Simple Contracts with Adverse Selection and Moral Hazard*

Daniel Gottlieb and Humberto Moreira†

First Version: August, 2011. This Version: January, 2015.

Abstract

We study a principal-agent model with both moral hazard and adverse selection. Risk-neutral agents with limited liability have arbitrary private information about the distribution of outputs and the cost of effort. We obtain conditions under which the optimal mechanism offers a single contract to all types. These conditions are always satisfied, for example, if output is binary or if the distribution of outputs is multiplicatively separable and ordered by FOSD (if it is not ordered, the optimal mechanism offers at most two contracts). If, in addition, the marginal distribution satisfies the monotone likelihood ratio property, this single contract is a debt contract. Our model suggests that offering a single contract may be optimal in environments with adverse selection and moral hazard, where offering flexible menus of contracts provides gaming opportunities to the agent.

The unavoidable price of reliability is simplicity. – Charles Antony Richard Hoare
Simplicity does not precede complexity, but follows it. – Alan Perlis

*We thank Eduardo Azevedo, Vinicius Carrasco, Sylvain Chassang, Alex Edmans, Faruk Gul, Philippe Jehiel, Jason Hartline, Bengt Holmstrom, Lucas Maestri, George Mailath, David Martimort, Roger Myerson, Stephen Morris, Luca Rigotti, Yuliy Sannikov, Jean Tirole, Rakesh Vohra, John Zhu, and seminar audiences at Arizona State University, FGV, HEC Montreal, Johns Hopkins University, Princeton University, PUC-Rio, Universidad de Chile, University of Pennsylvania, University of Pittsburg/Carnegie Mellon University, the Wharton School, and the BYU Computational Public Economics, 2013 LAMES, 2013 SBE, 2014 IWGTS, 2014 ESEM, and 2015 AEA meetings for comments and suggestions. Rafael Mourão provided outstanding research assistance. Gottlieb gratefully acknowledges financial support from the Dorinda and Mark Winkelman Distinguished Scholar Award. Moreira acknowledges CNPq for financial support.

†Gottlieb: The Wharton School, The University of Pennsylvania, dgott@wharton.upenn.edu. Moreira: FGV/EPGE, humberto@fgv.br.

Contents

1 Introduction

Most real-world contracts are much simpler than theory predicts. Differently from standard adverse selection models, contracting parties offer a limited number of contracts, often a single one. Unlike in standard moral hazard models, similar contracts are offered in fundamentally different environments. As *Hart and Holmstrom (1987)* and *Chiappori and Salanie (2003)* argue in their surveys of the literature:

> The extreme sensitivity to informational variables that comes across from this type of modeling is at odds with reality. Real world schemes are simpler than the theory would dictate and surprisingly uniform across a wide range of circumstances. *(Hart and Holmstrom, 1987, pp. 105)*

> The recent literature (...) provides very strong evidence that contractual forms have large effects on behavior. As the notion that "incentives matter" is one of the central tenets of economists of every persuasion, this should be comforting to the community. On the other hand, it raises an old puzzle: if contractual form matters so much, why do we observe such a prevalence of fairly simple contracts? *(Chiappori and Salanie, 2003 , pp. 34)*

In this paper, we propose an answer to this puzzle based on the interaction between adverse selection, moral hazard, and limited liability. Most contracting situations have both adverse selection and moral hazard. Managers, for example, take actions that affect the firm's profitability. At the same time, they usually have better knowledge about the efficacy of each action. Moreover, virtually all contracting parties have limited liability. Entrepreneurs raising capital from investors, for example, enjoy limited liability as the value of their equity cannot fall below zero. Also, anti-slavery laws enforce limited liability in employment contracts.

We consider a principal-agent relationship with bilateral risk neutrality, free disposal, and limited liability. The agent selects an unobserved "effort," which may consist of a single or multiple tasks. The agent also has private information, in an arbitrary way, about the distribution of outputs and about effort costs, resulting in a model where types and efforts are multidimensional (possibly infinite dimensional) and unordered. We show that the interaction between adverse selection, moral hazard, and limited liability imposes severe screening costs.

With binary outcomes, the optimal mechanism involves offering a single contract to all agents regardless of type space or the distribution of types. We generalize this result to settings with multiple outputs under a multiplicative separability condition. This condition – satisfied, for example, under the spanning condition of Grossman and Hart (1983) or if output is binary – is equivalent to assuming that agents rank the "power" of all contracts equally. The optimal mechanism involves either one or two contracts depending on whether the effort space is ordered by first-order stochastic dominance. Additionally, if the marginal distribution of output satisfies the monotone likelihood ratio property, the optimal contract consists of the principal taking a single debt contract or, equivalently, giving all agents the same call option.

More broadly, our paper identifies an important downside from offering flexibility to agents through menus of contracts: gaming. This is particularly stark in the model considered here, where both the principal and the agent are risk neutral. Then, the agent always selects the contract with the highest expected payment conditional on his effort, which is precisely the contract with the highest cost to the principal. That is, conditional on the effort that the agent choses, reducing the number of contracts always increases the principal's profits (for a fixed effort). In particular, when the principal can identify a contract with the highest power in a mechanism (i.e., when multiplicative separability holds), she can simultaneously reduce informational rents and increase efficiency by removing all other contracts.

Although the framework we study has been widely applied to financial contracting, it has many other applications. One such application is to procurement and regulation. Despite the central role that menus of contracts play in the theory of procurement and regulation, they are rarely observed in practice.[1] Accordingly, many papers try to identify conditions for simple procurement contracts to be close to optimal.[2]

We consider an extension of the classic model of Laffont and Tirole (1986, 1993), where effort is allowed to affect the regulated firm's costs stochastically and the firm is subject to limited liability. Our results then determine conditions for the optimal mechanism to offer a single contract and for this contract to be a price cap. Since limited liability constraints are a key aspect of most procurement contracts (see, e.g., Burguet et al., 2012), our model provides an explanation for the lack of menus of contracts in procurement.

Our result on the optimality of simple contracts is related to the robustness intuition of Holmstrom and Milgrom (1987). However, the notion of robustness in our static model is different from the one in their seminal paper. Here, offering a single contract is robust in that it reduces the agents' incentives to misrepresent their private information about the environment. In Holmstrom and Milgrom's model, linear contracts are robust in the sense that they prevent the agent from readjusting effort over time.[3] Moreover, as in their work, we also contribute to the

[1] For example, Bajari and Tadelis (2001) argue that "the descriptive engineering and construction management literature (...) suggests that menus of contracts are not used. Instead, the vast majority of contracts are variants of simple fixed-price (FP) and cost-plus (C+) contracts."

[2] Using the Laffont-Tirole framework, Rogerson (2003) and Chu and Sappington (2007) show that a pair of simple contracts can achieve a large fraction of the surplus under a certain range of parametric settings – 75 or 73 percent when costs follow either uniform or power distributions, respectively – for quadratic costs. Bajari and Tadelis (2001) assume that there is a fixed cost of specifying each state of nature in the contract to rationalize the simplicity of observed contracts.

[3] Edmans and Gabaix (2011) extend the linearity results to a model in which the realization of noise occurs before the action. Chassang (2013) introduces a class of calibrated contracts that are detail-free and approximate the performance of the best linear contract in dynamic environments when players are patient, while Carroll (Forthcoming) shows that the best contract for a principal who faces an agent with uncertain technology and

applied literature by identifying assumptions under which researchers can focus on a simpler set of contracts when solving their models. For example, under the standard assumptions from Innes (1990) (and, when there are more than two outputs, a multiplicative separability assumption), there is no loss of generality in assuming that the optimal mechanism involves a single debt contract even if there is adverse selection. It is then easy to obtain comparative statics results in this environment.

Related Literature

Our model consists of a principal-agent relationship with bilateral risk neutrality, free disposal, and limited liability. Starting with the canonical work of Innes (1990), many researchers have studied models with either moral hazard or adverse selection in this environment. Moral hazard models include, for example, Matthews (2001), Dewatripont et al. (2003), Poblete and Spulber (2012), and Chaigneau et al. (2014). Adverse selection models include Nachman and Noe (1994), Demarzo and Duffie (1999), DeMarzo (2005), and DeMarzo et al. (2005).[4]

Our work is related to a literature that identifies conditions for contracts to take the form of debt and for equilibria to have complete pooling. In a single-task moral hazard setting, Innes (1990) and Poblete and Spulber (2012) show that contracts take the form of debt if the distribution of output satisfies the monotonicity of the likelihood ratio property. We build on their environment by adding adverse selection in an arbitrary way and allowing effort to be multi-dimensional. Our main focus is on the lack of menus of contracts, which, of course, can only be addressed by introducing adverse selection. Nevertheless, our Theorem 3 is reminiscent of their main result.

In a signaling model of financial contracting, Nachman and Noe (1994) show that there is complete pooling if and only if firm types are strictly ordered by conditional stochastic dominance. When firms are ordered by stochastic dominance, capitalists face a lemons problem: while they would like to offer better terms to healthier firms, those who are more likely to accept the contract are precisely the least healthy firms. Our papers emphasize different forces that may lead to pooling. In Nachman and Noe (1994), pooling occurs when the distribution of types induces a market breakdown for all but the worse contract. In our model, pooling happens because of moral hazard and limited liability: giving flexibility to agents requires the principal to leave excessive rents. For example, when output is binary, complete pooling occurs in our model for any parameters of the model (i.e., regardless of whether types are ordered).[5]

Similarly, Demarzo and Duffie (1999) consider a signaling model of security design and show

evaluates contracts in terms of their worst-case performance is linear.

[4]For a general analysis of moral hazard models with limited liability, see Jewitt et al. (2008).

[5]See footnote 15 for a discussion of Nachman and Noe's requirement that types be ordered by conditional stochastic dominance in the context of our model.

that, under a uniform-worst case condition, equilibrium contracts take the form of debt. In this model, several authors studied whether intermediaries pool different assets in equilibrium. The conclusion depends on whether the security designed before or after firms learn about the asset's profitability (c.f., DeMarzo (2005), Biais and Mariotti (2005), and Farhi and Tirole (Forthcoming)).

In a one-dimensional adverse selection setting, Guesnerie and Laffont (1984) showed that optimal mechanisms are "non-responsive" when the first-best allocation is decreasing. This occurs because optimality clashes with incentive compatibility, which requires allocations to be non-decreasing. The reason for pooling in our model is quite different from non-responsiveness. For example, if the agent only has private information about the distribution of output, the first best is increasing and, therefore, implementable in a pure adverse selection environment. Nevertheless, with multiplicative separability, the principal offers a single contract to all types (see footnote 11). More related to our work, Ollier and Thomas (2013) substitute the traditional (interim) participation constraint by an ex-post constraint in a one-dimensional model with binary outcomes. They show that, under conditions that ensure that the first-order approach holds, there is no benefit from screening.

As argued previously, our application to procurement and regulation builds on Laffont and Tirole (1986, 1993). In their model, there is both adverse selection and moral hazard. However, because the link between effort, types, and output is deterministic, the model can be reduced to a pure adverse selection model.[6] We allow effort to affect the regulated firm's costs stochastically so the problem cannot be reduced to a pure adverse selection model. Picard (1987), Melumad and Reichelstein (1989), and Caillaud et al. (1992) also introduce noise in the relationship between output and effort and show that, under certain conditions, the principal can achieve the same utility as in the absence of noise. Therefore, unlike in our model, they find that there is no cost from moral hazard. Our model differs from theirs in two ways. First, we also allow the agent to have private information about the distribution of output, while they assume that all private information concerns the cost of effort. Introducing private information about the distribution makes moral hazard costly. Second, they do not assume that agents have limited liability. Limited liability also prevents the principal from eliminating moral hazard at no cost.

In Section 2, we present the model with two outputs and discuss the benchmark cases of pure adverse selection and pure model hazard. In Section 3, we generalize the results for multiple outputs and obtain conditions for the optimality of debt. Then, Section 4 concludes. The extension of our model to procurement and regulation, as well as all proofs, are in the appendix.

[6] For this reason, Laffont and Martimort (2002) refer to them as 'false moral hazard' models.

2 Two Outputs

2.1 Statement of the Problem

We start with the two-output model. There is a risk-neutral principal and a risk-neutral agent with limited liability. The agent has private information about the environment, captured by a type $\theta \in \Theta$. From the principal's perspective, types are distributed according to a distribution μ. We will discuss the assumptions on Θ and μ below.

The agent exerts an effort $e \in E$, which costs c_e^θ. The space of possible efforts E is a compact metric space. Effort can consist of a single task ($E \subset \mathbb{R}$) or multiple tasks ($E \subset \mathbb{R}^N$). The least-costly effort has a non-positive cost: $\min_{e \in E} c_e^\theta \leq 0$. This condition is satisfied in standard frameworks where the lowest effort costs zero, as well as in more general multi-task frameworks that allow the agent to derive private benefits from certain actions.[7]

The principal does not observe the effort chosen by the agent. She does, however, observe the output from the partnership $x \in \{x_L, x_H\}$, which is stochastically affected by the agent's effort. We refer to x_H as a high output or as success, to x_L as a low output or failure, and to $\Delta x := x_H - x_L > 0$ as the incremental output. Given effort e, a high output happens with probability p_e^θ.

The agent has private information about the distribution of outputs and the cost of effort. The type space Θ may be discrete or continuous, and types may be finite- or infinite-dimensional. Each type is fully characterized by the the pair of functions $\left(p_\cdot^\theta, c_\cdot^\theta\right)$ specifying the probability of success and the cost of each effort.[8] For example, if effort is binary, each type can be described by the four-dimensional vector $(p_0^\theta, p_1^\theta, c_0^\theta, c_1^\theta)$. If effort is continuous, each type is described by an infinite-dimensional function $\left(p_\cdot^\theta, c_\cdot^\theta\right) : E \to \mathbb{R}^2$. Our model does not require the agent to have private information about all of these dimensions, of course. The case where the cost of effort e is common knowledge, for example, is accommodated by letting c_e^θ be constant in θ. Note that we do not impose any order on the space of types and efforts. Success probabilities and costs may be non-monotone functions and, moreover, types and effort may be complements, substitutes, or neither in terms of probabilities and costs.

By the revelation principle, we can focus on direct mechanisms. A direct mechanism is a triple of $\mathcal{B}(\Theta)$-measurable functions $(w, b, e) : \Theta \to \mathbb{R}^2 \to E$, consisting of fixed payments w, bonuses b, and effort recommendations e. An agent who reports type θ agrees to exert effort $e(\theta)$ and receives $w(\theta)$ in case of failure and $w(\theta) + b(\theta)$ in case of success. A pair of payments $w(\theta)$ and $b(\theta)$ is called a *contract*.

[7] Allowing the cost of the lowest effort to be positive makes participation random as in Rochet and Stole (2002) and is beyond the scope of this paper.

[8] We write $p_\cdot^\theta$ to denote the function $e \longmapsto p_e^\theta$ that keeps θ constant and varies e. Similarly, $p_e^\cdot$ refers to $\theta \longmapsto p_e^\theta$. The same notation is used to all functions in the paper.

Given a mechanism (w, b, e), a type-θ agent gets payoff

$$U(\theta) := w(\theta) + p^{\theta}_{e(\theta)} b(\theta) - c^{\theta}_{e(\theta)}. \tag{1}$$

The mechanism must satisfy the following incentive compatibility (IC) and participation (IR) constraints:

$$U(\theta) \geq w\left(\hat{\theta}\right) + p^{\theta}_{\hat{e}} b\left(\hat{\theta}\right) - c^{\theta}_{\hat{e}}, \quad \forall \theta, \hat{\theta}, \hat{e}, \tag{IC}$$

$$U(\theta) \geq 0, \quad \forall \theta. \tag{IR}$$

It must also satisfy the following free disposal (FD) constraint:

$$b(\theta) \geq 0, \quad \forall \theta. \tag{FD}$$

Free disposal must be satisfied if the agent can costlessly reduce output or if the principal can secretly borrow from an outside lender to inflate output.

