Optimal monopoly regulation with flexible investments

Daniel Krähmer*

April 14, 2025

Abstract

I study the optimal regulation of a monopolist who can covertly invest in cost reductions and privately observes the resulting costs. Thus, there is moral hazard and adverse selection. The monopolist's investment is flexible: she can choose any cost distribution subject to an investment cost. I show that inefficiencies are driven by the convexity of the investment cost function. Optimal regulatory policies are simple, featuring finitely many options and inducing finite cost distributions. I identify conditions under which a fixed-price contract is optimal. A key step in the analysis is a characterization of implementable distributions.

Keywords: Moral Hazard, Adverse Selection, Limited Liability, Flexible Investment, Agency Costs

JEL: D42, D82, D86

*Universität Bonn (kraehmer@hcm.uni-bonn.de); I thank Helmut Bester, Görkem Celic, Siwen Liu, Nico Schutz, Roland Strausz, and my theory colleagues in Bonn for very valuable discussions. Special thanks go to Mathijs Janssen. I gratefully acknowledge financial support from the German Research Foundation (DFG) through Germany's Excellence Strategy EXC 2126/1–390838866 and CRC TR 224 (project B03).

1 Introduction

An important issue in the regulation of monopolies is to incentivize the regulated firm to make investments that reduce its production costs. In practice, this problem is aggravated by the fact that regulators frequently lack the expertise and the resources to monitor the regulated firms investment efforts and observe the realized cost reductions, that is, there is both moral hazard and adverse selection. This commonly arises in the regulation of utilities, where firms can covertly invest in infrastructure improvements, or in public procurement, where firms can invest in cost-reducing technologies, and in each case privately observe the resulting cost reductions.[1]

In this paper, I study the optimal regulation of a monopolist in the presence of these joint moral hazard and adverse selection problems. The regulator's objective is to design a contract that both induces appropriate investment and elicits the resulting cost reductions truthfully. The distinctive modelling feature of the paper is that the monopolist's investment is flexible, that is, she can choose *any* probability distribution of production costs (her "type") subject to an investment cost. While the set of possible cost types is assumed to be an interval, there is no restriction on the distributions the monopolist can choose. Investment costs are assumed to be monotone, convex, and smooth. Monotonicity, which means that investment costs increase if the monopolist chooses a cost distribution that is lower in the first order stochastic dominance sense, and convexity, which captures increasing marginal investment costs, are natural in the investment context considered here. Smoothness means that the investment cost function admits a Gateaux-derivative which makes the analysis amenable to familiar marginal utility reasoning.[2]

Two central insights emerge from my analysis. First, all inefficiencies that arise under an optimal regulation can, ultimately, be attributed to the convexity of the investment cost function. I demonstrate that if investment costs are strictly convex, optimal regulation induces some inefficiencies (in investments and/or production). In notable contrast, if investment costs are linear, the optimal regulation implements the first-best despite the presence of moral hazard and adverse selection.[3] Second, for a large class of natural investment cost functions, "simple" outcomes emerge: optimal contracts amount to finite menus and induce finite cost distributions. In

[1]Similar issues arise in other principal-agent relationships where the agent can make covert investments and privately observes her realized payoff type such as employment relationships, supply chains, or in consumers markets to name a few. With the appropriate re-interpretation of variables, my analysis carries over to these settings.

[2]Recall that the Gateaux-derivative is a functional derivative that generalizes the notion of a partial derivative from functions of vectors to functions of functions. Economically, the Gateaux-derivative evaluated at a given distribution and a given type measures the marginal cost of increasing the probability mass on this type.

[3]Linear costs reflect constant marginal costs and correspond to the case where the monopolist can generate stochastic investment outcomes only by randomizing over deterministic outcomes.

fact, in a variety of cases, the optimal contract is a fixed-price contract or a fixed-price-award-fee contract.[4] In other words, the simplicity of contracts is an optimal regulatory response that encourages investment and screens the monopolist. To the extent that I allow for a rich type space and that the flexible approach does not impose distributions to be continuous or discrete a priori, this provides a novel rationale for the simplicity of real world regulation and procurement contracts.[5]

