as severe as possible, and thus swiftly rules out conditioning it on the offense (e.g. small deviations from the state θ could be met with fines that commensurate). Setting the MPP aside, suppose that the principal instead uses some fine $\varphi \equiv \varphi(\widehat{m} - \theta)$ where $\varphi(0) = 0$. The agent expected utility then becomes $U = (1-p)v(t(\widehat{m})) + pv(t(\widehat{m}) - \varphi(\widehat{m} - \theta))$ and one can see that the truth-telling condition (3.2) turns into

$$v't'(\theta) = p'(0)\left[v(t(\widehat{m})) - v(t(\widehat{m}) - \varphi(\widehat{m} - \theta))\right]|_{\widehat{m}=\theta} + p(\widehat{m} - \theta; \cdot)v'\varphi'|_{\widehat{m}=\theta}$$

i.e. $v't'(\theta) = 0$. In other words, driving a wedge between the transfers when the agent reports truthfully and does not, is essential. That is, $\varphi(\cdot)$ must be discontinuous at 0. How large a wedge (discontinuity) is discussed above at some length. The MPP applies in this model as in many others because the audit generates no false negatives.

Audit technology

In this paper the accuracy of the audit is conditioned on the the magnitude of the misreporting. The literature has considered other approaches such as conditioning the precision of the audit on the message alone. Absent additional punishments or rewards (as in Border and Sobel, 1987 or MP) this *cannot* deliver separation, let alone truthful revelation. To see why,

[19] Noting that here truthful revelation would obtain immediately from the exogenous penalty.

rewrite the audit technology as $p(\omega;\alpha)$; this may also be interpreted either as a probability of running the audit, given some message ω, as in those papers. Then truth-telling requires $v(t(\theta)) = \max_{\omega\in\Theta} v(t(\omega))\,[1-p(\omega)]$, i.e. $v't'(\theta) = 0$; hence the need for fines or rewards in both these papers, and many others.

Furthermore, according to most accounting standards (e.g. US GAAP or the AASB in Australia), an audit seeks to provide a *reasonable* assurance that statements are free from material errors. As a result, a sampling procedure is usually adopted by financial auditors, who can verify the details of the transaction(s). This justifies the absence of type-II errors in the process.[20] Note also that the accuracy of this verification process cannot be connected to the message received, but rather to its veracity. Statistical sampling is also followed by ISO-accredited companies for the purpose of quality assurance.[21] In either case, the audit is *always* performed. The technology $p(\cdot;\alpha)$ displays exactly these characteristics.

Participation fee and binding constraint

The agent receives an *ex ante* rent in this model; the participation constraint fails to bind. This could be addressed with an *ex ante* participation fee, say ϕ. Then a contract entails a tariff $(t(\cdot),\phi)$ and the agent's expected utility reads $U = (1-p)v(t(\widehat{m})-\phi) + pv(-\phi)$ where $v(-\phi) < 0$. The truth-telling condition (3.2) becomes $v't'(\theta) = p'(0)\,[v(t(\widehat{m})-\phi) - v(-\phi)]$. Because $v(t(\cdot)$ is concave, $v(t-\phi) - v(-\phi) > v(t)$ for each t, so for a fixed transfer function the truth-telling condition holds for a larger set of states θ. That is, $-\phi$ acts like $-l$ (see the first paragraph of this discussion). When ϕ is not too large, the information revelation problem remains as in the main text.

[20] "If controls are assessed as appropriate and operating as expected then lower levels of substantive testing is expected. [...] appropriate sampling (either statistically -in total or stratified - or judgementally when a small number of items make up much of the volume) is performed and transactions and account balances verified. The steps involved include tracing transactions from the general ledger back to supporting documents or from initiating documents through to the ledger to ensure that they are appropriately included." Mark Pickering, Auditor at Deloitte Touche Tohmatsu, 1986-91

[21] ISO: International Organization for Standardization.

The main purpose of the fee ϕ is to render the participation constraint binding; suppose such a fee does exist. When $\mu > 0$ however the optimal transfer function still retains the same shape. The reason is that ϕ is paid *ex ante*, so *ex post* the agent still faces a gamble $\{p, 1-p\}$ over utilities $\{v(-\phi), v(t(\widehat{m}) - \phi)\}$ versus taking some really bad $v(t(\theta) - \phi)$.

If the participation constraint is made to bind the agent no longer receives an *ex ante* rent but an *ex post* information rent $U(t^O, \theta) = [1-p(\widehat{m}(\theta)-\theta)]v(t^O(\widehat{m}(\theta)))-v(t^O(\theta)) > 0$, $\forall \theta < \tilde{\theta}$ that is decreasing in the state θ.