Finally, the agent is protected by limited liability, which prevents payments from being negative. Since, by free disposal, bonuses are non-negative, we can write the agent's limited liability (LL) constraint as

$$w(\theta) \geq 0, \quad \forall \theta. \tag{LL}$$

An *optimal* mechanism maximizes the principal's expected profit

$$\int_{\Theta} \left\{ \, p^{\theta}_{e(\theta)} \left[x_H - \left(w(\theta) + b(\theta)\right)\right] + \left(1 - p^{\theta}_{e(\theta)}\right) \left[x_L - w(\theta)\right] \, \right\} d\mu(\theta) \tag{2}$$

among mechanisms that satisfy IC, IR, FD, and LL. To ensure the existence of an optimal mechanism, we make the following technical assumptions, which are satisfied by all standard agency models:

Assumption 1. Θ is a complete separable metric space. μ is a probability measure defined on the Borel σ-field of Θ, which we denote by $\mathcal{B}(\Theta)$. For each $\theta \in \Theta$, $p^{\theta}_{\cdot}$ and $c^{\theta}_{\cdot}$ are continuous functions and, for each $e \in E$, $p^{\cdot}_{e}$ and $c^{\cdot}_{e}$ are $\mathcal{B}(\Theta)$-mensurable functions.

2.2 Benchmarks

We first consider the benchmark cases of pure moral hazard and pure adverse selection. We show that, in both cases, the principal typically offers a different contract to each type. Therefore, the principal only prefers to offer the same contract to all types when moral hazard and adverse selection co-exist. For simplicity, we focus on the single-task case ($E \subset \mathbb{R}$) and assume that the probability of high output p^{θ}_{e} and the cost of effort are both non-decreasing in e (i.e., effort increases the probability of success at a cost). We normalize the lowest effort to zero.

Pure Moral Hazard

Suppose the principal observes the agent's type but does not observe effort. Without limited liability, the principal can implement the first best by "selling the firm" to each agent – i.e., paying a bonus equal to the incremental output $b(\theta) = \Delta x$ and offering a fixed wage that extracts the entire surplus $w(\theta) = c^{\theta}_{e(\theta)} - p^{\theta}_{e(\theta)}\Delta x$, where $e(\theta)$ is the first-best effort. With limited liability, the principal needs to leave positive rents to the agent if she wants to sell the firm. Then, it is profitable to distort the bonus downward, causing some types to exert less effort.[9] Moreover, limited liability binds, so the agent gets a zero fixed wage.

Optimal contracts with and without limited liability vary in opposite dimensions: While, without limited liability, they have the same bonus ($b = \Delta x$) and different fixed wages, optimal contracts with limited liability have the same fixed wage ($w = 0$) and different bonuses.[10] In both cases, however, the principal offers different contracts to different types. Moreover, these mechanisms are no longer feasible if types are unobservable. If offered contracts with the same bonus and different fixed wages, all types would select the one with the highest wage. Similarly, if offered the same fixed wage and different bonuses, they would all pick the contract with the highest bonus. The principal can still screen unobservable types by varying *both* the fixed wage and the bonus. In fact, without limited liability, this is typically optimal. Our main result shows that, with limited liability, the principal prefers not to offer a menu of contracts. Instead, the optimal mechanism offers a single contract despite the presence of many different types.

Pure Adverse Selection

Now suppose the principal observes the agent's effort but not his type. If effort costs are common knowledge (that is, θ only affects the conditional probability of success), the principal can implement the first best by fully reimbursing the cost of each effort. The agent would then be indifferent between all efforts and would, therefore, accept to pick the principal's preferred one.[11] The principal can no longer implement the first best, however, if effort costs are private information. If she offered to fully reimburse the effort costs of all types, they would all pretend

[9] To see this, note that if the principal wants to recommend the lowest effort, she will pay $w(\theta) = b(\theta) = 0$ and get expected payoff $p^{\theta}_0 x_H + \left(1 - p^{\theta}_0\right) x_L > x_L$. Suppose the principal pays $b(\theta) \geq \Delta x$ and the agent exerts effort e. The principal's payoff is then $p^{\theta}_e [x_H - b(\theta)] + \left(1 - p^{\theta}_e\right) x_L \leq p^{\theta}_e [x_H - \Delta x] + \left(1 - p^{\theta}_e\right) x_L = x_L$. Comparing these two inequalities, we can see that offering $b(\theta) \geq \Delta x$ is dominated by paying $w = b = 0$ and recommending $e = 0$. Then, because agents are paid a bonus lower than the incremental output, the efforts of all types $e(\theta)$ lie (weakly) below the first-best effort.

[10] For example, when there are only two efforts (say, 0 and 1) and the agent has limited liability, the principal offers a bonus $b(\theta) = \frac{c^{\theta}_1}{p^{\theta}_1 - p^{\theta}_0}$ if she recommends the high effort ($e(\theta) = 1$). This bonus is strictly increasing in c^{θ}_1 and p^{θ}_0 and strictly decreasing in p^{θ}_1.

[11] Note that, unlike in Guesnerie and Laffont (1984), the first-best allocation in this case is non-decreasing and is therefore feasible when the agent is not subject to moral hazard. Complete pooling in our model is due to the interaction between adverse selection, moral hazard and limited liability, not because of non-responsiveness.

to be the types with the highest costs.

If, on the other hand, the probability of success is common knowledge (i.e., θ only affects the cost of effort), the optimal mechanism posts a payment for each (observed) effort. Agents choose which effort to exert based on their privately-known costs.

As we show in the appendix, if the agent has private information about *both* probabilities and costs, the principal typically offers a menu of contracts to screen their private information. Agents with higher probability of success pick contracts with lower fixed payments and higher bonuses.

None of these mechanisms are feasible if effort is not observable. If the principal offered them, agents would pick the lowest effort and claim to have exerted the highest effort, violating incentive compatibility. In fact, as we show next, the principal's inability to observe effort limits her power to screen the agent's private information to the extent that it is optimal for her to offer a single contract.

2.3 Contract Simplicity

We can now state the simplicity result with binary outcomes, which establishes that the principal offers a single contract:

Theorem 1. *There exists an optimal mechanism that offers a single contract* (w, b) *to all types, with* $w = 0$ *and* $b < \Delta x$.

The optimal contract can be interpreted as a debt contract for the principal with face value $x_H - b \in (x_L, x_H)$. It can also be interpreted as giving the agent a call option on output with strike price $x_H - b$.

The proof is based on three lemmas. The first one shows that IC and LL imply that IR never binds. This follows from the fact that the agent can always guarantee himself a non-negative payoff by picking the lowest effort and collecting the non-negative payments. The second lemma shows that any mechanism that pays a bonus greater than the incremental output to some type cannot be optimal. Any such mechanism gives the principal a payoff that is lower than if she offered all types a constant payment of zero ($w = b = 0$).

The last lemma establishes that, for any mechanism with bonuses lower than the incremental output, if multiple contracts are being offered, the principal can improve by offering all types the contract with the highest bonus. Since IR never binds, all agents pick this single contract once all other contracts are removed. There are two effects from this migration to the highest-powered contract: a reduction in informational rents and an increase in efficiency. Since the agents are risk neutral, they pick the contract that maximizes expected payments conditional on their effort. Thus, holding effort fixed, reducing the set of contracts being offered decreases expected payments to the agents (rent extraction effect). Second, because agents now face a higher bonus, they choose an effort with a higher probability of success. This raises the principal's profit by $(p_{\tilde{e}} - p_e)[\Delta x - b]$, where e is the agent's old effort, $\tilde{e}$ is the agent's new effort, and $p_{\tilde{e}} \geq p_e$. Since the incremental output exceeds the bonus, both terms are positive. Hence, the efficiency effect from increasing effort is also positive. The proof then concludes by showing that an optimal mechanism exists.

Limited liability and risk neutrality play an important role in Theorem 1. Limited liability ensures that agents do not leave the mechanism if their contract is removed. Without it, the participation constraint would bind for some type. Then, removing low-powered contracts would induce some types to prefer not to participate. Risk neutrality implies that, holding effort fixed, the principal and the agent split a pie of a fixed size. Since the agent always picks the contract with the highest expected payment, providing more freedom of choice to the agent can only hurt the principal (holding effort fixed). With risk aversion, different bonuses also affect the size of the pie since lower bonuses insure the agent better. Then, removing all but the highest-powered contract improves efficiency but worsens risk sharing.

Theorem 1 greatly simplifies the analysis of the optimal mechanism by allowing us to rewrite the principal's program as a standard optimization problem with a single instrument $b \in [0, \Delta x]$. It is then straightforward to obtain comparative statics results. For example, using a supermodularity argument, we can show that the optimal bonus and the success probabilities of all types are increasing in the incremental output Δx. Moreover, the probability of success is distorted downwards relative to the first best. Formally, letting e^{FB} denote a first-best effort, we have:

$$e^{FB}(\theta) \in \arg\max_{e} x_L + \; p_e^\theta \Delta x - c_e^\theta \quad \text{and} \quad e(\theta) \in \arg\max_{e} \; p_e^\theta b - c_e^\theta \; .$$

Since $b < \Delta x$, it follows by revealed preferences that $p^\theta_{e(\theta)} \leq p^\theta_{e^{FB}(\theta)}$ with strict inequality for some type if the type space is sufficiently rich. Finally, notice that the theorem does not depend on the distribution of types or other parameters of the model.

It is straightforward to generalize the analysis above to the case of multiple outputs when contracts are restricted to take the form of two-part tariffs, where the fixed part corresponds to a wage that is paid independent of the output and the variable part corresponds to equity payments that are linear in the firm's output. The restriction to two-part tariffs can be motivated by arbitrage opportunities when the principal deals with multiple agents who can costlessly redistribute outputs between themselves. In the next section, we study optimal contracts without this restriction.

3 Multiple Outputs

We now generalize the model to allow for multiple outputs. Let $X \subseteq \mathbb{R}$ be a closed set of possible outputs, which may be discrete or continuous. The agent's private information is described by a type $\theta \in \Theta$, where Θ is a complete, separable, metric space. Types are distributed according to a probability measure μ defined on Borel σ-field of the type space $\mathcal{B}(\Theta)$. Notice that our formulation allows types to be finite- or infinite-dimensional, and their distribution may be discrete or continuous.

The agent chooses an unobservable effort e from the compact metric space E. A type-θ agent who exerts effort e produces output according to the cumulative distribution function $F_e^\theta(x)$. Let c_e^θ denote type θ's cost of of effort e. As before, we assume that the least-costly effort has a non-positive cost: $\min_e c_e^\theta \leq 0$ for all θ.

A *compensation contract* is a function that specifies a transfer to the agent conditional on each possible output. A mechanism specifies a compensation contract and an effort recommendation for each type. That is, a mechanism is a pair of $\mathcal{B}(\Theta)$-measurable functions $w : \Theta \times X \to \mathbb{R}$ and $e : \Theta \to E$. Therefore, a type-θ agent is recommended effort $e(\theta)$ and gets paid $w_\theta(x)$ in case of output x.

Given a mechanism (w, e), a type-θ agent gets expected payoff

$$U(\theta) := \int w_\theta(x)\, dF^\theta_{e(\theta)}(x) - c^\theta_{e(\theta)}.$$

As in the two-output case, the mechanism has to satisfy the following incentive compatibility, participation, and limited liability constraints:

$$U(\theta) \geq \int w_{\hat{\theta}}(x)\, dF^\theta_{\hat{e}}(x) - c^\theta_{\hat{e}}, \quad \forall \theta, \hat{\theta}, \hat{e}, \tag{IC}$$

$$U(\theta) \geq 0, \quad \forall \theta, \tag{IR}$$

$$w_\theta(x) \geq 0, \quad \forall \theta, x. \tag{LL}$$

It also has to satisfy the following bilateral free disposal (BFD) constraint:

$$0 \leq w_\theta(y) - w_\theta(x) \leq y - x, \tag{BFD}$$

for all θ, x, y with $y \geq x$. This constraint states that the payments of both the principal and the agent are non-decreasing. As Innes (1990) argues, BFD can be seen as an additional incentive constraint if the principal and the agent can costlessly reduce output or if they can borrow from outside lenders in order to inflate output. In the two-output framework from Section 2, BFD requires the bonus to lie between 0 and Δx. Since the optimal mechanism with two outputs always has a bonus lower than Δx, BFD is satisfied at the optimum – that is, the principal's free disposal constraint does not bind.

An *optimal* mechanism maximizes the principal's expected profit

$$\int_{\Theta}\int_{X}\left[x - w_{\theta}(x)\right] dF^{\theta}_{e(\theta)}(x)d\mu(\theta) \tag{3}$$

among mechanisms that satisfy IC, IR, LL, and BFD.

The following technical conditions, which generalize Assumption 1, are made to guarantee the existence of an optimal mechanism:

Assumption 2. i) For each $\theta \in \Theta$, $c^{\theta}_{\cdot}$ and $dF^{\theta}_{\cdot}(x)$ are continuous;

- **ii)** For each $e \in E$, $c^{\cdot}_{e}$ and $dF^{\cdot}_{e}(x)$ are $\mathcal{B}(\Theta)$-mensurable;[12]
- **iii)** For each $(\theta, e) \in \Theta \times E$, $dF^{\theta}_{e}(\cdot)$ is a probability measure on X; and
- **iv)** The expected output is μ-integrable on $\Theta \times E$. That is, $\left|\int x dF^{\theta}_{e}(x)\right| \leq \xi(\theta)$ for all θ, e, where $\xi : \Theta \to \mathbb{R}$ is an integrable function.

The main question we address in this section is whether there exists an optimal mechanism that offers the same compensation contract to all types, i.e., $w(\theta, x) = w(\tilde{\theta}, x)$ for all $\theta, \tilde{\theta}, x$. The example below shows that, without additional restrictions, the answer is no.

Example 1. There are two types $\Theta = \{A, B\}$ in equal proportion, two efforts $E = \{0, 1\}$, and three outcomes $X = \{L, M, H\}$ with $L < M < H$. Both types have the same effort costs: $c^{\theta}_{0} = 0$ and $c^{\theta}_{1} = 1$ for all θ. Their conditional probabilities are represented in the following table:

Type A			
	$x = L$	$x = M$	$x = H$
$e = 0$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
$e = 1$	$\frac{1}{5}$	$\frac{2}{5}$	$\frac{2}{5}$

Type B			
	$x = L$	$x = M$	$x = H$
$e = 0$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
$e = 1$	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{1}{3}$

Table 1: Conditional probabilities of each type.

Both types have the same probability distribution when they exert low effort. However, A observes both high and low outputs more frequently than B when he exerts high effort, whereas B observes the intermediate output more frequently than A. Notice that, for both types, the distribution of output with high effort first-order stochastically dominates the distribution with low effort. We assume that M and H are large enough for it to be optimal to recommend a high effort to both types.