To derive these results, I combine the classical regulation model of Baron and Myerson (1982) with the flexible moral hazard model of Georgiadis et al. (2024). The regulator commits to a contract that specifies a menu of production levels and reimbursement rates. The monopolist subsequently invests by covertly choosing a cost distribution and then selects terms from the menu after having observed her true costs. I assume that the monopolist is protected by limited liability and thus cannot sustain losses at the production stage. The flexible approach frees the analysis from restrictive parametric functional form assumptions, thus allowing to fully endogenize the monopolist's cost distribution. An important conceptual lesson of the paper is that the flexible approach also proves remarkably tractable and delivers rich insights where parametric approaches might be difficult to tackle.[6]

As an essential step in formulating the regulator's problem, I first characterize the set of outcomes that can be implemented by some contract. The characterization combines the restrictions that a contract has to induce truth-telling due to adverse selection and that investment choice is optimal for the monopolist due to moral hazard. Intuitively, given a contract, the monopolist chooses a distribution so that marginal benefits equal marginal costs from investment. With smooth costs, this is formally captured by a first order condition (see Georgiadis et al., 2024) requiring that the monopolist's investment distribution place probability mass only on cost types whose interim utility from the contract[7] is equal to the Gateaux-derivative of the cost function at

[4]These contracts are frequently used in infrastructure and defense procurement. Under an FPAF contract, the contractor is paid an award fee if she achieves certain objective, pre-determined performance criteria. These can include time to completion, and various reliability and performance targets (e.g. fuel consumption in a military vehicle).

[5]The gap between the complexity of regulatory schemes often predicted by standard theory and their simplicity as observed in practice has been a concern for the literature. Orthogonally to my findings, it has been proposed that simple menus, while not optimal, still achieve high revenues (e.g. Gami et al., 1999, Rogerson, 2003). Simple ("sparse") menus can also arise under specific conditions on the virtual surplus such as failure of single-peakedness (Anderson and Dana, 2009, Sandmann, 2024).

[6]Such a parametric approach would need to specify a parameterized, stochastically ordered family of distributions where the monopolist chooses the parameter at a cost. Finding an optimal contract is then difficult because the monopolist's moral hazard constraint might not be characterized by a first order condition, and ensuring truth-telling might necessitate the use of ironing techniques. None of these issues will arise in the flexible setting I consider.

[7]The interim utility assigns to every type the utility that this type obtains from selecting optimally from the menu at offer.

this type.[8]

As is well-known, the monopolist's interim utility is constrained by incentive compatibility to be convex in type. It is therefore not possible to implement every given distribution, that is, find an interim utility schedule that would validate the first order condition for this distribution. I show that a distribution is implementable if and only if its support is contained in the set of types where the Gateaux-derivative coincides with its own lower convex envelope. For example, if investment costs decrease in the riskiness of the distribution, then the Gateaux-derivative is concave (see Cerreia-Vioglio, 2017) so that it coincides with its lower convex envelope only on the boundaries of the type space, and thus, only distributions are implementable that put all mass on the smallest and largest cost type.

My characterization of implementability has the important implication that the interim utility that is required to induce the monopolist to choose an implementable distribution is pinned down—via the first order condition—by the marginal investment costs (in the form of the Gateaux-derivative). As is well-known, incentive compatibility, in turn, implies that the interim utility pins down the associated production schedule (by "revenue equivalence"). Therefore, the production schedule that can be implemented along with a given (implementable) distribution is fully determined by the marginal investment costs. This is a key difference to the case with an exogenous distribution where production distortions are determined by the hazard rate of the distribution.

The regulator's trade-off when implementing a distribution can be intuitively understood as trading off *virtual* production surplus against *virtual* investment costs. Virtual investment costs correspond to the expected interim utility the regulator has to provide to the monopolist to incentivize investment. The difference between virtual and actual investment costs amounts to agency costs due to moral hazard: in the "pure moral hazard" problem, when the monopolist's type becomes publicly known, the regulator would have to provide the same interim utility to induce the monopolist to select a given distribution because the interim utility is what determines her investment incentives. When there is, in addition, adverse selection, the regulator may need to distort the production level to provide the required interim utility so as to induce truth-telling. This results in a distorted, hence virtual, production surplus and the difference to the first-best production surplus amounts to agency costs due to adverse selection.