Other disclosure models

***M*-implementability (Green and Laffont, 1986).** These authors study the implementability of a social choice function when the agent may report a message from a set $M(\theta) \subset \Theta$, where $M(\cdot)$ is exogenous and publicly known. The idea is to allow for the agent to report small lies (the set $M(\theta)$) around the true state, and characterise the set of social choice functions that are truthfully implementable. It may sometimes be optimal for the principal to *not* induce truthful revelation. This is clearly a feature of the present paper, where the principal is better off with a contract that allows for reporting outside the type space, and where truthful revelation can never occur for at least some states.

Green and Laffont (1986) provide a necessary and sufficient condition – called the nested range condition (NRC) – for the agent to report her information truthfully. The NRC does not hold in this model, although it corresponds to a game of of "unidirectional distortions with an ordered space" (to use their words) – example a(2) in Green and Laffont (1986). Because the agent has a unique optimal deviation for each type, the set $M(\theta) = \{\omega | \omega \in \Theta, \omega \geq \theta\}$ (to use their notation) collapses to a singleton for each type, whence no nesting condition can possibly hold. This is because Green and Laffont (1986) exogenously allow for the agent's response to be a correspondence, whereas here the agent's optimal message is unique.

Almost cheap talk (Kartick et al., 2007; Kartik, 2009) Kartik (2009) builds on Crawford and Sobel's 1982 paper and introduces a lying cost k. He finds there cannot exist a completely separating equilibrium and that any equilibrium entails a measure of message inflation, which is decreasing in the lying cost. Pooling occurs because the high types "run out of messages to send", which is exactly the problem the principal faces if using a direct mechanism (see also Roger, 2012). In both papers, there exists a critical type, above which pooling occurs. This problem does not arise in Kartik et al. (2007) because the message space is unbounded; this is the approach I suggest in this paper too.

A large message inflation accompanies a small lying cost in Kartik (2009). This maps into a small probability of discovery in this paper, i.e. a poor audit technology. Message inflation is not problematic for Kartik (2009) because types are exogenous and the receiver anticipates inflation (and adjust his response). It is costly here because it hampers the *ex ante* incentives for effort.

7 Conclusion

When a principal cannot observe the outcome of his agent's action in a moral hazard framework and needs to elicit this information from that very agent, he faces a problem of *ex post* adverse selection as well. With limited instruments, this introduces a fundamental tension between *ex ante* incentive, for which a contingent transfer is necessary, and *ex post* incentives, best addressed with a state-independent transfer. Type separation (not necessarily truthful revelation) requires the use of an *ex post* audit and penalties.

The *ex post* adverse selection problem is costly to the principal in three ways: first, the agent is able to exaggerate her actual performance and thereby may receive an inflated transfer. The principal's response introduces a first set of distortions. Second, because penalties are weak, they act as an implicit limited liability constraint. As a result the participation constraint cannot bind (there are rents) and the contract resembles an option. Last, the very fact that the contract entails a region with constant transfer implies that complete truthful revelation can never arise in equilibrium. There may be partial truthful revelation below a threshold; that is, the agent misreports her information in the worse states because the incentive is the strongest and the cost the lowest.

A key result of this paper is that the audit investment and the level of transfer co-vary. That is, the stronger the incentives offered to the agent, the more she must be audited to be kept in check. In light of practical examples drawn from real life, this seems to be an important feature that was so far absent from our extensive literature on moral hazard and its applications.

8 Appendix

8.1 Preliminaries

I begin with a series of Lemmata that address the potential lack of smoothness of the agent's expected utility function U, and others that the will be useful throughout.

Lemma 4 *The function U is a.e. differentiable over Θ.*

Proof: By application of the Theorem of Lebesgue to a monotonically increasing function; i.e. by (3.2), U is monotonically increasing. ■

Then naturally:

Lemma 5 *Suppose a solution $m(t;\theta)$ of FOC (3.1) exists, then*

1. *this solution is unique;*
2. *$m(\theta)$ is a.e. differentiable and*
3. $\frac{dm}{d\theta} > 0$