As we show in the appendix, the optimal mechanism offers contract $w^{A} = (0, 0, 15)$ to type A and $w^{B} = (0, 6, 6)$ to type B. Referring to the incremental payment in each state as "the

[12] We use the weak topology on the space of distributions on X.

bonus," this mechanism pays type A a bonus of 15 in state $x = H$ and zero in other states. Type B gets a bonus of 6 in state M and zero in states L and H. The principal finds it profitable to separate them because their likelihood ratios are not ordered. Each type is paid a bonus in the state where his likelihood ratio is the highest ($x = H$ for type A and $x = M$ for type B).

The previous example illustrates the main problem in generalizing Theorem 1 to multiple outputs. With two outputs, the only way to incentivize effort is to pay a higher bonus. Thus, we can unequivocally rank the power of any pair of contracts by their bonuses.[13] With multiple outputs, there is one bonus associated with each incremental output, so contract power has, in general, only a partial order. Types with different probability distributions may be affected differently by bonuses paid in different states.

To rule out cases such as the one in Example 1, where types disagree over the effectiveness of incentives, we need to ensure that types order the power of different contracts in the same way. Formally, we need the following property to hold. Let w and $\tilde{w}$ be two contracts that satisfy FD and LL and let $e, \tilde{e} \in E$. If there exists θ for which

$$\int w(x) \left(dF_e^\theta(x) - dF_{\tilde{e}}^\theta(x)\right) dx \geq \int \tilde{w}(x) \left(dF_e^\theta(x) - dF_{\tilde{e}}^\theta(x)\right),$$

then, for all $\tilde{\theta}$,

$$\int w(x) \left(dF_e^{\tilde{\theta}}(x) - dF_{\tilde{e}}^{\tilde{\theta}}(x)\right) dx \geq \int \tilde{w}(x) \left(dF_e^{\tilde{\theta}}(x) - dF_{\tilde{e}}^{\tilde{\theta}}(x)\right).$$

This condition states that if one type has more incentives to exert a higher effort under w than under $\tilde{w}$, so do all other types. Of course, whether or not they will choose to exert the higher effort depends on the effectiveness and the cost of their efforts. But, because they all agree about which contracts have the highest power, the principal cannot increase effort from all types by offering contracts that pay higher bonuses in different states as in Example 1.

Although intuitive, this is a somewhat convoluted assumption. As we show in the appendix, however, this assumption is equivalent to the following *multiplicative separability* (MS) condition:

Definition 1. The distribution of outputs satisfies multiplicative separability (MS) if there exist functions $H : X \to \mathbb{R}$ and $I : E \times \Theta \to \mathbb{R}$ such that

$$F_{\tilde{e}}^\theta(x) + H(x)\, I(\tilde{e}, \theta) = F_e^\theta(x) + H(x)\, I(e, \theta) \quad \forall e, \tilde{e}, \theta, x. \tag{4}$$

Multiplicative separability always holds when there are only two outputs. It also holds under the Linearity of the Distribution Function Condition (Grossman and Hart (1983); Hart and

[13] Similarly, when the contract space is restricted to two-part tariffs, we can unequivocally rank the power of two contracts by their slopes.

Holmstrom (1987)), which is obtained by taking $E = [0,1]$ and $I(e,\theta) = e \times i(\theta)$:

$$F_e^\theta(x) = eF_1^\theta(x) + (1-e)\,F_0^\theta(x)\,. \tag{5}$$

This condition is commonly used in pure moral hazard models, along with the convexity of costs, to justify the first-order approach. Importantly, however, none of our results assume the validity of the first-order approach.[14]

We will say that a family of distributions $\mathcal{F}^\theta := \left\{F_e^\theta : e \in E\right\}$ is *ordered by first-order stochastic dominance (FOSD)* if it can be written as in equation (4) with $H(x) \geq 0$ for all x.[15] Under this condition, $F_e^\theta(x)$ first-order stochastically dominates $F_{\tilde{e}}^\theta(x)$ if and only if $I(e,\theta) > I(\tilde{e},\theta)$. That is, the family of distributions $\mathcal{F}^\theta$ is ordered by FOSD. Economically, in a family ordered by FOSD, the ranking of expected payments under different efforts is the same in all increasing compensation contracts. This is a usual assumption in single-task models, where we typically assume that effort orders the output distribution in terms of FOSD. In multi-task models, this assumption is less compelling. It fails, for example, if the agent allocates effort to a safe and a risky task, and effort allocated to the risky task increases both the mean and the variance of output. Notice that any distribution is ordered by FOSD if there are only two outputs.

The next theorem presents our main result. It shows that, whenever multiplicative separability holds, the optimal mechanism offers at most two contracts. If, in addition, the family of distributions is ordered by FOSD, the optimal mechanism offers a single contract to all types. Despite being offered the same contract, types may choose different efforts depending on their effort costs and output distributions.[16]

Theorem 2. *Suppose MS holds. There exists an optimal mechanism that offers at most two compensation contracts to all types. Moreover, if $\mathcal{F}^\theta$ is ordered by FOSD for all θ, there exists*

[14]There is no relationship between MS and Holmstrom's (1979) sufficient statistic result. First, equation (4) does not imply that x is a sufficient statistic for e given (x,θ). To see this, suppose the effort space is $[0,1]$ and let x be continuously distributed with a differentiable p.d.f. f_e^θ. Then, the likelihood ratio is $\frac{\frac{\partial f_e^\theta}{\partial e}(x)}{f_e^\theta(x)} = \left[1 - \frac{1}{\frac{\partial H}{\partial x}(x)}\frac{f_0^\theta(x)}{\frac{\partial I}{\partial e}(e,\theta)}\right]^{-1}$, which depends on θ unless $\frac{f_0^\theta(x)}{\frac{\partial I}{\partial e}(e,\theta)}$ is constant in θ. Second, because types also affect the cost of effort, the optimal pure-moral-hazard contract is also a function of θ. For example, in the model of Innes (1990), the strike price of the agent's option is decreasing in the cost of effort (Chaigneau et al. (2014)). Thus, even when θ is uninformative about effort, it still affects the optimal contract directly through the cost of effort.

[15]The relevant assumption is that $H(x)$ has the same sign for all x. We can always renormalize a distribution with $H(x) \leq 0$ for all x as one with $H(x) \geq 0$ by switching the sign of I. In the pure adverse selection model of Nachman and Noe (1994), project types are assumed to be ranked in terms of first-order stochastic dominance, which is fundamentally different than assuming that, for each type θ, effort orders outputs in terms of FOSD. We do not require types to be ordered in any way.

[16]Notice that Example 1 shows that MS is important for Theorem 2, as the distributions in that example are ordered by FOSD. MS can be slightly weakened since it does not need to hold for all efforts, only the ones that principal may want to implement. For example, Theorem 2 remains unchanged if (4) fails at points where $F_\cdot^\theta(x)$ and $c_\cdot^\theta$ are locally concave (since such efforts are not implementable).

an optimal mechanism that offers the same compensation contract to all types.

When $\mathcal{F}^{\theta}$ is ordered by FOSD, the proof of Theorem 2 is similar to the one from Theorem 1. Recall that, with two efforts, substituting all contracts by the one with the highest bonus had two effects. First, holding effort fixed, it reduced the expected payment to all types. Second, it increased the probability of success, which raised the principal's profit because the bonus was lower than the incremental output.

The first effect above remains unchanged with multiple outputs: holding effort fixed, the principal and the agent split a pie of fixed size (since they are both risk neutral). Therefore, for any fixed effort, removing a contract from the mechanism cannot hurt the principal. Moreover, because of MS, contracts can be ranked in terms of their incentives. Thus, replacing all contracts by the one with the highest incentives weakly improves the output distribution of all types (in the sense of FOSD). Since payments cannot grow faster than outputs (BFD), improving the output distribution is always beneficial. Thus, the second effect is also positive when MS holds.

When $\mathcal{F}^{\theta}$ is not ordered, the principal may not want to encourage "higher efforts" from all types, so the second step of the previous argument fails. Instead of offering only the contract with the highest incentives, the principal must also include the contract with the lowest incentives. The proof then verifies that, because of MS, all agents with a positive benefit from effort pick the high-powered contract and the ones with a negative benefit from effort pick the low-powered contract.

Debt Contracts

Next, we consider the optimality of debt contracts. The principal gets a debt contract if his payments $x - w(x)$ equal $\min\{x, \bar{x}\}$ for some face value $\bar{x}$, or, equivalently, if the agent is paid a call option $w(x) = \max\{x - \bar{x},\ 0\}$. To avoid rounding issues, we assume that output is continuously distributed on an interval $X \subset \mathbb{R}_+$. Let the probability density function $f_e^{\theta}(x)$ with full support on X denote type θ's probability of output x conditional on effort e.

Incentive compatible mechanisms cannot have two different debt contracts, since agents would never pick the one with the highest face value. Therefore, for a mechanism to offer debt contracts only, it needs to offer the same contract to all types. Accordingly, we assume that MS holds and $\mathcal{F}^{\theta}$ is ordered by FOSD for all θ, so the principal offers a single contract (Theorem 2).

The existing literature has established that, in the single-task version of the model ($E \subset \mathbb{R}$) with pure moral hazard, the monotone likelihood ratio property (MLRP) is sufficient for the optimality of debt. A distribution f_e^{θ} satisfies MLRP if, for any $e_L, e_H \in \mathbb{R}$ with $e_L < e_H$, the ratio $\frac{f_{e_L}^{\theta}(x)}{f_{e_H}^{\theta}(x)}$ is decreasing in x. Intuitively, MLRP means that the "evidence" in favor of higher efforts increases with output.[17] Here, because all types are offered the same contract, we

[17]MLRP plays an important role on the monotonicity of contracts (Holmstrom (1979), Grossman and Hart

do not need to impose MLRP on *all* types. Let $\bar{f}_e(x) := \int f_e^\theta(x)\,d\mu(\theta)$ denote the marginal distribution of output conditional on e. That is, if all types choose effort $e \in E$, output is distributed according to $\bar{f}_e$. Similarly, let $\bar{I}(e) := \int I(e,\theta)d\mu(\theta)$. The marginal distribution satisfies the monotone likelihood ratio property with respect to $\bar{I}$ if, whenever $\bar{I}(e_L) < \bar{I}(e_H)$, the ratio $\frac{\bar{f}_{e_L}(x)}{\bar{f}_{e_H}(x)}$ is decreasing in x.[18]

Our next result establishes that when, in addition to MS, the family of distributions is ordered by FOSD and the marginal distribution satisfies MLRP, not only is it optimal to offer only one contract, but this contract takes the standard form of debt for the principal or, equivalently, a call option to the agent.

Theorem 3. *Suppose MS holds, the family of distributions is ordered by FOSD, and the marginal distribution satisfies MLRP. Then, the optimal mechanism gives the principal a single debt contract.*

The intuition for the optimality of debt is reminiscent of Innes (1990). With MLRP, higher outputs are more "indicative" of higher effort. Therefore, transferring payments from lower to higher outputs relaxes the IC constraints. In the canonical pure moral hazard model, there is a single IC, preventing the agent from choosing a lower effort.

Here, because of adverse selection, there is one IC for each type, which prevents each type from picking a different effort. However, because of multiplicative separability and ordering by FOSD, the ICs of all types are aligned, in the sense that if a perturbation raises one type's incentives to exert effort, it must also raise the incentives of all types. This observation allows us to summarize all the constraints of the program into a single one, as in the pure moral hazard model. Moreover, since the principal offers a single contract (Theorem 2), the optimality of debt depends on the likelihood ratio of the marginal distribution, and not the conditional distribution of each type. In fact, it is straightforward to follow the approach from the proof of Theorem 3 to characterize optimal non-debt contracts when marginal distributions do not satisfy MLRP.[19]

(1983)).

[18]Under MS, any family of distributions that satisfies MLRP for each type θ is both ordered by FOSD and has a marginal distribution $\bar{F}_e$ that satisfies MLRP. Therefore, MLRP for each type is a stronger requirement than being ordered by FOSD and having a marginal that satisfies MLRP. Notice that the linear distribution (5) satisfies MS, is ordered by FOSD, and has a marginal that satisfies MLRP as long as $F_1^\theta(x) \le F_0^\theta(x)$ for all x, θ (i.e., the highest effort first-order stochastically dominates the lowest effort).

[19]In the appendix, we present a partial converse of Theorem 3 showing that, if the marginal distribution does not satisfy MLRP and satisfies a weak condition, there exists a cost function for which debt is not optimal. This weak condition rules out distributions that are so extreme that the principal would implement the same effort for any cost function (in which case the optimal contract would pay zero in all states).

4 Conclusion

The observation that many contracts are simple and relatively uniform across different sectors is an old puzzle in contract theory. While standard adverse selection models predict that agents will be offered large menus of contracts, contracting parties typically offer a limited number of contracts, often a single one. While standard moral hazard models predict that contracts should be fine tuned to the likelihood ratio of output, similar contracts are offered in different environments.

We argue that these two features endogenously emerge in a general model of moral hazard and adverse selection if contracts must satisfy two common contractual constraints: limited liability and free disposal. With binary outcomes, the principal always offers a single debt-like contract regardless of any parameters of the model. The joint presence of moral hazard and adverse selection is key for this result. When either types or effort are observed, the principal typically prefers to offer different contracts to different types. With multiple outputs, it is optimal to offer a single contract if either arbitrage opportunities restrict the space of contracts to two-part tariffs, or if the distribution of output is ordered by FOSD and satisfies a separability condition. Moreover, if the marginal distribution satisfies MLRP, this single optimal contract is a debt contract for the principal.

Our paper shows that gaming may be an important downside from giving flexibility to agents by offering menus of contracts. This is particularly stark with bilateral risk neutrality where, holding effort fixed, the agent always selects the most expensive contract to the principal. Then, reducing the number of contracts offered to the agent always increases the principal's profits (for a fixed effort). If, in addition, we can identify a most efficient contract from the menu (such as when output is binary or when the distribution is multiplicatively separable), the principal can always improve by eliminating other contracts.

The simplicity result relies on the presence of limited liability and bilateral risk neutrality. Limited liability ensures that increasing the power of a contract will not force the agent to abandon the mechanism. Bilateral risk neutrality means that, holding effort fixed, principal and agent are perfectly misaligned so that the principal always benefits by reducing the agent's flexibility. With risk aversion, their misalignment is no longer perfect because there are potential gains from risk-sharing. While risk neutrality is a reasonable assumptions in many settings (such as the optimal compensation of wealthy managers, procurement contracts, or the regulation of large companies), there are many other settings where they are not (for example, insurance contracts or sharecropping). In these cases, our results no longer hold. In our companion paper (Gottlieb and Moreira 2014), we study optimal mechanisms in the binary-outcome model when agents are risk averse and when there are no limited liability constraints. While we obtain some simplicity results, optimal mechanisms are considerably more complex than they are here.

Appendix

A. Procurement and Regulation

In this appendix, we adapt our framework to a model of procurement and regulation that builds on the setup of Laffont and Tirole (1986, 1993). The key distinguishing feature is that, in their model, effort affects the regulated firm's cost deterministically so the model can be reduced to a pure adverse selection model. For this reason, it is often called a model of 'false moral hazard.' Our model incorporates a dimension of moral hazard in their model by allowing effort to affect the firm's cost stochastically. We also allow private information to be multidimensional, as opposed to the standard single-dimensional model.