The agency cost perspective is useful to understand why the degree of investment cost convex-

[8] In particular, this means that the "first order approach" is always valid with flexible investments. This is a key tractability advantage over settings with parameterized investments where the validity of the first order approach is well known to be restrictive (see, e.g., Jewitt, 1988).

ity is the source of inefficiency implied by optimal contracting. A special case of convexity arises when costs are linear. Economically, linear investment costs mean that marginal investment costs are constant. I show that when costs are strictly convex, virtual costs exceed actual costs while they are the same for linear costs.[9] The reason is familiar from standard marginal benefit vs cost reasoning: given some contract, the monopolist invests "up to the margin" until marginal benefits equal marginal investment cost. In the present context, marginal benefits are constant (since expected utility is linear in probabilities), and when costs are strictly convex (resp. linear), marginal investment costs are strictly increasing (resp. constant). Thus, "below the margin", the monopolist accumulates strictly positive (resp. zero) marginal utility. In other words, when investment costs are strictly convex (resp. linear) the agency costs due to moral hazard are positive (resp. zero). Moreover, I show that the efficient investment distribution is implementable by a contract which displays efficient production levels. Therefore, when investment costs are linear, the efficient outcome can be implemented without any agency costs and is thus optimal for the regulator. By contrast, when investment costs are strictly convex, the gap between virtual and actual investment costs leads to distortions of some kind.

In the last part of the paper, I provide detailed insights into the distortions induced under optimal regulation by distinguishing investment cost functions according to their risk properties. The risk properties capture the costs of making the investment outcome more or less risky in the mean preserving spread sense. Investment costs that are decreasing in risk correspond to situations where "swinging for the fences" is cheaper than making safe, incremental cost reductions. For example, when cost-reducing investments are the result of R&D efforts, pursuing radical innovations might be less costly than fine-tuning existing technologies in young and dynamic industries but more costly in mature ones.

When investment costs are decreasing in risk, only two-type distributions that are supported on the most extreme cost types are implementable. Consequently, the regulator's problem simplifies and can be fully solved. I show that an optimal contract induces under-investment but may stipulate efficient production levels. Optimal contracts can generally be implemented as fixed-price-award-fee contracts where the firm is awarded an additional fee if it supplies the higher of two pre-specified production levels. Moreover, in the case that some investment is necessary for production to be efficient, a fixed-price contract (without award fee) is optimal.

When investment costs are increasing (or constant) in risk, my characterization of the set of implementable distributions implies that any distribution is implementable. Absent any con-

[9] This observation appears in a somewhat different context in Krähmer (2024) as well as in the "pure moral hazard" model of Georgiadis et al. (2024).

straints, the regulator's problem becomes intractable in general. To make progress, I additionally assume that investment costs are moment-based, that is, depend only on K-many (generalized) moments of the distribution. In this case, the regulator's problem can be written as a linear problem with K-many moment constraints. Extreme point arguments based on Winkler (1988) then imply that there is an optimal distribution in the class of discrete distributions with at most $K+1$ mass points, and it can be implemented by a menu consisting of at most $K+1$ menu items. In this sense, optimal contracts remain simple also in this case. While it is difficult to identify the direction of distortions in general, when investment costs depend on one moment only, production levels are downward distorted but the optimal and the efficient distribution cannot be stochastically ranked, in general.

Literature

My paper is related to various literatures. In the literature on principal agent problems with moral hazard and adverse selection, the setting I consider is analysed in Laffont and Martimort's text book (2002, Chapter 7.3.3) for the case with two possible types where the agent can invest in the probability to be the more efficient type at a cost. Despite its fundamental nature, I am not aware of work that addresses my research question in a more general parametric setting.[10]

A number of papers study (parametric) settings where, reversely to my paper, the agent first observes her type and then chooses effort. In such a setting, the interplay between ex post payment constraints and the need to provide effort incentives can lead to the optimality of simple pooling contracts consisting of few, or even a single remuneration schedule (see Gottlieb and Moreira, 2022, Ollier and Thomas, 2013, Martimort et al., 2025).[11] By contrast, simple contracts in my paper arise because simple type distributions emerge endogenously.[12]

Within the literature that studies flexible investments, I build on Georgiadis et al. (2024), who consider a flexible moral hazard problem, by adding a screening stage after the agent's investment. While any distribution is implementable by some wage contract when there is only moral hazard, I characterize how the presence of adverse selection restricts implementability, and delineate the agency costs this adds to the principal's problem. In Krähmer (2024), I consider a hold-up problem

[10]A large literature studies the provision of cost reduction incentives in the seminal framework of Laffont and Tirole (1986). In this literature, the firm's pre-investment costs are privately and post-investment costs are publicly known whereas in my model it is the other way round. Moreover, the agent's investment effort can, effectively, be perfectly monitored, whereas in my model, monitoring is impossible.