Proof: Directly from the sorting condition $\frac{\partial^2 U}{\partial t \partial \theta} = v'p' > 0$, we know that condition (3.1) admits a unique maximiser when it binds. That $m(\theta;t)$ is increasing in θ is immediate from observing that the agent's optimisation problem is supermodular. I will need more that this statement though. Continuity of the solution $m(t;\theta)$ follows from the Theorem of the Maximum. To show that $m(\theta,t)$ is monotonically increasing, re-arrange (3.1) as $v't'/v = p'/1-p$, i.e. $d\ln(v(t(m)))/dm = -d\ln(1-p)/dm$. Take some $\theta' > \theta$ and suppose $m(\theta') \leq m(\theta)$. Then $p'(m(\theta')-\theta')/1-p\prime(m(\theta')-\theta') < p'(m(\theta)-\theta)/1-p\prime(m(\theta)-\theta)$, so that $d\ln(v(t(m(\theta'))))/dm < d\ln(v(t(m(\theta))))/dm$. Therefore $v(t(m(\theta'))) > v(t(m(\theta)))$ and since $v(\cdot)$ and $t(\cdot)$ are monotone increasing, $m(\theta') > m(\theta)$, a contradiction. The same can be shown if taking some $\theta' < \theta$ and supposing that $m(\theta') \geq m(\theta)$. It follows that $m(\theta,t)$ is a.e. differentiable, by application of the Theorem of Lebesgue, except at most for a finite set of points. Differentiate (3.1) with respect to θ and rearrange. ■

In spite of Lemma 4, there may still exist problematic discontinuities, especially at $\tilde{\theta}$, and this point is one of particular interest.

Lemma 6 *The function U is continuous and differentiable at $\tilde{\theta}$ when $\tilde{\theta} \in (\underline{\theta},\overline{\theta})$*

Proof: I show that U cannot be discontinuous at $\tilde{\theta}$ and that by Condition (3.2) it must be also differentiable. The proof is written for U concave but also applies with obvious adjustments when it is convex. Suppose $v(t(\cdot))$ is at least weakly concave; since only upward deviations are of concern, the trouble is that we may have $v(t(\tilde{\theta})) < [1-p(m(\tilde{\theta}-\varepsilon)-(\tilde{\theta}-\varepsilon))]v(t(m(\tilde{\theta}-\varepsilon)))$ for $\varepsilon > 0$, $\varepsilon \to 0$. Suppose so, then truth-telling cannot be an optimal response at $\tilde{\theta}$. So there must exist some value $\theta_0 < \tilde{\theta}$ (possibly $\underline{\theta}$) such that $v(t(\tilde{\theta})) \geq [1-p(m(\theta)-\theta)]v(t(m(\theta)))$ for $\theta \in [\theta_0,\tilde{\theta})$. Let $\theta \to \tilde{\theta}$, this is exactly the definition of continuity. Now notice that

$$v't'(\tilde{\theta}) = v(\tilde{\theta})p'(0;\alpha) \Leftrightarrow \frac{\partial}{\partial\theta}v(t(\theta))|_{\tilde{\theta}} = \frac{\partial}{\partial\theta}[1-p(m(\theta)-\theta)]v(t(m(\theta)))|_{\tilde{\theta}}$$

or $\frac{\partial}{\partial\theta}U|_R = \frac{\partial}{\partial\theta}U|_L$ at $\tilde{\theta}$. So U is differentiable. Condition (3.2) is a pasting condition at $\tilde{\theta}$. ■

Lemma 7 *The mapping $m : \Theta \mapsto \mathcal{M}$ is piece-wise weakly convex in θ.*

Proof: Take first $\tilde{\theta} \in (\underline{\theta},\overline{\theta})$. $m(\theta)$ is increasing and a.e. differentiable by application of Lemma 1, with $m(\underline{\theta}) > \underline{\theta}$ for any $\tilde{\theta} > \underline{\theta}$. Because U is continuous and differentiable,

$\lim_{\theta \uparrow \tilde{\theta}} m(\theta) = \theta$. Suppose now that $m(\theta) - \theta$ were increasing; then $dm(\theta)/d\theta > 1$ and $\lim_{\theta \uparrow \tilde{\theta}} m(\theta) \neq \theta$; so $m(\theta) - \theta$ must be decreasing, and consequently, $dm(\theta)/d\theta < 1$. Therefore $m(\theta)$ is convex when $\tilde{\theta} \in (\underline{\theta}, \overline{\theta})$. Now extend $\tilde{\theta}$ to $\overline{\theta}$ to obtain Case 3. ■

Lemma 8 *The message space $\widehat{\mathcal{M}}$ is the optimal message space.*

Proof: Proposition 1 of Roger (2012) shows that the principal is at least weakly better off extending the message space from Θ to $\widehat{\mathcal{M}}$ (strictly when truthful revelation is impossible, i.e. when $\tilde{\theta} \geq \overline{\theta}$). Lemma 2 of the same paper extends the Revelation Principle: there is no gain by using richer message spaces than $\widehat{\mathcal{M}}$. It is immediate that restricting $\widehat{\mathcal{M}}$ by truncating it from the bottom (say, $\min m > \widehat{m}(\underline{\theta})$) does not help. It induces a measure of types to overstate their report beyond what is privately optimal (given by Condition (3.1)), for which they have to be compensated (the agent is risk-averse). Last, considering a grid of messages of the form $M = \{m_0, m_1, ..., m_i, ...m_N\}$ also does not help.