A regulated firm produces an indivisible project at a random monetary cost $C \in \mathcal{C}$. The firm's manager chooses a cost-reducing effort $e \in E \subset \mathbb{R}$, which is not observed by the regulator. Let F_e^θ denote the distribution of the firm's cost C conditional on type θ and effort e. By exerting effort, the manager improves the distribution of the firm's cost in the sense of FOSD:

$$e_H > e_L \implies F_{e_H}^\theta(x) \leq F_{e_L}^\theta(x) \quad \forall x. \tag{FOSD}$$

The firm's manager has cost of effort c_e^θ, with $\min_e c_e^\theta \leq 0$. As argued in the text, this is satisfied if the lowest effort costs zero or if the manager gets private benefits out of some activities (so that $c_e^\theta < 0$ for some e). The firm's manager has private information about both the ability to cut costs (i.e., the distribution of costs F_e^θ) and the cost of effort c_e^θ.

The project generates a consumer surplus of $S > 0$. The regulator observes the monetary cost C incurred by the firm but not the manager's private effort e. As an accounting convention, we assume that the regulator reimburses the firm's monetary costs in addition to paying the firm an amount conditional on the realized cost.

A *procurement contract* is a function that specifies a transfer to the firm conditional on each possible cost C. A mechanism is a pair of functions $w : \mathcal{C} \times \Theta \to \mathbb{R}$ and $e : \Theta \to \mathbb{R}$ specifying, for each reported type, a payment conditional on the cost realization and an effort recommendation. If this mechanism is truthful, a type-θ firm gets payoff

$$U(\theta) \equiv \int w_\theta(C)\, dF_{e(\theta)}^\theta(C) - c_{e(\theta)}^\theta. \tag{6}$$

The manager is protected by limited liability (LL) so that payments are non-negative:

$$w_\theta(C) \geq 0. \tag{LL}$$

There is also bilateral free disposal (BFD), which requires the compensation for reducing costs to be positive and not to exceed the amount of cost savings:

$$0 \leq w_\theta(C) - w_\theta(C+\epsilon) \leq \epsilon \tag{BFD}$$

for all $\epsilon > 0$. BFD must be satisfied, for example, if the manager can freely inflate costs (so that payments are non-decreasing) and if the manager can freely borrow from an outside party to inflate firm earnings (so that payments cannot grow faster than the amount of cost savings).

Since, by the accounting convention described above, the regulator fully reimburses the firm's cost realization, the regulator's expected payment conditional on type θ effort e equals

$$\int [C + w_\theta(C)]\, dF^\theta_{e(\theta)}(C).$$

As in Laffont and Tirole (1986, 1993), we assume that the government has to revert to distortionary taxation to raise funds and, therefore, the regulator faces a shadow cost of public funds $\lambda > 0$. Thus, the net surplus of consumers/taxpayers is

$$S - (1+\lambda)\int [C + w_\theta(C)]\, dF^\theta_{e(\theta)}(C). \tag{7}$$

A utilitarian regulator maximizes the sum of the expected utility of the firm's manager (6) and the consumers' net surplus (7):

$$S - (1+\lambda)\int C dF^\theta_{e(\theta)}(C) - c^\theta_{e(\theta)} - \lambda \int w_\theta(C)\, dF^\theta_{e(\theta)}(C).$$

Notice that, because taxation is distortionary ($\lambda > 0$), leaving rents the the regulated firm is costly. Moreover, because each dollar reimbursed to the firm has an additional cost of λ due to distortionary taxation, cutting the firm's cost increases social surplus by $1+\lambda$. The first-best effort, therefore, minimizes $(1+\lambda)\int C dF^\theta_e(C) + c^\theta_e$.[20]

As in Section 3, we impose the technical conditions from Assumption 2 (with C instead of x). Recall that MS is automatically satisfied if there are only two possible costs.

The main difference between this model and the one from Section 3 that, while the principal only cares about her own payoff, a utilitarian regulator also cares about the regulated firm's payoffs. Therefore, their preferences are no longer perfectly misaligned. Because of distortionary

[20] With pure moral hazard and without LL and BFD, the first best can be implemented by making the firm the residual claimant of the social gains from cutting costs and extracting the firm's entire surplus:

$$w_\theta(C) = (1+\lambda)\left[\int C dF^\theta_{e^{FB}_\theta(\theta)}(C) - C\right],$$

where $e^{FB}_\theta(\theta)$ is the first-best effort. This violates both LL and BFD. It violates LL because w is non-degenerate and has mean zero (so there must make negative payments), and it violates BFD because $\frac{\partial w_\theta}{\partial C} = -(1+\lambda) < -1$.

taxation, the regulator would still like to leave as little rents as possible to the firm. However, because the regulator also internalizes the manager's effort cost, reducing effort distortions may not increase the regulator's payoffs. We therefore need conditions to ensure that effect distortions have a monotonic effect on social surplus. For this reason, we assume that social surplus net of firm payments has a unique maximum and that it is increasing in the range of efforts below the maximum. Formally, let $\bar{e}_\theta$ be the effort that maximizes

$$S - (1+\lambda)\int CdF_e^\theta(C) - c_e^\theta, \tag{8}$$

which we assume to be unique. We assume that

$$\frac{d}{de}\left\{(1+\lambda)\int CdF_e^\theta(C) + c_e^\theta\right\} \leq 0$$

for all $e \leq \bar{e}_\theta$. These assumptions state that reducing effort below the efficient level lowers the net surplus. It is automatically satisfied if the net social surplus (8) is a quasi-concave function of effort e. In particular, it holds if $F_\cdot^\theta$ and $c_\cdot^e$ are convex functions.

We can now state the simplicity result:

Proposition 1. *Suppose MS holds. There exists an optimal mechanism that offers a single procurement contract to all types.*

Next, we turn to the optimal contractual form. Adapting Theorem 3, we obtain the following proposition:

Proposition 2. *Suppose MS holds and the marginal distribution satisfies MLRP. There exists an optimal mechanism that offers to all types the procurement contract* $w(C) = \max\left\{\bar{C} - C;\ \ 0\right\}$, *for some* $\bar{C} \in \mathbb{R}$.

The optimal procurement contract specifies a "reasonable cost" $\bar{C}$ and reimburses the regulated firm for any cost cuts beyond $\bar{C}$. Since, by our accounting convention, the regulator pays the firm's cost directly, the firm's revenues under this contract equal:

$$w(C) + C = \max\left\{\bar{C}, C\right\}.$$

This reimbursement rule consists of a standard price cap except that, because of limited liability, the regulator must bailout firms with cost realizations above the cap. As previously mentioned, price caps are the most common form of incentive regulation. It is used, for example, by the U.S. Federal Communications Commission (FCC) to regulate the telephone industry. Price caps are often used in procurement as well. For example, prospective reimbursement systems commonly used in health care specify an amount $\bar{C}$ based on what a service should cost, and let providers keep cost savings $\bar{C} - C$ to themselves. They are used, for example, by Medicare.

B. Proofs

Proof of Theorem 1

We first verify that participation is implied by incentive compatibility, limited liability, and free disposal:

Lemma 1. *Let (w, b, e) be a mechanism that satisfies IC, LL, and FD. Then, it satisfies IR.*

Proof. The IC preventing θ from deviating to e states that $U(\theta) \geq w(\theta) + p_e^\theta b(\theta)$. Then, because $w(\theta)$ and $b(\theta)$ are non-negative (LL and FD) and $\min_{e \in E} c_e^\theta \leq 0$, it follows that $U(\theta) \geq 0$. □

The proof of the theorem will use the following result, which states that the bonus does not exceed the incremental output:

Lemma 2. *Consider a mechanism (w, b, e) in which $b(\theta) > \Delta x$ for some θ. Then, the mechanism is not optimal.*

Proof. Notice that, in any optimal mechanism, the limited liability constraint must bind for some type. Otherwise, reducing the fixed payment to all types by a uniform amount maintains feasibility and increases the principal's payoff.

Let (w, b) be an optimal mechanism such that $b(\theta) > \Delta x$ for some type θ. If all types obtain $b(\theta) > \Delta x$, the principal's payoff is strictly lower than the one she obtains by offering $w = b = 0$ to all types, which contradicts optimality.

Suppose there exists a type $\tilde{\theta}$ who picks a contract with $b(\tilde{\theta}) \leq \Delta x$. Because LL must bind for some type, it must bind for a type that picks $b_H > \Delta x$ (otherwise, the contract $(0, b_H)$ with $b_H \leq \Delta x$ would pay less than the contract that pays $w \geq 0$ and $b > \Delta x$). Therefore, the mechanism also includes at least one contract (w_L, b_L) with $b_L \leq \Delta x$ and contract $(0, b_H)$ with $b_H > \Delta x$. We claim that this mechanism gives the principal a lower profit on all types than offering $w = b = 0$ to all types.

For each θ, let $\bar{e}(\theta) \in \arg\min_{e \in E} c_e^\theta$. The principal's profit from a type $\hat{\theta}$ choosing contract $(0, b_H)$ and exerting effort $e_H \in E$ is

$$x_L + p_{e_H}^{\hat{\theta}} \left[\Delta x - b_H\right].$$

Replacing this contract by $w = b = 0$ changes the principal's payoff to $x_L + p_{\bar{e}(\hat{\theta})}^{\hat{\theta}} \Delta x$, which gives a profit of

$$p_0^{\hat{\theta}} \Delta x - p_{e_H}^{\hat{\theta}} \left[b_H - \Delta x\right] > 0.$$

The principal's profit from a type θ choosing (w_L, b_L) and exerting effort $e_L \in E$ is

$$x_L - w_L + p_{e_L}^\theta \left[\Delta x - b_L\right].$$

By the incentive compatibility constraint of a type who picks (w_L, b_L) and the fact that $b_H > \Delta x$,

$$p^\theta_{e_L} \Delta x < p^\theta_{e_L} b_H \le w_L + p^\theta_{e_L} b_L.$$

Adding x_L to both extremes of this inequality and rearranging, gives

$$x_L + p^\theta_{e_L} \Delta x - \left(w_L + p^\theta_{e_L} b_L\right) < x_L.$$

Add $p^\theta_{\bar{e}(\theta)} \Delta x \ge 0$ to the expression on the right to obtain:

$$x_L + p^\theta_{e_L} \Delta x - \left(w_L + p^\theta_{e_L} b_L\right) < x_L + p^\theta_{\bar{e}(\theta)} \Delta x.$$

The term on the right is the principal's profit from the constant-wage contract ($w = b = 0$) whereas the term on the left is the profit from the original contract. Hence, this replacement also raises profits from any type who chooses a contract with $b(\theta) \le \Delta x$. $\square$

The next lemma, which is the main step for proving Theorem 1, shows that any feasible mechanism is weakly dominated by a mechanism that offers a single contract to all types:

Lemma 3. *Let (w, b, e) be a mechanism satisfying IC, IR, FD, and LL. There exists a mechanism that offers a single contract $(0, b^*)$ to all types and gives the principal a (weakly) greater payoff than (w, b, e).*

Proof. Let (w, b, e) be a mechanism that satisfies IC, IR, FD, and LL. By Lemma 2, for any mechanism that offers a bonus greater than Δx for some type, there exists another mechanism offering bonuses lower than Δx to all types that gives the principal a higher payoff. Thus, there is no loss of generality in assuming that $b(\hat{\theta}) \le \Delta x$ for all $\hat{\theta} \in \Theta$. Fix a type $\theta \in \Theta$.

Let $b^* \equiv \sup \left\{b(\hat{\theta}) : \hat{\theta} \in \Theta\right\}$ and $w^* \equiv \inf \left\{w(\hat{\theta}) : \hat{\theta} \in \Theta\right\}$ denote the "highest" bonus and the "lowest" fixed payment in the mechanism. If $w^* > 0$, reducing all fixed payments uniformly by w^* would keep all the constraints satisfied and increase the principal's payoff. Therefore, we can assume that $w^* = 0$. Moreover, either there exists $\hat{\theta}$ such that $w\left(\hat{\theta}\right) = 0$, $b\left(\hat{\theta}\right) = b^*$ (i.e., $(0, b^*)$ is offered in the mechanism), or $(0, b^*)$ is a limit point of $\left\{\left(w\left(\hat{\theta}\right), b\left(\hat{\theta}\right)\right) ; \hat{\theta} \in \Theta\right\}$.

Consider the alternative mechanism that offers the contract $(0, b^*)$ to all types and let

$$e^*(\theta) \in \arg\max_{e \in E} p^\theta_e b^* - c^\theta_e,$$

which exists because the objective function is continuous and E is compact. The principal's payoff from type θ in the original mechanism is

$$x_L - w(\theta) + p^\theta_{e(\theta)} \left[\Delta x - b(\theta)\right]. \tag{9}$$

Her payoff in the alternative mechanism is

$$x_L + p^{\theta}_{e^*(\theta)} \left[\Delta x - b^*\right]. \tag{10}$$

Since the contract $(0, b^*)$ either belongs to, or is a limit point of the original mechanism, no agent can be better off by switching to $(0, b^*)$ while holding the recommended effort fixed:

$$w(\theta) + p^{\theta}_{e(\theta)} b(\theta) - c^{\theta}_{e(\theta)} \geq p^{\theta}_{e(\theta)} b^* - c^{\theta}_{e(\theta)}.$$

(This inequality follows from type θ's IC constraint while holding $e(\theta)$ fixed). Summing $c^{\theta}_{e(\theta)}$ to both sides, it follows that the expected payment in the alternative mechanism cannot exceed the one from the original mechanism if effort is held constant:

$$w(\theta) + p^{\theta}_{e(\theta)} b(\theta) \geq p^{\theta}_{e(\theta)} b^*. \tag{11}$$

That is, if the agent chooses not to change effort ($e^*(\theta) = e(\theta)$), the principal obtains a higher payoff in the alternative mechanism (10) than in the original one (9). Allowing effort to change, the principal's payoff in the alternative mechanism minus the payoff in the original mechanism equals

$$\begin{aligned} p^{\theta}_{e^*(\theta)} \left[\Delta x - b^*\right] - p^{\theta}_{e(\theta)} \left[\Delta x - b(\theta)\right] + w(\theta) &= \left(p^{\theta}_{e^*(\theta)} - p^{\theta}_{e(\theta)}\right) \left[\Delta x - b^*\right] + w(\theta) + p^{\theta}_{e(\theta)} b\left(\theta\right) - p^{\theta}_{e(\theta)} b^* \\ &\geq \left(p^{\theta}_{e^*(\theta)} - p^{\theta}_{e(\theta)}\right) \left[\Delta x - b^*\right] \end{aligned}, \tag{12}$$

where the inequality follows from (11).

By assumption, $\Delta x \geq b^*$. We claim that $p^{\theta}_{e^*(\theta)} \geq p^{\theta}_{e(\theta)}$. To wit, by the IC of type θ in the original mechanism,

$$w(\theta) + p^{\theta}_{e(\theta)} b(\theta) - c^{\theta}_{e(\theta)} \geq w(\theta) + p^{\theta}_{e^*(\theta)} b(\theta) - c^{\theta}_{e^*(\theta)},$$

and, by the definition of $e^*(\theta)$,

$$p^{\theta}_{e^*(\theta)} b^* - c^{\theta}_{e^*(\theta)} \geq p^{\theta}_{e(\theta)} b^* - c^{\theta}_{e(\theta)}.$$

Rearranging both inequalities, we can write them as:

$$\left(p^{\theta}_{e^*(\theta)} - p^{\theta}_{e(\theta)}\right) b^* \geq c^{\theta}_{e^*(\theta)} - c^{\theta}_{e(\theta)} \geq \left(p^{\theta}_{e^*(\theta)} - p^{\theta}_{e(\theta)}\right) b(\theta)$$

$$\therefore \left(p^{\theta}_{e^*(\theta)} - p^{\theta}_{e(\theta)}\right) \left[b^* - b(\theta)\right] \geq 0.$$

Since b^* is the supremum of bonuses, $b^* \geq b(\theta)$. Therefore, $p^{\theta}_{e^*(\theta)} \geq p^{\theta}_{e(\theta)}$. Hence, (12) implies that the principal's payoff from each type θ is higher in the alternative mechanism even when the agent chooses a different effort. $\square$

The proof concludes by verifying that an optimal mechanism exists. The proof of existence follows arguments similar to Page (1992) and is given in the supplementary appendix.