[11]Castro-Pires et al. (2024) identify conditions so that the screening and the effort provision problem can be decoupled in a setting with a risk averse agent.

[12]The timing I consider is also considered in the literature on mechanism design with ex ante investments where multiple agents can make productive investments before participating in the mechanism (see, e.g., Rogerson, 1992, Piccione and Tan, 1996, Arozamena and Cantillon, 2004, Krähmer and Strausz, 2007). In contrast to my paper, these papers consider parametric investment technologies and often take the mechanism, notably an auction, as given.

with flexible investments where, unlike in the current paper, the principal offers a contract after the agent has invested. As here, the convexity of the cost function allows the agent to extract a rent, but in the hold-up game, this enhances, while in the present paper, it decreases efficiency.[13]

My paper is also related to a literature that, instead of investment, studies information acquisition in principal agent models. The most closely related among these is Mensch and Ravid (2024) who extend the Mussa and Rosen (1978) model by a moral hazard stage where the agent acquires a flexible, costly signal about her type.[14] The key difference to flexible investment is that the agent's distribution (of posterior means) then needs to be Bayes consistent with a prior. Methodologically, this necessitates a duality-based approach very different from my approach based on Gateaux-differentiability. Mensch and Ravid (2024) find that production levels are always strictly downward distorted, whereas I identify various instances where production levels are efficient. Mensch and Ravid (2024) also find conditions under which optimal menus consist of only finitely many items, using extreme point arguments. Two other papers study screening with flexible information acquisition. In Mensch (2022), the principal has linear preferences and at most one unit is traded. Thereze (2024) studies the case when the type space is binary. All of the papers consider cost functions that are linear ("posterior-separable") and increasing in risk, while I can allow for more general ones, as I do not require Bayes consistency.

The paper is organized as follows. The next section introduces the regulation model with flexible investments. Section 3 characterizes the set of feasible outcomes. Section 4 derives and analyses the regulator's problem. Section 5 considers specific cost structures. Section 6 concludes. All proofs are in the appendix.

2 Model

I consider the classical regulation framework where a regulator (the principal) contracts with a firm (the agent) to produce a quantity/quality $x \geq 0$ of a service/good in exchange for a transfer $t \in \mathbb{R}$. I refer to x as "allocation". The agent has quasi-linear utility with constant marginal production costs $\theta \in \Theta = [\underline{\theta}, \overline{\theta}]$, $0 \leq \underline{\theta} < \overline{\theta}$. I refer to θ as the agent's type. The novelty of my model is that the distribution (cdf) F of the agent's type is not exogenous, but chosen by the

[13]Condorelli and Szentes (2020) study a hold-up problem with flexible investments where the seller can observe the investment choice. Other work that studies optimal contract design with flexible effort choice but specific cost functions include Diamond (1998), Hébert (2018), or Barron et al. (2020).

[14]Various papers study information acquisition in the standard screening model using parametric approaches. See, e.g., Crémer and Khalil (1992), Crémer et al. (1998a,b), Szalay (2009).

agent at cost $C(F)$. Thus, the agent's profit is

$$\pi = t - \theta x - C(F). \tag{1}$$

The principal seeks to maximize a weighted sum of (post-tax) "consumer surplus" $\beta(x) - t$ and the agent's profit with the latter receiving welfare weight $\alpha \leq 1$. The function β is increasing and strictly concave and twice differentiable. The principal's objective is thus[15]

$$\beta(x) - t + \alpha\pi = \beta(x) - \theta x - C(F) - (1-\alpha)\pi. \tag{2}$$

My analysis focuses on the case that both the agent's investment choice F and her type θ are her private information, that is, there is moral hazard and adverse selection. Moreover, I assume that the agent, as a firm, is protected by limited liability, that is, her post investment profit $t - \theta x$ must be non-negative.[16] A convenient way to capture this is to assume that the agent has an outside option of zero after observing her type.[17] Thus, the timing is as follows.

1. The principal commits to a contract specifying terms of trade x and t.
2. The agent covertly chooses F.
3. The agent privately observes the realization θ of F.
4. The agent accepts or rejects the contract.
   (a) If the agent accepts, the contract is enforced.
   (b) If the agent rejects, both parties receive their outside option of 0.