First take as given that the transfer function must be monotone increasing, and consider an arbitrary set M and two arbitrary points $m_i, m_{i+1} \in M$.[22] (Using any such grid will be shown to be dominated by using an interval $\widehat{\mathcal{M}}$, so we need not worry whether that grid is optimal.) By monotonicity of t, $t(m_{i+1}) \geq t(m_i)$ and agent θ reports m_{i+1} over m_i if and only if

$$v(t(m_{i+1}))[1 - p(m_{i+1} - \theta)] \geq v(t(m_i))[1 - p(m_i - \theta)].$$

Because M is a grid, there exists some θ_i such that

$$v(t(m_{i+1}))[1 - p(m_{i+1} - \theta_i)] = v(t(m_i))[1 - p(m_i - \theta_i)],$$

and for types to the "left" of θ_i, $t = t(m_i)$, while to the right of θ_i, $t = t(m_{i+1})$. Similarly for pairs of messages m_{i-1}, m_i and m_{i+1}, m_{i+2} and so on. That is, offering the agent a grid amounts to offering a transfer scheme that is a step function of the type. With this, the agent's *ex post* expected utility

$$U = v(t(m_i))[1 - p(m_i - \theta)]$$

[22] This is established in a companion paper, Roger (2012).

is no longer monotonic in θ; it reaches a local maximum at $\theta = m_i$ and local minima at θ_{i-1} and θ_i. In contrast

$$U = v(t(\widehat{m}(\theta)))[1 - p(\widehat{m}(\theta) - \theta)]$$

is monotone (increasing) by application of the Envelop Theorem when $\widehat{m}(\theta)$ is continuous.

Suppose there exists a scheme $(t(m_i), m_i),\ m_i \in M$ that is optimal and induces effort $\bar{a}$. Let $w(\theta) = v(t(m_i(\theta)))[1 - p(m_i(\theta) - \theta)]$, where $m_i(\theta) \in \arg\max_{m_i \in M} U$, so that a solution to this problem can be represented as $(w(\theta), \bar{a})$. The *ex ante* expected utility reads

$$\mathbb{E}[U(w)] = \int_{\Theta} w(x) dF(x|a)$$

By Proposition 1 of Carlier and Dana (2005), there exists a *non-decreasing* function $z(\theta)$ (the non-decreasing re-arrangement of w) such that, for any fixed a,

$$\mathbb{E}[U(w)] = \int_{\Theta} w(x) dF(x|a) = \int_{\Theta} z(x) dF(x|a) = \mathbb{E}[U(z)]$$

By Lemma 2 of Carlier and Dana (2005), $(z(\theta), \bar{a})$ also represents a solution; here is can be constructed as $z(\theta) = v(t(\widehat{m}(\theta)))[1 - p(\widehat{m}(\theta) - \theta)]$. Furthermore, for the principal,

$$\int_{\Theta} S(z(\theta), \theta) dF(x|\bar{a}) \geq \int_{\Theta} S(w(\theta), \theta) dF(x|\bar{a})$$

with a strict inequality if $w(\theta)$ is not monotonic–which is the case here.

It is immediate that the claim extends to any other modification of $\widehat{\mathcal{M}}$, such as restricting it to be disjoint intervals or combinations of intervals and points. ■

8.2 Proofs

Proof of Lemma 1: By pointwise optimization of Problem 1. Below $\tilde{\theta}$, $m(\theta) > \theta$, so the transfer $t^{SB} \equiv t(m(\theta))$, while above $\tilde{\theta}$, $t^S \equiv t(\theta)$. Notice that $\theta_a \leq \tilde{\theta}$, otherwise there exists an interval $[\tilde{\theta}, \theta_a]$ where t^O is constant and the agent reports truthfully. But this cannot be optimal by (3.1). ■

Proof of Lemma 2: Fix the transfer schedule t; by optimality of the message $\widehat{m}(\theta)$, $(1-p)v(t(\widehat{m}(\theta))) > v(t(\theta))$, $\forall \theta < \tilde{\theta}$ and $\widehat{m}(\theta)$, $(1-p)v(t(\widehat{m}(\theta))) \geq v(t(\theta))$, $\forall \theta \geq$

$\tilde{\theta}$. So for any given action a, $\int_\Theta (1-p)v(t(\widehat{m}(\theta)))dF \geq \int_\Theta v(t(\theta))dF$ and therefore $\int_\Theta (1-p)v(t(\widehat{m}(\theta)))dF_a \leq \int_\Theta v(t(\theta))dF_a$ by concavity in a. These latter two inequalities become strict as soon as $\tilde{\theta} > \underline{\theta}$. ■