Proofs of Theorems 2 and 3

In the proofs below, we will use the following facts:

1. Any contract that satisfies BFD is a Lipschitz function.
2. Lipschitz functions are absolutely continuous (and, hence, almost everywhere differentiable) and we can apply the Fundamental Theorem of Calculus to them.
3. Integration by parts is a valid procedure between integrable and absolutely continuous functions.
4. If a contract satisfies IC and LL, it also satisfies IR (because $\min_e c_e^\theta \leq 0$).

Proof of Theorem 2

We first introduce some notation. Let $\underline{x} := \inf X$ (possibly $-\infty$). Since any contract $w : X \to \mathbb{R}_+$ that satisfies BFD is absolutely continuous, we can apply the fundamental theorem of calculus and write $w(x) = w(\underline{x}) + \int_{\underline{x}}^{x} \dot{w}(y)dy$. Hence, any such contract w is characterized by a "fixed wage" $w(\underline{x}) \in \mathbb{R}_+$ and "bonus" $\dot{w} \in L_+^\infty(X)$.[21] Using integration by parts, the payoff of a type-θ agent who exerts effort e and gets contract $(w(\underline{x}), \dot{w})$ equals

$$v_e^\theta(w(\underline{x}), \dot{w}) := w(\underline{x}) + \int \dot{w}(x)\left(1 - F_e^\theta(x)\right) dx - c_e^\theta. \tag{13}$$

Similarly, the principal's payoff is

$$u_e^\theta(w(\underline{x}), \dot{w}) := -w(\underline{x}) + \int [1 - \dot{w}(x)]\left(1 - F_e^\theta(x)\right) dx. \tag{14}$$

By MS, a type-θ agent who switches from effort $\tilde{e}$ to e while keeping the same contract $(w(\underline{x}), \dot{w})$ gains

$$v_e^\theta(w(\underline{x}), \dot{w}) - v_{\tilde{e}}^\theta(w(\underline{x}), \dot{w}) = [I(e, \theta) - I(\tilde{e}, \theta)] \int \dot{w}(x)H(x)dx + c_{\tilde{e}}^\theta - c_e^\theta. \tag{15}$$

[21] $L^\infty(X)$ is the space of all real-valued, measurable, and essentially bounded functions with domain X. $L_+^\infty(X)$ is the subset of non-negative functions in $L^\infty(X)$.

In turn, this switch affects the principal's payoff by

$$u_e^\theta(w(\underline{x}), \dot{w}) - u_{\tilde{e}}^\theta(w(\underline{x}), \dot{w}) = [I(e,\theta) - I(\tilde{e},\theta)] \int [1 - \dot{w}(x)]H(x)dx. \tag{16}$$

Notice that the principal gains from shifting effort towards an effort associated with a higher $I(e,\theta)$ if and only if

$$\int [1 - \dot{w}(x)]H(x)dx \geq 0. \tag{17}$$

In this proof, we will verify that the principal benefits by substituting all contracts associated with $\int[1 - \dot{w}(x)]H(x)dx > 0$ by the contract that encourages the highest $I(e,\theta)$ and by substituting all contracts associated with $\int[1 - \dot{w}(x)]H(x)dx < 0$ by the one that encourages the lowest $I(e,\theta)$. To establish this result, not only do we need to verify that the principal benefits by this substitution, but also that each type picks the contract designed for him.

Recall that a mechanism offers a contract $w_\theta(x)$ and an effort recommendation e_θ to each type $\theta \in \Theta$. We will write $\dot{w}_\theta(x) := \frac{d}{dx} w_\theta(x)$. Let (w, e) be a feasible mechanism. By BFD, the set of all bonuses in this mechanism,

$$\mathcal{M} := \{\dot{w}_\theta(x) : \ \theta \in \Theta\},$$

is well defined and is composed of uniformly bounded functions (with a lower bound of 0 and an upper bound of 1). The Banach-Alaoglu Theorem (Rudin (1991), p. 68) implies that its closure, $\bar{\mathcal{M}}$, is weak* compact in $L^\infty(X)$.

For each type $\theta \in \Theta$, there are two possibilities:

$$\int [1 - \dot{w}_\theta(x)]H(x)dx \geq 0 \text{ or } \int [1 - \dot{w}_\theta(x)]H(x)dx \leq 0.$$

In the first one, the principal would like to encourage an effort associated with a higher I; the reverse is true in the second case.

Case 1) $\int[1-\dot{w}_\theta(x)]H(x)dx \geq 0$. Let $\dot{w}^+$ be a solution of the following maximization program:

$$\begin{aligned} &\max_{\dot{w}\in\bar{\mathcal{M}}} \int \dot{w}(x)H(x)dx, \\ &\text{subject to } \int[1 - \dot{w}(x)]H(x)dx \geq 0. \end{aligned} \tag{18}$$

Because the constraint of problem (18) is satisfied for w_θ, the set of bonuses in $\bar{\mathcal{M}}$ that satisfy this constraint is non-empty. Since $\bar{\mathcal{M}}$ is weak* compact, the objective function is a continuous linear functional on $L^\infty(X)$ and the constraint defines a weak* closed set. Therefore, (18) has a solution $\dot{w}^+ \in \bar{\mathcal{M}}$. In order to complete the definition of contract w^+, we need to specify $w^+(\underline{x})$.

Let $(\dot{w}_{\theta_n})$ be a sequence in $\mathcal{M}$ that satisfies the constraint of (18) and weak*-converges to $\dot{w}^+$. We claim that $(w_{\theta_n}(\underline{x}))$ is a bounded sequence in $\mathbb{R}_+$. Indeed, fix a type $\bar{\theta} \in \Theta$. By $\bar{\theta}$'s IC,

$$v^{\bar{\theta}}_{e(\bar{\theta})}(w_{\bar{\theta}}(\underline{x}), \dot{w}_{\bar{\theta}}) \geq v^{\bar{\theta}}_{e(\bar{\theta})}(w_{\theta_n}(\underline{x}), \dot{w}_{\theta_n})$$

which, by LL and assumption A2 (iv), gives

$$0 \leq w_{\theta_n}(\underline{x}) \leq w_{\bar{\theta}}(\underline{x}) + \xi(\theta),$$

since $\int \dot{w}_{\bar{\theta}}(x) \left(1 - F^{\bar{\theta}}_{e(\bar{\theta})}(x)\right) dx \leq \int \left(1 - F^{\bar{\theta}}_{e(\bar{\theta})}(x)\right) dx = \int x dF^{\bar{\theta}}_{e(\bar{\theta})}(x)$. Take a convergent subsequence of $(w_{\theta_n}(\underline{x}))$ and let $w^+(\underline{x})$ denote its limit. Notice that w^+ satisfies LL and BFD.

We now verify that replacing w_θ by w^+ increases the principal's payoff. Let $e^+(\theta)$ be an effort that maximizes the agent's payoff under contract w^+ (for existence, see the online appendix):

$$v^{\theta}_{e^+(\theta)}(w^+(\underline{x}), \dot{w}^+) \geq v^{\theta}_{e(\theta)}(w^+(\underline{x}), \dot{w}^+),$$

which, by MS, can be written as

$$\left[I(e^+(\theta), \theta) - I(e(\theta), \theta)\right] \int \dot{w}^+(x) H(x)\, dx \geq c^{\theta}_{e^+(\theta)} - c^{\theta}_{e(\theta)}. \tag{19}$$

Similarly, because $e(\theta)$ is his effort choice with contract $w_\theta(x)$,

$$\left[I(e^+(\theta), \theta) - I(e(\theta), \theta)\right] \int \dot{w}_\theta(x) H(x)\, dx \leq c^{\theta}_{e^+(\theta)} - c^{\theta}_{e(\theta)}. \tag{20}$$

Combining (19) and (20), we obtain

$$\begin{aligned} \left[I(e^+(\theta), \theta) - I(e(\theta), \theta)\right] \textstyle\int \dot{w}_\theta(x) H(x)\, dx &\leq c^{\theta}_{e^+(\theta)} - c^{\theta}_{e(\theta)} \\ &\leq \left[I(e^+(\theta), \theta) - I(e(\theta), \theta)\right] \textstyle\int \dot{w}^+(x) H(x)\, dx. \end{aligned} \tag{21}$$

Since $\dot{w}^+$ solves program (18), we have

$$\int \dot{w}^+(x) H(x) dx \geq \int \dot{w}_\theta(x) H(x) dx.$$

Therefore, it follows from (21) that $I(e^+(\theta), \theta) \geq I(e(\theta), \theta)$.

We now establish that replacing contract w_θ by w^+ increases the principal's payoff from type θ. As in the proof of Theorem 1, we first show that, holding effort fixed, the principal is better off with the substitution of contracts. Since w^+ is the weak*-limit of sequence in $\mathcal{C}$, the agent's utility is continuous in the weak* topology, and the original mechanism is incentive compatible, it follows that

$$v^{\theta}_{e(\theta)}(w_\theta(\underline{x}), \dot{w}_\theta) \geq v^{\theta}_{e(\theta)}(w^+(\underline{x}), \dot{w}^+). \tag{22}$$

Substitute the expression for the agent's payoff and multiply both sides by -1:

$$-w^+(\underline{x}) - \int \dot{w}^+(x)\left(1 - F^{\theta}_{e(\theta)}(x)\right) dx \geq -w_\theta(\underline{x}) - \int \dot{w}_\theta(x)\left(1 - F^{\theta}_{e(\theta)}(x)\right) dx.$$

Add $\int \left(1 - F^{\theta}_{e(\theta)}(x)\right) dx$ to both sides of this inequality:

$$-w^+(\underline{x}) + \int \left[1 - \dot{w}^+(x)\right]\left(1 - F^{\theta}_{e(\theta)}(x)\right) dx \geq -w_\theta(\underline{x}) + \int \left[1 - \dot{w}_\theta(x)\right]\left(1 - F^{\theta}_{e(\theta)}(x)\right) dx, \tag{23}$$

which states that, holding effort $e(\theta)$ fixed, the principal gets a higher profit with contract $(w^+(\underline{x}), \dot{w}^+)$ than with $(w_\theta(\underline{x}), \dot{w}_\theta)$.

Next, we show that the change in effort also benefits the principal. Because $I(e^+(\theta), \theta) > I(e(\theta), \theta)$, and because $\dot{w}^+$ solves program (18), the following inequality holds:

$$\left[I\left(e^+(\theta), \theta\right) - I\left(e(\theta), \theta\right)\right] \int \left[1 - \dot{w}^+(x)\right] H(x)\, dx \geq 0.$$

Using MS, rewrite this inequality as

$$\int \left[1 - \dot{w}^+(x)\right]\left(1 - F^{\theta}_{e^+(\theta)}(x)\right) dx \geq \int \left[1 - \dot{w}^+(x)\right]\left(1 - F^{\theta}_{e(\theta)}(x)\right) dx, \tag{24}$$

which shows that the principal gains fro the change in effort.

Combining (23) and (24), establishes that the principal's profit from θ with the new contract exceeds her profit with the original contract:

$$\underline{x} - w^+(\underline{x}) + \int \left[1 - \dot{w}^+(x)\right]\left(1 - F^{\theta}_{e^+(\theta)}(x)\right) dx \geq \underline{x} - w_\theta(\underline{x}) + \int \left[1 - \dot{w}^+(x)\right]\left(1 - F^{\theta}_{e(\theta)}(x)\right) dx.$$

Case 2) $\int [1 - \dot{w}_\theta(x)] H(x) dx \leq 0$. Let w^- be the solution of the minimization program:

$$\begin{aligned} &\min_{\dot{w} \in \bar{\mathcal{M}}} \int \dot{w}(x) H(x) dx, \\ &\text{subject to } \int [1 - \dot{w}(x)] H(x) dx \geq 0. \end{aligned} \tag{25}$$

As in problem (18), let $\dot{w}^- \in \bar{\mathcal{M}}$ be a solution of problem (25) and define $w^-(0)$ following the same procedure as used to obtain $w^+(0)$. Then, w^- satisfies LL and BFD.

By the same arguments as in Case 1, incentive compatibility implies that

$$
\begin{aligned}
\left[I(e^{-}(\theta),\theta)-I(e(\theta),\theta)\right]\textstyle\int \dot{w}(\theta,x)H(x)\,dx \quad &\leq c^{\theta}_{e^{-}(\theta)}-c^{\theta}_{e(\theta)} \\
&\leq \left[I(e^{-}(\theta),\theta)-I(e(\theta),\theta)\right]\textstyle\int \dot{w}^{-}(x)H(x)\,dx
\end{aligned}
\tag{26}
$$

Moreover, since $\dot{w}^{-}$ solves (25),

$$
\int \dot{w}^{-}(x)H(x)dx \leq \int \dot{w}_{\theta}(x)H(x)dx,
$$

so that, by (26), $I(e^{-}(\theta),\theta) \leq I(e(\theta),\theta)$. Thus, by (16), the principal (weakly) gains from replacing w_{θ} by w^{-}.

Consider the mechanism $(\bar{w},\bar{e})$:

$$
\bar{w}_{\theta}(x)=\begin{cases} w^{+}(x), \text{ if } \int[1-\dot{w}_{\theta}(x)]H(x)dx \geq 0 \\ w^{-}(x), \text{ if } \int[1-\dot{w}_{\theta}(x)]H(x)dx < 0 \end{cases}
$$

and

$$
\bar{e}(\theta)=\begin{cases} e^{+}(\theta), \text{ if } \int[1-\dot{w}_{\theta}(x)]H(x)dx \geq 0 \\ e^{-}(\theta), \text{ if } \int[1-\dot{w}_{\theta}(x)]H(x)dx < 0 \end{cases},
$$

where, as before, $e^{\pm}(\theta) \in \arg\max_{e} v^{\theta}_{e}(w^{\pm}(\underline{x}),\dot{w}^{\pm})$. As previously shown, this mechanism raises the principal's payoff pointwise (i.e., for it raises the payoff conditional on each type) and satisfies LL, and BFD. Moreover, if it satisfies incentive compatibility, it must also satisfy IR, since the expected utility of type θ at contract $w^{\pm}$ who chooses effort e is

$$
v^{\theta}_{e}(w^{\pm}(\underline{x}),\dot{w}^{\pm})=w^{\pm}(\underline{x})+\int \dot{w}(x)(1-F^{\theta}_{e}(x))dx-c^{\theta}_{e}.
$$

Since $w^{\pm}(0)\geq 0$, $\dot{w}^{\pm}\geq 0$, and the agent must weakly prefer the recommended effort to the one with $c^{\theta}_{e}\leq 0$, IR must hold.