I next state the assumptions on the cost function:

1. C is continuous, monotone, and convex.[18]

[15] In the price regulation model of Baron and Myerson (1982), the agent is a monopolist who operates in a product market with inverse demand $P(\cdot)$, and the social value is the total surplus $\beta(x) = \int_0^x P(y)dy$. For $\alpha = 0$, the model corresponds to a "pure" procurement problem. It is straightforward to include social costs of public funds as in Laffont and Tirole (1986) by assuming that consumer surplus is $\beta(x) - \rho t$ where $\rho \geq 1$ captures the social cost of administering one unit of the transfer. The principal's objective is then $\beta(x) - \rho\theta x - \rho C(F) - (\rho - \alpha)\pi$. While the formal analysis is identical to the case with $\rho = 1$, the interpretation of some results changes.

[16] As explained below, the model is not interesting if the agent can sustain losses, since in this case the first-best outcome can be implemented.

[17] As is usual, by re-interpreting variables, the model applies equally well to a price discrimination context as in Mussa and Rosen (1982) where the principal is a monopolistic seller who commits to a menu of price-quality options, and the agent is a buyer who can invest to increase her marginal valuation for the product before deciding whether to choose an item from the menu or to abstain.

[18] Continuity refers to the weak topology.

2. C is smooth in the sense that C is Gateaux-differentiable with continuous and differentiable Gateaux derivative $c_F : \Theta \to \mathbb{R}$, that is, for $F, \tilde{F}$, we have

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \left[C(F + \epsilon(\tilde{F} - F)) - C(F) \right] = \int_{\Theta} c_F(\theta)\, d(\tilde{F} - F). \tag{3}$$

3. The cost of the "highest cost" distribution, which places mass 1 on the highest cost $\overline{\theta}$, referred to as F_0, is normalized to 0: $C(F_0) = 0$.

Continuity is a technical condition that ensures the existence of various maximizers below. Monotonicity means that smaller (stochastic) cost reductions are cheaper, that is, $C(F) \leq C(G)$ if F first order stochastically dominates G. Convexity captures increasing marginal costs of investment. For example, consider a weighted average of a "low-cost" and a "high-cost" distribution, that is, the latter first order stochastically dominates the former. Convexity then implies that marginally increasing the weight on the low-cost distribution gets more costly the higher the weight. Both, monotonicity and convexity are natural assumptions in the investment context considered here.

Smoothness captures a notion of differentiability which will make the analysis tractable. As is well-known, the Gateaux-derivative is a functional derivative that generalizes the notion of a partial derivative from functions of vectors to functions of functions. Economically, the Gateaux-derivative $c_F(\theta)$ evaluated at a type θ measures the marginal cost of increasing the probability mass assigned to this type given F. The final assumption is a normalization that ensures that "not investing" has no cost.

I will make use of the well-known fact that for smooth costs, monotonicity is characterized by monotonicity of the Gateaux-derivative, that is, monotonicity is equivalent to c_F being decreasing in θ for all F (see Cerreia-Vioglio et al., 2017).

Next, I discuss properties of cost functions that I do not impose throughout but will play important roles at various points in the paper. C is linear (in F) if there is a differentiable function c so that $C(F) = \int c(\theta) dF$. The Gateaux derivative $c_F = c$ does then not depend on F. Moreover, my normalization $C(F_0) = 0$ implies that $c(\overline{\theta}) = 0$. With linear costs the agent can generate a stochastic distribution F only through a "mixed strategy" that randomizes over deterministic outcomes θ, each costing $c(\theta)$, according to F, and the costs of doing so is "expected costs". (When costs C are convex, such a mixed strategy is more costly than choosing F directly.)

I shall sometimes distinguish cost functions according to their "risk properties". C is called decreasing (resp. increasing) in risk if $C(F) \leq C(G)$ (resp. $C(F) \geq C(G)$) whenever F is a mean-preserving spread of G. The risk properties capture the costs of allowing or avoiding dispersion in the investment outcome. For example, when marginal production costs θ are the result of investment in R&D, then investment costs that are decreasing in risk correspond to situations where "swinging for the fences" is cheaper than making safe, incremental production cost reductions.