Proof of Lemma 3: The existence, sufficiency and uniqueness of such contract is shown in Jewitt, Kadan and Swinkels [17] (in particular, they show the multipliers μ, λ exist and are non-negative). To construct the contract, fix some action a^O and take the first-order condition. We know $\mu = 0$ necessarily, so below θ_a the transfer must be such that $1/v'$ remains non-negative. To show that the multiplier λ is positive, fix some a. Integrate $1/v'$ over Θ:

$$\mathbb{E}_\theta \left[\frac{1}{v'(t^O)} \right] = \kappa \int_{\underline{\theta}}^{\overline{\theta}} dF(x|a) + \lambda \int_{\theta_a}^{\overline{\theta}} \frac{f_a}{f} dF(x|a) = \kappa + \lambda \int_{\theta_a}^{\overline{\theta}} f_a(x|a)dx.$$

where $\kappa \geq 0$. That is,

$$0 < \mathbb{E}_\theta \left[\frac{1}{v'(t^O)} \right] - \frac{1}{v'(t^O(\theta))}|_{\theta \leq \theta_a} = \lambda \int_{\theta_a}^{\overline{\theta}} f_a(x|a)dx.$$

(unless $v' = \infty$ for some t and that t is a constant). For any increasing t^O on some measure of Θ, the inequality must hold as $1/v'$ is increasing. Because $f_a/f \geq 0$ on $\left[\theta_a, \overline{\theta}\right]$ and strictly for at least a positive measure, $\lambda > 0$ necessarily. ■

Proof of Proposition 1: Fix a. Rewrite the first-order condition as $v'(t^O) = (\kappa + \lambda f_a/f)^{-1}$; let $h \equiv (v')^{-1}$. The function $h(\cdot)$ is continuous because v' is also continuous, so $t^O \equiv h\left([\kappa + \lambda f_a/f]^{-1}\right)$ is a continuous function. To show continuity at θ_a, recall that $\lambda \frac{f_a}{f}|_{\theta_a} = 0$ and f_a/f is continuous in θ, so continuity at θ_a follows. For the second part of the Proposition, restrict attention to $\theta \geq \theta_a$ and define $\tau(\theta) \equiv t^O \circ m(\theta)$. Then rewrite the FOC as $v'(\tau) - \left(\kappa + \lambda \frac{f_a}{f}\right)^{-1} = 0$, where $\tau(\theta)$ is a.e. differentiable; differentiate w.r.t. θ to find $v''\tau' + \lambda \frac{d}{d\theta}\left(\frac{f_a}{f}\right) / \left(\kappa + \lambda \frac{f_a}{f}\right)^2 = 0$. This verifies $\tau' > 0$ and therefore $t' > 0$ as required since $\frac{dm}{d\theta} > 0$. Re-arrange this expression and redefine the variables

$$\tau' = -\lambda \underbrace{\frac{1}{v''}}_{Y} \underbrace{\frac{\frac{d}{d\theta}\left(\frac{f_a}{f}\right)}{\left(\kappa + \lambda \frac{f_a}{f}\right)^2}}_{X}$$

Then $\tau'' \geq 0 \Leftrightarrow \left(\frac{dY}{d\theta}X + \frac{dX}{d\theta}Y\right) \leq 0$. With $Y < 0$, rewrite the second condition as

$$\frac{dY}{d\theta}X \leq -\frac{dX}{d\theta}Y \Leftrightarrow \frac{d}{d\theta}\ln -Y \leq \frac{d}{d\theta}\ln X,$$

$$\frac{d}{d\theta}\ln -\frac{1}{v''} \leq \frac{d}{d\theta}\ln\left(\frac{\frac{d}{d\theta}\left(\frac{f_a}{f}\right)}{\left(\kappa + \lambda\frac{f_a}{f}\right)^2}\right)$$

Since the ratio $\frac{f_a}{f}$ is increasing concave, the RHS is negative. It is immediate to verify by differentiation that the LHS is positive, so the necessary and sufficient condition cannot hold. Hence $\tau'' < 0$ (where it is differentiable), that is, the effective transfer $\tau(\theta)$ is concave in the type. To show it is concave in the *message*, call on Lemma 7 and observe that τ is the composition of the function $t(\cdot)$ and the convex function $m(\theta)$. Therefore $t(\cdot)$ must be concave in m. For the last item, observe that at $\tilde{\theta}$, $m(\tilde{\theta}) = \tilde{\theta}$ by (3.2) – the agent is truthful. Thus, under $t^O(\cdot)$:-

$$\begin{aligned} v(t^O(\tilde{\theta})) &= [1 - p(m(\tilde{\theta}) - \tilde{\theta})]v(t^O(m(\tilde{\theta}))) = v(t^O(m(\tilde{\theta}))) \\ \Leftrightarrow \quad t^O(\tilde{\theta}) &= t^O(m(\tilde{\theta})) \end{aligned} \tag{8.1}$$