It only remains to show that this mechanism satisfies IC. Since $e^{+}(\theta)$ and $e^{-}(\theta)$ maximize type θ's payoff conditional on each contract, we only need to verify that types with recommended contract w^{+} do not benefit from deviating to w^{-}:

$$
\int[1-\dot{w}_{\theta}(x)]H(x)dx>0 \implies v^{\theta}_{e^{+}(\theta)}(w^{+}(\underline{x}),\dot{w}^{+})\geq v^{\theta}_{e^{-}(\theta)}(w^{-}(\underline{x}),\dot{w}^{-}), \tag{27}
$$

and types with recommended contract w^{-} do not wish to deviate to w^{+}:

$$
\int[1-\dot{w}_{\theta}(x)]H(x)dx<0 \implies v^{\theta}_{e^{-}(\theta)}(w^{-}(\underline{x}),\dot{w}^{-})\geq v^{\theta}_{e^{+}(\theta)}(w^{+}(\underline{x}),\dot{w}^{+}). \tag{28}
$$

In order to verify condition (27), let θ and $\hat{\theta}$ be types with

$$\int [1 - \dot{w}_\theta(x)] H(x) dx \geq 0 \geq \int [1 - \dot{w}_{\hat{\theta}}(x)] H(x) dx.$$

Since type θ chose effort $e(\theta)$ in the original mechanism, incentive compatibility of the original mechanism gives

$$v^\theta_{e(\theta)}(w_\theta(\underline{x}), \dot{w}_\theta) \geq v^\theta_{e^-(\theta)}(w_{\hat{\theta}}(\underline{x}), \dot{w}_{\hat{\theta}}).$$

Let $(\hat{\theta}_n)$ be a sequence such that $\int [1 - \dot{w}_{\hat{\theta}_n}(x)] H(x) dx \leq 0$ for all n, $(\dot{w}_{\hat{\theta}_n})$ weak*-converges to $w^-(x)$ and $(w_{\hat{\theta}_n}(\underline{x}))$ converges to $w^-(\underline{x})$. Again, incentive compatibility of the original mechanism gives

$$v^\theta_{e(\theta)}(w_\theta(\underline{x}), \dot{w}_\theta) \geq v^\theta_{e^-(\theta)}(w^-(\underline{x}), \dot{w}^-), \tag{29}$$

where we are using the continuity of $v^\theta_{e^-(\theta)}(\cdot, \cdot)$.

Now let (θ_n) be a sequence such that $\int [1 - \dot{w}_{\theta_n}(x)] H(x) dx \geq 0$ for all n, $(\dot{w}_{\hat{\theta}_n})$ weak* converges to $w^+(x)$, and $(w_{\theta_n}(\underline{x}))$ converges to $w^+(\underline{x})$. Since E is a compact metric space, $(e(\theta_n))$ has a converging subsequence. Let $\tilde{e} \in E$ denote its limit and, with some abuse of notation, let $(e(\theta_n))$ denote the subsequence itself. We claim that

$$\lim_{n \to \infty} v^\theta_{e(\theta_n)}(w_{\theta_n}(\underline{x}), \dot{w}_{\theta_n}) = v^\theta_{\tilde{e}}(w^+(\underline{x}), \dot{w}^+).$$

Indeed, by the continuity of $c^\theta_\cdot$ and the fact that $\lim_{n \to \infty} w_{\theta_n}(\underline{x}) = w^+(\underline{x})$, we only need to show the convergence of the integral term in (13). But notice that

$$\begin{aligned} \int \dot{w}_{\theta_n}(x) \left(1 - F^\theta_{e(\theta_n)}(x)\right) dx - \int \dot{w}^+(x) \left(1 - F^\theta_{\tilde{e}}(x)\right) dx &= \int \dot{w}_{\theta_n}(x) \left(F^\theta_{\tilde{e}}(x) - F^\theta_{e(\theta_n)}(x)\right) dx \\ &\quad + \int \left[\dot{w}_{\theta_n}(x) - \dot{w}^+(x)\right] F^\theta_{\tilde{e}}(x) dx. \end{aligned}$$

By MS, the first term on the right hand side equals

$$[I(e(\theta_n), \theta) - I(\tilde{e}, \theta)] \int \dot{w}_{\theta_n}(x) H(x) dx,$$

which converges to zero since $I(\cdot, \theta)$ is a continuous function. The second term on the right hand side also converges to zero because $(\dot{w}_{\theta_n})$ weak* converges to $\dot{w}^+$.

Since $e^+(\theta)$ is an optimal effort choice for type θ under contract w^+,

$$\begin{aligned} v^\theta_{e^+(\theta)}(w^+(\underline{x}), \dot{w}^+) &\geq v^\theta_{\tilde{e}}(w^+(\underline{x}), \dot{w}^+) \\ &= \lim_{n \to \infty} v^\theta_{e(\theta_n)}(w_{\theta_n}(\underline{x}), \dot{w}_{\theta_n}). \end{aligned}$$

Substituting this inequality in (29) for contract $(w_{\theta_n}(\underline{x}), \dot{w}_{\theta_n})$ and taking the limit, we obtain

$$\begin{aligned} v^{\theta}_{e^+(\theta)}(w^+(0), \dot{w}^+) & \geq v^{\theta}_{\bar{e}}(w^+(\underline{x}), w^+) \\ & \geq v^{\theta}_{e^-(\theta)}(w^-(\underline{x}), \dot{w}^-) \end{aligned},$$

verifying that (27) holds. The proof of (28) is analogous. Hence, the mechanism $(\bar{w}, \bar{e})$ is IC.

To conclude the proof, we need to verify that an optimal mechanism exists. The proof, which follows arguments similar to Page (1992), is given in the supplementary appendix.

Proof of Theorem 3

The agent's IC constraint is

$$e(\theta) \in \arg\max_{e} \int \dot{w}(x) \left(1 - F^{\theta}_{e}(x)\right) dx - c^{\theta}_{e}.$$

Use MS to write the IC as

$$\left[I\left(e(\theta), \theta\right) - I(e, \theta)\right] \int \dot{w}(x) H(x) dx \geq c^{\theta}_{e(\theta)} - c^{\theta}_{e}.$$

Notice that if w implements effort $e(\cdot)$, then so does any other $\tilde{w}$ with $\int \dot{w}(x) H(x) dx = \int \dot{\tilde{w}}(x) H(x) dx$.

The principal's expected cost from the contract w is

$$\int \int_{\Theta} \dot{w}(x) \left(1 - F^{\theta}_{e(\theta)}(x)\right) d\mu(\theta) dx.$$

Let (w^*, e) be an optimal mechanism and let $K \equiv \int \dot{w}^*(x) H(x) dx$. Then, w^* must also solve the following program:

$$\min_{0 \leq \dot{w}(x) \leq 1} \int \int_{\Theta} \dot{w}(x) \left(1 - F^{\theta}_{e(\theta)}(x)\right) d\mu(\theta) dx. \tag{30}$$

s.t.

$$\int \dot{w}(x) H(x) dx = K.$$

This is a more restricted program than searching for the optimal contract to implement $e(\cdot)$ since it only looks at contracts that have a fixed $\int \dot{w}(x) H(x) dx$. However, as noted above, any such contract satisfies IC. Since the optimal contract is feasible in this program (by our choice of K), it must be a solution.

The optimality conditions are:

$$\dot{w}(x) = 0 \implies \xi(x) < 0$$

$$\dot{w}(x) = 1 \implies \xi(x) > 0$$

where

$$\xi(x) \equiv -\int_{\Theta} \left(1 - F^{\theta}_{e(\theta)}(x)\right) d\mu(\theta) + \lambda H(x). \tag{31}$$

For each $x \in X$, there are two possibilities: $H(x) = 0$ or $H(x) > 0$. In the first case, $\xi(x) \le 0$. Hence, we can only have $\xi(x) > 0$ if $H(x) > 0$.

Suppose that $H(x) > 0$, so that $\xi(x) \ge 0$ if and only if $\frac{\xi(x)}{H(x)} \ge 0$. Rearranging (31), gives

$$\frac{\xi(x)}{H(x)} \ge 0 \iff \frac{1 - \bar{F}_{e(\theta)}(x)}{H(x)} \le \lambda.$$

The expression on the RHS is not a function of x. Thus, for a debt contract to be optimal, it suffices to show that the LHS is decreasing in x. Differentiating this expression, gives:

$$\frac{d}{dx}\left(\frac{1 - \bar{F}_{e(\theta)}(x)}{H(x)}\right) = -\frac{1 - \bar{F}_{e(\theta)}(x)}{H(x)}\left[\frac{H'(x)}{H(x)} + \frac{\bar{f}_{e(\theta)}(x)}{1 - \bar{F}_{e(\theta)}(x)}\right].$$

Hence, the LHS is decreasing if $\frac{H'(x)}{H(x)} + \frac{\bar{f}_{e(\theta)}(x)}{1-\bar{F}_{e(\theta)}(x)} \ge 0$ for all θ. The following lemma concludes the proof by showing that MLRP implies that this condition holds.

Lemma 4. *Suppose that the family of distributions is ordered by FOSD, and the marginal distribution satisfies MLRP. Then,* $\frac{1-\bar{F}_e(x)}{H(x)}$ *is decreasing in* x *for all* e.

Proof. Recall that, under MS, the family of distributions satisfies MLRP if the ratio

$$\frac{\bar{f}_{e_L}(x)}{\bar{f}_{e_H}(x)} = 1 + \frac{H'(x)}{\bar{f}_{e_H}(x)}\left[\bar{I}(e_H) - \bar{I}(e_L)\right],$$

is decreasing for any $e_L, e_H \in E$ with $I(e_L) < I(e_H)$. Hence, MLRP is equivalent to $\frac{H'(x)}{\bar{f}_e(x)}$ being decreasing for all $e \in E$. Therefore, for any $x_1 > x_0$,

$$\frac{H'(x_1)}{\bar{f}_e(x_1)} \le \frac{H'(x_0)}{\bar{f}_e(x_0)} \therefore H'(x_1)\,\bar{f}_e(x_0) \le H'(x_0)\,\bar{f}_e(x_1).$$

Integrate this on $x_1 \in [x_0, \bar{x}]$ (where $\bar{x} = \sup X$):

$$\int_{x_0}^{\bar{x}} H'(x_1)\,\bar{f}_e(x_0)\,dx_1 \le \int_{x_0}^{\bar{x}} H'(x_0)\,\bar{f}_e(x_1)\,dx_1.$$

Using the fact that $H(\bar{x}) = 0$ and $\bar{F}_e(\bar{x}) = 1$, we obtain:

$$-\bar{f}_e(x_0) H(x_0) \leq H'(x_0)\left(1 - \bar{F}_e(x_0)\right) \therefore \frac{H'(x_0)}{H(x_0)} + \frac{\bar{f}_e(x_0)}{1 - \bar{F}_e(x_0)} \geq 0.$$

Since x_0 is arbitrary, this is equivalent to $\frac{1-\bar{F}_e(x)}{H(x)}$ being decreasing in x. $\square$

Proof of Proposition 1

In order to rewrite this model in the same terms as in Section 3, perform the change of variables:

$$x := S - (1+\lambda)C.$$

We will write contracts in terms of the taxpayer's net surplus x, instead of the firm's production cost C by letting $W_\theta(x) := w_\theta\left(\frac{S-x}{1+\lambda}\right)$. Note that BFD can be rewritten as

$$0 \leq \frac{\partial W_\theta}{\partial x}(x) \leq \frac{1}{1+\lambda}$$

at all points of differentiability of $W_\theta(\cdot)$. The regulator's payoff is

$$\int [x - \lambda W_\theta(x)]\, dG^\theta_{e(\theta)}(x) - c^\theta_{e(\theta)}, \tag{32}$$

where $G^\theta_e(x) := F^\theta_e\left(\frac{S-x}{1+\lambda}\right)$.

Let (W, e) be a feasible mechanism. Construct $W^*(x)$ and $e^*(\theta)$ as in the proof of Theorem 2, and recall that $I(e(\theta), \theta) \leq I(e^*(\theta), \theta)$. Fix a type, say θ. The regulator's payoff from type θ in mechanism (W, e) is

$$\int [x - \lambda W_\theta(x)] dG^\theta_{e(\theta)}(x) - c^\theta_{e(\theta)},$$

and her payoff from θ in the new mechanism is

$$\int [x - \lambda W^*(x)] dG^\theta_{e^*(\theta)}(x) - c^\theta_{e^*(\theta)}.$$

By MS and FOSD, $I(\cdot, \theta)$ is an increasing function and, therefore, $e(\theta) \leq e^*(\theta)$, for all θ. The agent's IC constraint and weak* approximation (see the proof of Theorem 2) imply that

$$\int W_\theta(x) dG^\theta_{e(\theta)}(x) - c^\theta_{e(\theta)} \geq \int W^*(x) dG^\theta_{e^*(\theta)}(x) - c^\theta_{e^*(\theta)},$$

which can be rewritten as

$$\int W_\theta(x)dG^\theta_{e(\theta)}(x) - \int W^*(x)dG^\theta_{e^*(\theta)}(x) \geq c^\theta_{e(\theta)} - c^\theta_{e^*(\theta)}. \tag{33}$$

Using the regulator's payoff (33), we can see that the regulator obtains a gain from θ when she replaces the mechanism if and only if

$$\int x\left(dG^\theta_{e^*(\theta)}(x) - dG^\theta_{e(\theta)}(x)\right) \geq \lambda\left(\int W^*(x)dG^\theta_{e^*(\theta)}(x) - \int W_\theta(x)dG^\theta_{e(\theta)}(x)\right) + \left(c^\theta_{e^*(\theta)} - c^\theta_{e(\theta)}\right),$$

which, from inequality (33), holds if

$$\int x\left(dG^\theta_{e^*(\theta)}(x) - dG^\theta_{e(\theta)}(x)\right) \geq (1+\lambda)\left(c^\theta_{e^*(\theta)} - c^\theta_{e(\theta)}\right). \tag{34}$$

Using integration by parts, the effort that maximizes net-surplus (8) is

$$\bar{e}(\theta) \in \arg\max_e (1+\lambda)^{-1}\int \left(1 - G^\theta_e(x)\right)dx - c^\theta_e.$$

By our assumption that $\bar{e}(\theta)$ is unique and (8) is increasing in $e < \bar{e}(\theta)$, it suffices to show that $\bar{e}(\theta) \geq e^*(\theta)$ in order to establish that (34) holds.

Since e^* maximizes the payoff of the firm's manager, if must give him a higher payoff than $\bar{e}$:

$$\int \dot{W}^*(x)\left(1 - G^\theta_{e^*(\theta)}(x)\right)dx - c^\theta_{e^*(\theta)} \geq \int \dot{W}^*(x)\left(1 - G^\theta_{\bar{e}(\theta)}(x)\right)dx - c^\theta_{\bar{e}(\theta)}.$$

Similarly, because $\bar{e}$ maximizes net surplus, we have

$$(1+\lambda)^{-1}\int\left(1 - G^\theta_{\bar{e}(\theta)}(x)\right)dx - c^\theta_{\bar{e}(\theta)} \geq (1+\lambda)^{-1}\int\left(1 - G^\theta_{e^*(\theta)}(x)\right)dx - c^\theta_{e^*(\theta)}.$$

Using these two inequalities, we obtain

$$\int\left[(1+\lambda)^{-1} - \dot{W}^*(x)\right]\left(G^\theta_{e^*(\theta)}(x) - G^\theta_{\bar{e}(\theta)}(x)\right)dx \geq 0.$$

By BFD, the term inside the first brackets are non-negative for all x. Since G^θ_e is ordered by FOSD, the term inside the second brackets have a constant sign, which is positive if $\bar{e}(\theta) \geq e^*(\theta)$ and negative if $\bar{e}(\theta) \leq e^*(\theta)$. Then, the inequality above implies that $\bar{e}(\theta) \geq e^*(\theta)$, concluding the proof.