directly from (3.2). From Lemma 1, $t^O(m(\theta)) = t^{SB}(m(\theta))$ for $\theta \leq \tilde{\theta}$ and $t^O(\theta) = t^S(\theta)$ for $\theta > \tilde{\theta}$. Both these transfer functions are continuous on their respective domains. Thus by (8.1) I have shown that $\lim_{\theta\uparrow\tilde{\theta}} t(m(\theta)) = t^O(m(\tilde{\theta})) = t^O(\tilde{\theta}) = \lim_{\theta\downarrow\tilde{\theta}} t(\theta)$, which is the definition of continuity. Next, the right-derivative of t^O at $\tilde{\theta}$ can be denoted $\frac{dt^O}{d\theta}|_{\tilde{\theta}}$, while the left-derivative is $\frac{dt^O}{dm}\frac{dm}{d\theta}|_{\tilde{\theta}}$, where $dm/d\theta|_{\tilde{\theta}} = 1$ since $m(\theta) = \theta$ at this point. Using this one more time, $\frac{dt^O}{dm}\frac{dm}{d\theta}|_{\tilde{\theta}} = \frac{dt^O}{d\theta}|_{\tilde{\theta}}$; i.e. the left- and right-derivative are identical at $\tilde{\theta}$, which defines differentiability. Last, any amount lower than zero is not binding. Take t^O to be zero below θ_{a^O}. Then by application of (3.2) and Lemma 6, $\tilde{\theta} > \theta_{a^O}$. (There is a kink at θ_{a^O}, so Lemma 6 precludes $\tilde{\theta} = \theta_{a^O}$. Economically, the LHS of (3.2) is the marginal benefit of misreporting and the RHS the marginal cost; at θ_{a^O} the former is positive but the latter is 0, so it cannot be the point of indifference.) All things otherwise equal, having $\tilde{\theta}$ interior is costly to the principal in that the expected transfer is higher (otherwise the agent would not misreport) and so is the agent's optimal action. So the principal may have incentives to lower $\tilde{\theta}$. The smallest possible change, $d\theta$, requires a fixed $\gamma > 0$ to be paid for *all* types (not just below θ_{a^O}). So the increase in expected cost is $\gamma > 0$, and because $d\theta$ has measure zero, it alters neither the agent's moral hazard constraint (4.3) nor her information revelation problem (4.1). Calling on continuity completes the argument for any measure $\int d\theta$. ■

Proof of Corollary 2: Take any two $\theta_1 < \theta_2$ and suppose that truthful revelation holds at θ_1, i.e. $v't'(\theta_1) \leq p'(0)v(t(\theta_1)$. Because t^O is everywhere non-decreasing and concave (and so is $v(\cdot)$), it must therefore be that $v't'(\theta_2) \leq v't'(\theta_1) \leq p'(0)v(t(\theta_1) \leq p'(0)v(t(\theta_2)$. Therefore the agent also reveals herself truthfully at θ_2; she does not jump away from truth-telling. ■

Proof of Proposition 2: Construct the Lagrangian with the objective function and the constraints (4.1)-(4.3). Apply the Envelop Theorem to the first constraint. Because $\tilde{\theta} \equiv \tilde{\theta}(\alpha, t)$, Leibnitz rule gives an additional term (e.g. $p(m(\tilde{\theta}) - \tilde{\theta}; \alpha)t(m(\tilde{\theta}))f(\tilde{\theta}|a)\frac{d\tilde{\theta}}{d\alpha})$. But it is naught at $\tilde{\theta}$, where $m(\tilde{\theta}) = \tilde{\theta}$. This gives the first-order conditions found in Lemma 1, as well as (4.8). When $\tilde{\theta} = \underline{\theta}$, this latter condition is meaningless. In this case the level of investment is determined by (3.2) at $\underline{\theta}$, i.e. (4.7). ■

Proof of Proposition 3: The following will be useful in several instances. Let a^* solve the agent's moral hazard constraint (4.3). Differentiate (4.3) with respect to t:

$$\begin{aligned} 0 &= \int_{\underline{\theta}}^{\tilde{\theta}} v'[1-p]dF_a(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v'dF_a(x|a) \\ &+ \left[\int_{\underline{\theta}}^{\tilde{\theta}} v(t(m(x)))[1-p(m(x)-x)]dF_{aa}(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t(x))dF_{aa}(x|a) - c''(a)\right]\frac{da^*}{dt} \end{aligned} \quad (8.2)$$