Proof of Proposition 2

The proof follows the same steps as the proof of Theorem 3. Use MS to write the IC constraint as

$$[I(e(\theta),\theta) - I(e,\theta)] \int \dot{w}(C)\, H(C)dC \geq c^{\theta}_{e(\theta)} - c^{\theta}_{e}.$$

Thus, if w implements effort $e(\cdot)$, then so does any other $\tilde{w}$ with $\int \dot{w}(C)\, H(C)dC = \int \dot{\tilde{w}}(C)\, H(C)dC$. Let (w^*, e) be an optimal mechanism and let $K \equiv \int \dot{w}^*(C)\, H(C)dC$. Then, w^* must also solve the following program:

$$\min_{-1 \leq \dot{w}(C) \leq 0} \int \int_{\Theta} \dot{w}(C) \left(1 - F^{\theta}_{e(\theta)}(C)\right) d\theta dC$$

subject to

$$\int \dot{w}(C)\, H(C)dC = K,$$

where we removed from the objective function all terms that are not affected by $\dot{w}$. As in the proof of Theorem 3, this is a restricted program since it takes effort as fixed and only considers contracts with a fixed $\int \dot{w}(C)H(C)dC$. Notice that this is exactly the same program as (30), except that now $\dot{w}$ is in $[-1, 0]$ instead of $[0, 1]$. The same argument establishes that there exists $\bar{C}$ such that $\dot{w}(C) = -1$ if $C \leq \bar{C}$ and $\dot{w} = 0$ otherwise (recall that we assumed that e orders F by FOSD in this model). Integrating and using the binding LL constraint, we obtain $w(C) = \max\left\{\bar{C} - C; 0\right\}$.

C. Examples

Screening with Pure Adverse Selection

This appendix shows that, with pure adverse selection, it is typically sub-optimal to offer a single contract even when there are only two outputs. Thus, moral hazard is important for the simplicity result from Theorem 1. We consider a simple counterexample. There are two states (H and L, which will be referred to as "high" and "low" outputs), two efforts (0 and 1, or "low" and "high" efforts), and two types (A and B). The effort costs are $c^A_1 = 1$, $c^B_1 = \frac{2}{3}$, and $c^A_0 = c^B_0 = 0$.

Given a high effort, the probability of success for type A is $p^A_1 = \frac{2}{3}$ and for type B is $p^B_1 = \frac{1}{3}$. We assume that the project fails with a high enough probability if they exert low effort and we take $x_H - x_L$ to be large enough for the principal to want to implement high effort from both types. Then, the optimal mechanism must solve the following program:

$$\begin{array}{ll} \min\limits_{w_H^i, w_L^i \geq 0} & 2w_H^A + w_L^A + w_H^B + 2w_L^B \\ \text{subject to} & 2w_H^A + w_L^A \geq 2w_H^B + w_L^B \\ & w_H^B + 2w_L^B \geq w_H^A + 2w_L^A \\ & 2w_H^A + w_L^A \geq 3 \\ & w_H^B + 2w_L^B \geq 2. \end{array}$$

The two first constraints require A and B to prefer to report their types truthfully (IC constraints). Because effort is observable, the principal does not need to worry about deviations on effort. Then, it is no longer the case that LL implies IR. The last two constraints are precisely the IR constraints.

It is straightforward to show that the unique solution offers the following payments: $w_H^A = \frac{3}{2}$, $w_L^A = 0$ and $w_H^B = w_L^B = \frac{2}{3}$. Moreover, this mechanism is no longer feasible if effort is not observable. In fact, both types would choose $e = 0$ if offered these contracts and effort was unobservable.

Example 1 (continuation)

Let $w^\theta = (w_H^\theta, w_M^\theta, w_L^\theta)$ denote the vector of payments to the agent. Their conditional probability distributions are given in the text. There are six ICs constraints that must be checked.

- ICs preventing deviations in effort for a fixed contract:

$$\frac{2w_H^A + 2w_M^A + w_L^A}{5} - 1 \geq \frac{w_H^A + w_M^A + w_L^A}{3}, \tag{35}$$

$$\frac{2w_H^B + 3w_M^B + w_L^B}{6} - 1 \geq \frac{w_H^B + w_M^B + 3w_L^B}{3}. \tag{36}$$

- ICs preventing deviations in both contracts and efforts:

$$\frac{2w_H^A + 2w_M^A + w_L^A}{5} - 1 \geq \frac{w_H^B + w_M^B + w_L^B}{3}, \tag{37}$$

$$\frac{2w_H^B + 3w_M^B + w_L^B}{6} - 1 \geq \frac{w_H^A + w_M^A + w_L^A}{5}. \tag{38}$$

- ICs preventing deviations in contract for a fixed effort:

$$2w_H^A + 2w_M^A + w_L^A \geq 2w_H^B + 2w_M^B + w_L^B, \tag{39}$$

$$2w_H^B + 3w_M^B + w_L^B \geq 2w_H^A + 3w_M^A + w_L^A. \tag{40}$$

The principal's objective is to minimize her cost

$$\frac{2w_H^A + 2w_M^A + w_L^A}{5} + \frac{2w_H^B + 3w_M^B + w_L^B}{6}$$

subject to the constraints above. This is a standard linear program, and it is straightforward to check that the solution is $w^A = (0, 0, 15)$ and $w^B = (0, 6, 6)$, which gives the principal a cost of 11.

To verify that we can't implement the optimal mechanism with a single contract, consider the principal's program if we impose that she gives the same contract (w_H, w_M, w_L) to both types. The principal's cost is then

$$\frac{2w_H + 2w_M + w_L}{5} + \frac{2w_H + 3w_M + w_L}{6} = \frac{11w_L + 27w_M + 22w_H}{30}.$$

Since there is only one contract, the only relevant IC is the one preventing each type from choosing a different effort:

$$\frac{2w_H + 2w_M + w_L}{5} - 1 \geq \frac{w_H + w_M + w_L}{3},$$

and

$$\frac{3w_M + w_L}{6} - 1 \geq \frac{w_M + 3w_L}{3}.$$

It is straightforward to verify that the solution is $w_L = 0,\ w_M = 6,\ w_H = 9$, which gives the principal a cost of 12. Thus, offering the best single-contract mechanism yields a higher cost than offering the optimal menu of contracts in this example.

D. Multiplicative Separability

Recall the ordering condition presented in the text, which we will refer to as the ordering (O) condition:

Definition 2. The distribution of outputs satisfies the ordering (O) condition if given contracts w and $\tilde{w}$ that satisfy BFD and LL and $e, \tilde{e} \in E$, if there exists $\theta \in \Theta$ for which

$$\int w(x) \left(dF_e^\theta(x) - dF_{\tilde{e}}^\theta(x)\right) dx \geq \int \tilde{w}(x) \left(dF_e^\theta(x) - dF_{\tilde{e}}^\theta(x)\right),$$

then, for all $\tilde{\theta}$,

$$\int w(x) \left(dF_e^{\tilde{\theta}}(x) - dF_{\tilde{e}}^{\tilde{\theta}}(x)\right) dx \geq \int \tilde{w}(x) \left(dF_e^{\tilde{\theta}}(x) - dF_{\tilde{e}}^{\tilde{\theta}}(x)\right).$$

Substitution establishes that MS implies O. The next lemma shows that the reverse is also true, so these conditions are equivalent to each other:

Lemma 5. *MS and O are equivalent.*

Proof. First consider the case $E = \{0, 1\}$ and let $\theta \in \Theta$. Since $F_1^\theta(x) - F_0^\theta(x) \to 0$, as $x \to -\infty$ and $x \to \infty$, integration by parts gives

$$\int w(x) \left(dF_1^\theta(x) - dF_0^\theta(x)\right) = \int \dot{w}(x) \left(F_0^\theta(x) - F_1^\theta(x)\right) dx$$

for any contract $w : X \to \mathbb{R}_+$ that satisfies BFD and LL, i.e., $\dot{w} \in L^\infty(X)$ and $0 \leq \dot{w}(x) \leq 1$, for almost all $x \in X$. Define the linear functional $\varphi^\theta : L^\infty(X) \to \mathbb{R}$ by $\varphi^\theta(\dot{w}) = \int \dot{w}(x) \left(F_0^\theta(x) - F_1^\theta(x)\right) dx$. Condition O implies that

$$\varphi^\theta(\dot{w}) = 0, \quad \forall \dot{w} \in L_+^\infty(X) \text{ if and only if } \varphi^{\tilde{\theta}}(\dot{w}) = 0, \quad \forall \dot{w} \in L_+^\infty(X),$$

for all $\theta, \tilde{\theta} \in \Theta$. Since $\dot{w}(x) = \dot{w}_+(x) - \dot{w}_-(x)$, where $\dot{w}_+(x) = \max\{\dot{w}(x), 0\}$ and $\dot{w}_-(x) = \max\{-\dot{w}(x), 0\}$, the previous equivalence is also true replacing $L_+^\infty(X)$ by $L^\infty(X)$. Hence, functionals φ^θ and $\varphi^{\tilde{\theta}}$ are equivalent, for all $\theta, \tilde{\theta} \in \Theta$, i.e., there exist constants $\lambda^\theta, \lambda^{\tilde{\theta}} > 0$ and linear functional φ in $L^1(X)$ such that $\varphi^\theta = \lambda^\theta \varphi$ and $\varphi^{\tilde{\theta}} = \lambda^{\tilde{\theta}} \varphi$. Indeed, we have that the null spaces of φ^θ and $\varphi^{\tilde{\theta}}$ are the same, which we denote by N. By the Rank-Nullity Theorem, there exists $v \in L^\infty(X) \backslash N$ such that $L^\infty(X) = [v] \oplus N$, where $[v]$ is the subspace generated by vector v and $\oplus$ represents the direct sum between vector spaces. Let φ be the unique linear functional such that $\varphi(v) = 1$ and $\varphi(n) = 0$ for all $n \in N$. Define $\lambda^\theta = \varphi^\theta(v)$ and $\lambda^{\tilde{\theta}} = \varphi^{\tilde{\theta}}(v)$. Notice that λ^θ and $\lambda^{\tilde{\theta}}$ have the same sign because φ^θ and $\varphi^{\tilde{\theta}}$ "point" in the same direction according to O, which we can assume to be positive without loss of generality (otherwise, we define φ such that $\varphi(v) = -1$). The result then follows immediately. For a general set E, we can apply the same argument to show that the linear functionals $\varphi^{\theta,e,\hat{e}}$ defined by $\varphi^{\theta,e,\hat{e}}(\dot{w}) = \int \dot{w}(x) \left(F_{\hat{e}}^\theta(x) - F_e^\theta(x)\right) dx$ are equivalent, for all $\theta \in \Theta$ and $e, \hat{e} \in E$. $\square$

E. Converse of Theorem 3

We will write **1** to denote the indicator function. Let $\bar{I}(e) := \int I(e, \theta) \mu(d\theta)$, and let $\varphi(e, x) := \frac{1 - \bar{F}_e(x)}{H(x)} + \bar{I}(e)$. By MS, we have

$$\frac{1 - F_{\tilde{e}}^\theta(x)}{H(x)} + I(\tilde{e}, \theta) = \frac{1 - F_e^\theta(x)}{H(x)} + I(e, \theta) \quad \forall \tilde{e}, e, \theta, x.$$

Integrating over θ, gives $\varphi(\tilde{e},x) = \varphi(e,x)$ for all $e, \tilde{e}$. Therefore, φ depends only on x and we will drop the dependency on e. The following result determines characterizes MLRP:[22]

Lemma 6. *Suppose F_e^θ satisfies MS and is ordered by FOSD. F_e^θ satisfies MLRP if and only if $\varphi'(x) \leq 0$ for all $x \in X$.*

Therefore, the distribution fails MLRP if and only if $\varphi'(x_0) > 0$ for some x_0. Other than MS, we will assume that the distribution satisfies the following condition:

Definition 3. Let F_e^θ satisfy MS. F_e^θ is regular if

$$\left[1 + \frac{\int \mathbf{1}_{\varphi(x_0) \geq \varphi(x)} H(x)dx}{\int H(x)dx}\right] \bar{I}(e^*) \geq \bar{I}(e_*) + \frac{\int \mathbf{1}_{\varphi(x_0) \geq \varphi(x)} \varphi(x) H(x)dx}{\int H(x)dx}$$

for some $e^*, e_* \in E$ and $x_0 \in X$ with $\varphi'(x_0) > 0$.

The regularity condition ensures that there exist cost functions for which the contract that always pays zero is not optimal. The following proposition presents the partial converse to Theorem 3:

Proposition 3. *Suppose the output distribution satisfies MS, is ordered by FOSD, and is regular. If the marginal distribution does not satisfy MLRP, there exists a cost function $c : \Theta, E \to \mathbb{R}_+$ for which the optimal mechanism does not give the principal a debt contract.*

References

Bajari, P. and S. Tadelis (2001): "Incentives Versus Transaction Costs: A Theory of Procurement Contracts," *RAND Journal of Economics*, 32, 287–307.

Biais, B. and T. Mariotti (2005): "Strategic Liquidity Supply and Security Design," *Review of Economic Studies*, 72, 615–649.

Burguet, R., J.-J. Ganuza, and E. Hauk (2012): "Limited Liability and Mechanism Design in Procurement," *Games and Economic Behavior*, 76, 15–25.

Caillaud, B., R. Guesnerie, and P. Rey (1992): "Noisy Observation in Adverse Selection Models," *Review of Economic Studies*, 59, 595–615.

Carroll, G. (Forthcoming): "Robustness and Linear Contracts," *American Economic Review.*

Chaigneau, P., A. Edmans, and D. Gottlieb (2014): "The Value of Informativeness for Contracting," Tech. rep., HEC Montreal, LBS, and Wharton.

[22] The proofs of all results in this section are in the supplementary appendix.

Chassang, S. (2013): "Calibrated Incentive Contracts," *Econometrica*, 81, 1935–1971.

Chiappori, P.-A. and B. Salanie (2003): "Testing Contract Theory: A Survey of Some Recent Work," in *Advances in Economics and Econometrics*, ed. by M. Dewatripont, L. P. Hansen, and S. T. Turnovsky, Cambridge: Cambridge University Press, vol. 1.

Chu, L. Y. and D. Sappington (2007): "Simple Cost-Sharing Contracts," *American Economic Review*, 97, 419–428.

Demarzo, P. and D. Duffie (1999): "A Liquidity-based Model of Security Design," *Econometrica*, 67, 65–99.

DeMarzo, P. M. (2005): "The Pooling and Tranching of Securities: A Model of Informed Intermediation," *Review of Financial Studies*, 18, 1–35.

DeMarzo, P. M., I. Kremer, and A. Skrzypacz (2005): "Bidding with Securities: Auctions and Security Design," *American Economic Review*, 95, 936–959.

Dewatripont, M., P. Legros, and S. A. Matthews (2003): "Moral Hazard and Capital Structure Dynamics," *Journal of the European Economic Association*, 1, 890–930.

Edmans, A. and X. Gabaix (2011): "Tractability in Incentive Contracting," *Review of Financial Studies*, 24, 2865–94.

Farhi, E. and J. Tirole (Forthcoming): "Liquid Bundles," *Journal of Economic Theory.*

Gottlieb, D. and H. Moreira (2014): "Simultaneous Adverse Selection and Moral Hazard," Tech. rep., Wharton School and EPGE/FGV.

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

Guesnerie, R. and J.-J. Laffont (1984): "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm," *Journal of Public Economics*, 25, 329–369.

Hart, O. D. and B. Holmstrom (1987): "The Theory of Contracts," in *Advances in Economic Theory, Fifth World Congress*, ed. by T. Bewley, Cambridge: Cambridge University Press.

Holmstrom, B. (1979): "Moral Hazard and Observability," *The Bell Journal of Economics*, 10, 74–91.

Holmstrom, B. and P. Milgrom (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55, 303–28.

Innes, R. D. (1990): "Limited Liability and Incentive Contracting with Ex-Ante Action Choices," *Journal of Economic Theory*, 52, 45–67.

Jewitt, I., O. Kadan, and J. M. Swinkels (2008): "Moral Hazard with Bounded Payments," *Journal of Economic Theory*, 143, 59–82.

Laffont, J.-J. and D. Martimort (2002): *The Theory of Incentives - Part I*, Princeton University Press.

Laffont, J.-J. and J. Tirole (1986): "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 94, 614–641.

——— (1993): *A Theory of Incentives in Procurement and Regulation*, MIT press.

Matthews, S. A. (2001): "Renegotiating Moral Hazard Contracts under Limited Liability and Monotonicity," *Journal of Economic Theory*, 97, 1–29.

Melumad, N. D. and S. Reichelstein (1989): "Value of Communication in Agencies," *Journal of Economic Theory*, 47, 334–368.

Nachman, D. C. and T. H. Noe (1994): "Optimal design of securities under asymmetric information," *Review of Financial Studies*, 7, 1–44.

Ollier, S. and L. Thomas (2013): "Ex Post Participation Constraint in a Principal–Agent Model with Adverse Selection and Moral Hazard," *Journal of Economic Theory*, 148, 2383–2403.

Page, F. H. (1992): "Mechanism Design for General Screening Problems with Moral Hazard," *Economic Theory*, 2, 265–281.

Picard, P. (1987): "On the Design of Incentive Schemes under Moral Hazard and Adverse Selection," *Journal of Public Economics*, 33, 305–31.

Poblete, J. and D. Spulber (2012): "The Form of Incentive Contracts: Agency with Moral Hazard, Risk Neutrality, and Limited Liability," *RAND Journal of Economics*, 43, 215–34.

Rochet, J.-C. and L. A. Stole (2002): "Nonlinear Pricing with Random Participation," *Review of Economic Studies*, 69, 277–311.

Rogerson, W. P. (2003): "Simple Menus of Contracts in Cost-Based Procurement and Regulation," *American Economic Review*, 93, 919–26.

Rudin, W. (1991): *Functional Analysis*, McGraw-Hill, Inc, 2nd edition ed.

Online Appendix

Existence with Two Outcomes

From Lemma 3 in the proof of Theorem 1, there is no loss of generality in restricting the set of mechanisms to those consisting of a single contract that pays zero if output is low and a bonus if the output is high. Each mechanism consists of a bonus $b \in [0, \Delta x]$ and effort recommendation (which is the agent's preferred effort given the offered bonus).

Let $v_e^\theta(b) := p_e^\theta b - c_e^\theta$ denote the payoff of type θ who chooses effort e. By Assumption 1, $v_\cdot^\theta(\cdot) : [0, \Delta x] \times E \to \mathbb{R}$ is continuous and, for each $(b, e) \in [0, \Delta x] \times E$, $v_e^\cdot(b) : \Theta \to \mathbb{R}$ is $\mathcal{B}(\Theta)$-measurable.

Since $v_\cdot^\theta(b)$ continuous and E is compact, the IC constraint gives, for each fixed bonus, at least one optimal effort $e(\theta)$. Let $V^\theta(b) := \max_{e \in E} v_e^\theta(b)$ be the payoff of type θ under bonus b. Let $(\theta, b) \to e^*(\theta, b)$ denote the non-empty, compact-valued mapping specifying the optimal efforts of type θ under bonus b:

$$e^*(\theta, b) = \{e \in E;\ v_e^\theta(b) \geq V^\theta(b)\}.$$

By Berge's Maximum Theorem, $b \to V^\theta(b)$ is continuous for each θ and $b \to e^*(\theta, b)$ is upper semicontinuous on $[0, \Delta x]$ for each θ.

Let $u_e^\theta(b) := p_e^\theta(\Delta x - b)$ denote the principal's payoff from a type $\theta \in \Theta$ who is offered bonus $b \in [0, \Delta x]$ and exerts effort $e \in E$. By Assumption 1, for each $\theta \in \Theta$, $u_\cdot^\theta(\cdot) : [0, \Delta x] \times E \to \mathbb{R}$ is continuous and, for each $(b, e) \in [0, \Delta x] \times E$, $u_e^\cdot(b) : \Theta \to \mathbb{R}$ is $\mathcal{B}(\Theta)$-measurable. Notice that p is μ-integrable on $\Theta \times E$ because $|p_e^\theta| \leq 1$ on $\Theta \times E$.

Since $u_\cdot^\theta(b)$ continuous and $e^*(\theta, b)$ is compact, for each (θ, b), the problem:

$$\max_{e \in e^*(\theta, b)} u_e^\theta(b)$$

has a solution, say e_b^θ. Thus, for each bonus b, the principal can provide the agent an effort recommendation function $e_b^\theta : \Theta \to E$, which specifies an effort to each type. Since $e_b^\theta \in e^*(\theta, b)$ for all θ, no type has an incentive to deviate from the recommended effort.

To choose the best effort recommendation, the principal selects, for each $(\theta, b) \in \Theta \times [0, \Delta x]$, the effort that maximizes her payoff. Let

$$U^\theta(b) := \max_{e \in e^*(\theta, b)} u_e^\theta(b)$$

denote the principal's maximum payoff among all efforts that are incentive compatible to type θ when the bonus equals b. Since u is $\mathcal{B}(\Theta) \times \mathcal{B}([0, \Delta x]) \times \mathcal{B}(E)$-measurable and continuous on E, and $e^*(\cdot, \cdot)$ is $\mathcal{B}(\Theta) \times \mathcal{B}([0, \Delta x])$-measurable (see Nowak (1984), Lemma 1.10) and compact-valued, it follows that U is $\mathcal{B}(\Theta) \times \mathcal{B}([0, \Delta x])$-measurable (see Himmelberg et.al. (1976), Theorem 2).

Moreover, since $e^*(\theta, \cdot)$ is upper semicontinuous on $[0, \Delta x]$, it follows from Berge (1963), Theorem 2, that $U^\theta(\cdot)$ is upper continuous on $[0, \Delta x]$ for each θ. Finally, since the principal's payoff u is integrably bounded, $U : \Theta \times [0, \Delta x] \to \mathbb{R}$ is also integrably bounded on $\Theta \times [0, \Delta x]$.

Given these observations, the principal's problem consists of maximizing a continuous function over a compact interval,

$$\max_{b \in [0, \Delta x]} \int U^\theta(b) d\mu(\theta),$$

which must then have a solution.

Existence with Multiple Outputs

The proof follows similar steps as in the two-output case. Let

$$\mathcal{W} := \{w : X \to [0, \Delta x] \text{ satisfying } BFD,\ LL \text{ and } w(0) = 0\}$$

denote the space of contracts. As noticed in the proof of Theorem 2, for each $w \in \mathcal{W}$, $\dot{w} \in L^\infty(X)$. Consider the following topology in $\mathcal{W}$: a net (w_n) converges to w in $\mathcal{W}$ if and only if $(\dot{w}_n)$ converges to $\dot{w}$ in the $(L^\infty(X), L^1(X))$-weak* sense. Under this topology, $\mathcal{W}$ is metrizable and compact (see Rudin (1991), Theorem 3.16). Let $\mathcal{B}(\mathcal{W})$ be the Borel σ-field in $\mathcal{W}$ with respect to the metric induced by this topology.

Let $v_e^\theta(w) := \int w(x) dF_e^\theta(x) - c_e^\theta$. By Assumption 2, $v_\cdot^\theta(\cdot)$ is continuous on $\mathcal{W} \times E$, and, for each $(w, e) \in \mathcal{W} \times E$, $v_e^\cdot(w)$ is $\mathcal{B}(\Theta)$-measurable. Since E is compact and $v_\cdot^\theta(w)$ continuous, the agent's effort choice problem has a solution. Let $V^\theta(w) := \sup_{e \in E} v_e^\theta(w)$ denote the utility of type θ. Let the non-empty, compact-valued mapping $(\theta, w) \to e^*(\theta, w)$ denote the "reaction function" of type θ:

$$e^*(\theta, w) := \{e \in E : \ v_e^\theta(w) \geq V^\theta(w)\}.$$

By Berge's Maximum Theorem, $w \to V^\theta(w)$ is continuous for each θ and $w \to e^*(\theta, w)$ is upper semicontinuous on $\mathcal{W}$ for each θ.

Let $u_e^\theta(w) := \int (x - w(x)) dF_e^\theta(x)$ denote the principal's payoff. By Assumption 2, for each $\theta \in \Theta$, $u_\cdot^\theta(\cdot)$ is continuous on $\mathcal{W} \times E$, and for each $(w, e) \in \mathcal{W} \times E$, $u_e^\cdot(w)$ is $\mathcal{B}(\Theta)$-measurable. By Assumption 2, u is μ-integrable on $\Theta \times E$ (i.e., $|u_e^\theta(w)| \leq \xi(\theta)$ on $\Theta \times \mathcal{W} \times E$ where ξ is an integrable real-valued function defined on Θ).

Since $e^*(\theta, w)$ is compact and $u_\cdot^\theta(w)$ continuous, for each (θ, w), the problem $\sup_{e \in e^*(\theta, w)} u_e^\theta(w)$ has a solution, say e_w^θ. Thus, for each contract w, the principal can provide the agent with a list, $\{e_w^\theta; \theta \in \Theta\}$, of recommended efforts. Since $e_w^\theta \in e^*(\theta, w)$ for each $\theta \in \Theta$, a type-θ agent has no incentive to take an effort other than the one requested by the principal e_w^θ (i.e., the agent is obedient). To choose the best list of requests, the principal must therefore solve the problem

$$\sup_{e\in e^*(\theta,w)} u_e^\theta(w)$$

for each $(\theta,w)\in\Theta\times\mathcal{W}$. Let $U^\theta(w):=\sup_{e\in e^*(\theta,w)} u_e^\theta(w)$. Since u is $\mathcal{B}(\Theta)\times\mathcal{B}(\mathcal{W})\times\mathcal{B}(E)$-measurable and continuous on E, and since $e^*(\cdot,\cdot)$ is $\mathcal{B}(\Theta)\times\mathcal{B}(\mathcal{W})$-measurable (see Nowak (1984), Lemma 1.10) and compact-valued, it follows that U is $\mathcal{B}(\Theta)\times\mathcal{B}(\mathcal{W})$-measurable (see Himmelberg et.al. (1976), Theorem 2). Moreover, since $e^*(\theta,\cdot)$ is upper semicontinuous on $\mathcal{W}$, it follows from Theorem 2 of Berge (1963) that $U^\theta(\cdot)$ is upper semicontinuous on $\mathcal{W}$ for each θ. Finally, since the principal's payoff u is integrably bounded, $U:\Theta\times\mathcal{W}\to\mathbb{R}$ is also integrably bounded on $\Theta\times\mathcal{W}$.

With these observations, we can write the principal's program as

$$\sup_{w\in\mathcal{W}}\int U^\theta(w)d\mu(\theta)$$

which has a solution since $\mathcal{W}$ is compact.

Proofs from Appendix E

Proof of Lemma 6.

Sufficiency is immediate. To establish the necessity part, let $e\in E$ and $x_0\in X$ be such that the derivative of $\frac{1-\bar{F}_e(x)}{H(x)}$ at $x=x_0$ is positive:

$$-\frac{\bar{f}_e(x_0)}{H(x_0)}-\frac{(1-\bar{F}_e(x_0))H'(x_0)}{H(x_0)^2}>0,$$

Because $H(x)\geq 0$ (ordering by FOSD), we have

$$\frac{H'(x_0)}{\bar{f}_e(x_0)}<-\frac{H(x_0)}{1-\bar{F}_e(x_0)}\leq 0,$$

which shows that $\frac{H'(x_0)}{\bar{f}_e(x_0)}<0$, contradicting MLRP. □

Proof of Proposition 3.

Fix $e^*\in E$ and $x_0\in X$ that satisfy the regularity condition. Construct the contract w^* by setting $w^*(0)=0$ and

$$\dot{w}^*(x)=\begin{cases}1, & \text{if } \varphi(x_0)\geq\varphi(x)\\ 0, & \text{if otherwise}\end{cases}. \tag{41}$$

Let $e(\theta) = e^*$. Pick the following cost function:

$$c_e^\theta = \begin{cases} 0 & \text{if } e = e_* \\ [I(e^*, \theta) - I(e_*, \theta)] \int \dot{w}^*(x) H(x) dx & \text{if } e = e^* \\ \infty & \text{if otherwise} \end{cases}.$$

Type θ's IC constraint is

$$[I(e(\theta), \theta) - I(e, \theta)] \int \dot{w}^*(x) H(x) dx \geq c_{e(\theta)}^\theta - c_e^\theta, \quad \forall e.$$

By our choice of the cost function, the IC holds with equality at $e = e_*$ and inequality at other efforts. From the proof of Theorem 3, w^* is the cost-minimizing contract that implements $e = e^*$.

The principal's payoff from implementing $e = e_*$ is

$$\int x d\bar{F}_{e_*}(x) = \int \left[1 - \bar{F}_{e_*}(x)\right] dx.$$

Her payoff from implementing $e = e^*$ is

$$\int [x - w^*(x)] d\bar{F}_{e^*}(x) dx = \int [1 - \dot{w}^*(x)] \left[1 - \bar{F}_{e^*}(x)\right] dx.$$

The difference in the principal's payoff from implementing e^* and e_* equals

$$[\bar{I}(e^*) - \bar{I}(e_*)] \int H(x) dx - \int \dot{w}^*(x) \left[1 - \bar{F}_{e^*}(x)\right] dx$$

$$= [\bar{I}(e^*) - \bar{I}(e_*)] \int H(x) dx - \int_{[\varphi(x_0) - \varphi(\cdot) \geq 0]} \left[1 - \bar{F}_{e^*}(x)\right] dx \geq 0,$$

where the equality uses (41) and the inequality follows by the regularity condition. Thus, it is optimal for the principal to implement $e(\theta) = e^*$. Next, we verify that w^* is not a debt contract. There are two possible cases:

- Suppose x_0 is the only solution to the equation $\varphi(x) = \varphi(x_0)$ in the interior of X. Then $w^*(x) = \min\{x, x_0\}$, which is a debt contract to the agent so the principal's payment ($\max\{x - x_0,\ 0\}$) is not a debt contract (it is a call option).

- Now suppose that $\varphi(x) = \varphi(x_0)$ has more than one solution in the interior of X. Then, there exists a non-degenerate interval $[a, b] \subset \operatorname{int}(X)$ with $\dot{w}^*(x) = 1$ for $x \in [a, b]$ and $\dot{w}^*(x) = 0$ for some $x > b$, which shows that the principal does not get a debit contract. $\square$

References

Berge, C. (1963). Topological Spaces. MacMillan.

Himmelberg, C.J., Parthasarathy, T. and VanVleck, F.S. (1976), "Optimal plans for dynamic programming problems," Mathematics of Operations Research, Mathematics of Operations Research, 1, pp. 390-394.

Nowak, A.S. (1984), "On zero-sum stochastic games with general state space.I," Probability and Mathematical Statistics, 4(1), pp. 13-32.

Rudin, W. (1991). Functional Analysis. McGraw-Hill, Inc, 2nd edition.