Since the term in the brackets is the agent's second-order condition, it is negative. Therefore $\frac{da^*}{dt} > 0$. To prove item *(i)*, consider two distributions $F^1(\theta|a)$ and $F^2(\theta|a)$, where F^2 is a mean-preserving spread of F^1 (see Rothschild and Stiglitz [31]). Fix t; because F^1 dominates F^2 in the second order sense, it follows from (4.3) that at a^*

$$\int_{\underline{\theta}}^{\tilde{\theta}} v[1-p]dF_a^2 + \int_{\tilde{\theta}}^{\overline{\theta}} vdF_a^2 < \int_{\underline{\theta}}^{\tilde{\theta}} v[1-p]dF_a^1 + \int_{\tilde{\theta}}^{\overline{\theta}} vdF_a^1 \quad (8.3)$$

by application of the envelop theorem (to the messages). Now define the following variable $\theta_2 = \theta_1 + \epsilon$, where $\theta_2 \sim F^2$ and $\theta_1 \sim F^1$ (so θ_2 is more risky than θ_1, and (8.3) follows).

Consider again (4.3), as under F^1, and differentiate with respect to ϵ at $\epsilon = 0$:

$$\left[\int_{\underline{\theta}}^{\tilde{\theta}} v(t(m(x)))[1-p(m(x)-x)]dF_{aa}^1(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t(x))dF_{aa}^1(x|a) - c''(a)\right]\frac{da}{d\epsilon}$$
$$+\frac{d}{d\epsilon}\left[\int_{\underline{\theta}}^{\tilde{\theta}} v[1-p]dF_a^1 + \int_{\tilde{\theta}}^{\overline{\theta}} vdF_a^1\right] \quad = \quad 0$$

By (8.3) the last term is negative, so from (8.2) $\frac{da}{d\epsilon} < 0$. Letting $\frac{da}{d\epsilon} \equiv \frac{da}{dt}\frac{dt}{d\epsilon}$, $\frac{dt}{d\epsilon} < 0$ as claimed. To show (ii), consider a family of utility functions $v(t;r)$ parametrized by r; risk aversion (i.e. the concavity of $v(\cdot;\cdot)$) increases in r. Suppose for simplicity that $v(t;r)$ is continuous and differentiable in r (as well as t). For a fixed action a, we know that

$$\frac{d}{dr}\left[\int_{\underline{\theta}}^{\tilde{\theta}} v(t;r)[1-p]dF(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t;r)dF(x|a)\right] < 0$$

using the envelop theorem again. That is, equivalently, for any two $r_2 > r_1$, $\int_{\underline{\theta}}^{\tilde{\theta}} v(t;r_2)[1-p]dF(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t;r_2)dF(x|a) < \int_{\underline{\theta}}^{\tilde{\theta}} v(t;r_1)[1-p]dF(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t;r_1)dF(x|a)$. It then follows from (4.3) that $a^*(r_2) < a^*(r_1)$; equivalently, differentiating (4.3)

$$\begin{aligned} 0 \quad &= \quad \frac{d}{dr}\left[\int_{\underline{\theta}}^{\tilde{\theta}} v(t;r)[1-p]dF_a(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t;r)dF_a(x|a)\right] \\ &+ \quad \frac{da}{dr}\left[\int_{\underline{\theta}}^{\tilde{\theta}} v(t;r)[1-p]dF_{aa}(x|a) + \int_{\tilde{\theta}}^{\overline{\theta}} v(t;r)dF_{aa}(x|a) - c''(a)\right] \end{aligned} \tag{8.4}$$

Because the first term of (8.4) is negative it follows that $\frac{da}{dr} < 0$ as well. Making use of the fact that $\frac{da}{dt} > 0$ completes the argument. To prove (iii), consider two cost functions $c_1(a), c_2(a)$ such that $\forall a, c_2 > c_1$. Because $c_i', c_i'', c_i''' > 0, c_2 > c_1\ \forall a$ implies $c_2' > c_1'\ \forall a$. Fix t, from (4.3) we have that $a^*(c_2) < a^*(c_1)$. By (8.2) therefore $t(\theta, c_2) < t(\theta, c_1)\ \forall \theta$ (with obvious notation). For the last item, suppose the principal's payoff is some increasing function $\pi(\theta)$. From (4.6) it follows that a^O increases, and from (8.2) so does the transfer t. To complete the proof, apply Lemma 2. ■

References

[1] Araujo A. and H. Moreira (2001) "A general Lagrangian approach for non-concave moral hazard problems. " *Journal of Mathematical Economics*, Vol. 35, pp. 17-39

[2] Arya, A., Jonathan Glover and Shyam Sunder (1998) "Earnings management and the Revelation Principle. " *Review of Accounting Studies*, Vol. 3, pp. 7-34

[3] Baron D. and David Besanko "Regulation, Asymmetric Information, and Auditing." *The RAND Journal of Economics*, Vol. 15, No. 4, pp. 447-470

[4] Border, K. and J. Sobel (1987) "Samuarai accounting: a Theory of Auditing and Plunder" *The Review of Economic Studies*, Vol. 54, pp. 525-540

[5] Bushman, R. and Chandra Kanodia (1996) "A Note on Strategic Sampling in Agencies." *Management Science*, Vol.42, pp. 151-156

[6] Conlon, John R. (2009) "Two New Conditions Supporting the First-Order Approach to Multisignal Principal-Agent Problems." *Econometrica*, Vol. 77(1), pp. 249-278

[7] Crawford, V. and Joel Sobel (1982) "Strategic Information Transmission." *Econometrica*, Vol. 50, pp. 1431-1451

[8] Demski, J. and Ronald Dye (1999) "Risk, Returns and Moral Hazard" *Journal of Accounting Research*, Vol. 37, No. 1, pp. 27-55

[9] Demski, J. and Hans Frimor (1999) "Performance measure garbling under renegotiation in multi-period agencies." *Journal of Accounting Research*, Vol. 37, Supplement, pp. 187-214

[10] Dye, R. and Sanjay Sridhar "Moral Hazard Severity and Contract Design" *The RAND Journal of Economics*, Vol. 36, No. 1, pp. 78-92

[11] Gromb, D. and David Martimort (2007), "Collusion and the organization of delegated expertise" *Journal of Economic Theory*, Vol. 137 (1), pp. 271-299

[12] Grossman, S. and Oliver Hart (1983), "An Analysis of the Principal-Agent Problem" *Econometrica*, vol. 51(1), pages 7-45.

[13] Holmström, B. (1979) "Moral hazard and observability." *The Bell Journal of Economics*, Vol. 10, pp. 74-91

[14] Holmström, B. and Paul Milgrom (1987) "Aggregation and Linearity in the Provision of Intertemporal incentives." *Econometrica*, Vol. 55, pp. 303-328

[15] Holmström, B. and Paul Milgrom (1991) "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Description." *Journal of Law, Economics, & Organization*, Vol. 7, pp. 24-52

[16] Jewitt, I. (1988) "Justifying the first-order approach to principal-agent problems." *Econometrica*, Vol. 56, pp. 1177-1190

[17] Jewitt, I. Ohad Kadan and Jeroen Swinkels "Moral hazard with bounded payments", *Journal of Economic Theory* Vol. 143, pp. 59-82

[18] Kanodia, C. (1985) "Stochastic Monitoring and Moral Hazard." *Journal of Accounting Research*, Vol.23, pp. 175-194

[19] Kartik, N (2009) "Strategic communication with lying costs", *Review of Economic Studies*, 76(4), pp. 1359-1395.

[20] Kedia, S. and Thomas Philippon (2006) "The economics of fraudulent accounting", *NYU working paper*

[21] Khalil, F. (1997) "Auditing without commitment. "*Rand Journal of Economics*, Vol. 28, pp. 629-640

[22] Krähmer D. and Roland Strausz "Optimal procurement contract with pre-project planning.", *Review of Economic Studies*, 78(3), pp. 1015-1041.

[23] Laffont J.-J. and David Martimort (2002) "The Theory of Incentive – the Principal-Agent model." *Princeton University Press*

[24] Levitt, S. and Christopher Snyder (1997) "Is no News Bad News? Information Transmission and the Role of "Early Warning" in the Principal-Agent Model." *Rand Journal of Economics*, Vol. 28 (4), pp. 641-661

[25] Malcolmson, J. (2009) "Principal and expert agent." *The B.E. Journal of Theoretical Economics*, Vol. 9 (1), Article 17.

[26] Mookherjee D. and Ivan Png (1989) "Optimal Auditing, Insurance, and Redistribution." *The Quarterly Journal of Economics*, vol. 104(2), pp. 399-415

[27] Reinganum J. and Louis Wilde (1985) "Income Tax Compliance in a Principal-Agent Framework." *The Journal of Public Economics*, vol. 26, pp. 1-18.

[28] Reinhart, C. and Kenneth Rogoff "From financial crash to debt crisis." *American Economic Review*, vol. 101 (5) pp. 1677-1706.

[29] Roger, G. (2012) "Moral hazard with soft information." *Working paper, UNSW*

[30] Rogerson, W. (1985b) "The first-order approach to principal-agent problems." *Econometrica*, Vol. 53, pp. 1357-1368.

[31] Rothschild, M. and Joseph Stiglitz (1970) "Increasing Risk I: a Definition. " *Journal of Economic Theory*, Vol. 2, pp. 315-329