Importantly, the risk properties are characterized by the shape of the Gateaux derivative c_F. This is well understood when costs are linear: in this case, when c is concave (resp. convex), then $C(F) = \int c\, dF$ is decreasing (resp. increasing) in risk. Analogously, for general cost functions C, being decreasing (resp. increasing) in risk is equivalent to c_F being concave (resp. convex) for all F (see Cerreia-Vioglio et al., 2017).

Notice that this implies that if c_F is affine for all F, then C is both decreasing and increasing in risk (in fact, C then only depends on the mean of F). For my purposes it will be convenient to define C to be strongly decreasing in risk if c_F is concave and non-affine for all F. Finally, note that convexity of C and the risk properties of C are distinct properties.

I conclude the section with discussing the "first-best" benchmark, where the principal can mandate F and x and observe θ. In this case, the optimal transfer compensates the agent for her production and investment costs, $t = \theta x + C(F)$, and the principal receives $\beta(x) - \theta x - C(F)$. Thus, given θ, the first-best allocation, $x^{FB}(\theta)$, maximizes the production surplus

$$S(x, \theta) = \beta(x) - \theta x. \tag{4}$$

I assume that $x^{FB}(\theta)$ is uniquely given by the first order condition $\beta'(x^{FB}(\theta)) - \theta = 0$. Let $S^{FB}(\theta) = S(x^{FB}(\theta), \theta)$. The first-best value of investment F is therefore

$$V^{FB}(F) = \int S^{FB}(\theta)\, dF - C(F). \tag{5}$$

I refer to a distribution F^{FB} that maximizes V^{FB} as a first-best distribution.[19] Standard arguments imply that the principal can attain the first-best outcome if, rather than at stage 3 in the timeline above, the agent had to accept or reject the contract immediately after stage 1 and was not protected by limited liability.[20]

[19]Compactness of the set of all cdfs and continuity of C implies that the problem $\max_F V^{FB}(F)$ has a solution.

[20]Indeed, since x^{FB} is decreasing, the transfer $t(\theta) = S^{FB}(\theta) - \theta x^{FB}(\theta) + k$ elicits the agent's type truthfully at stage 3, and by setting k appropriately the principal can extract the entire surplus at stage 1. This way, the agent effectively becomes the residual claimant and chooses first-best investment F^{FB}.

3 Implementability

Before I study optimal contracts, I characterize in this section the set of distributions that can be implemented by some contract. This set will constitute the feasible set for the principal's contracting problem.

In general, a contract consists of a message set and a message-contingent allocation-transfer assignment. By the revelation principle, it is without loss to restrict attention to direct, incentive compatible, and individually rational contracts where the message set is the entire set of valuations Θ. In what follows, instead of transfers, I work with the agent's interim (or indirect) utility function

$$U(\theta) = t(\theta) - \theta x(\theta), \tag{6}$$

and refer to $(x, U) : \Theta \to \mathbb{R}_+ \times \mathbb{R}$ as a contract. It is well-known that (x, U) is incentive compatible and individually rational ("IC and IR") if and only if

$$U \text{ is convex}, \quad U'(\theta) = -x(\theta) \text{ (whenever the derivative exists)}, \quad U(\overline{\theta}) \geq 0. \tag{7}$$

I say F is implementable if there is an IC and IR contract (x, U) that induces the agent to choose F:

$$F \in \arg\max_{F'} \int_\Theta U(\theta)\, dF'(\theta) - C(F'). \tag{8}$$

In this case, (x, U) is said to implement F. A combination (F, x, U) is feasible if (x, U) is IC and IR and implements F.

The next lemma characterizes feasible outcomes. To state it, let $\breve{c}_F$ be the lower convex envelope of c_F:

$$\breve{c}_F(\theta) = \sup\{g(\theta) \mid g(\tau) \leq c_F(\tau) \text{ for all } \tau \in \Theta, g \text{ convex }\}. \tag{9}$$

Lemma 1. *Let (x, U) be IC and IR.*

(i) (x, U) implements F if and only if there is $\lambda \in \mathbb{R}$ so that

$$U(\theta) = c_F(\theta) + \lambda \quad \forall \theta \in supp(F), \tag{10}$$

$$U(\theta) \leq c_F(\theta) + \lambda \quad \forall \theta. \tag{11}$$

(ii) If (x,U) implements F, then[21]

$$U \text{ is differentiable on } supp(F)\cap(\underline{\theta},\overline{\theta}) \text{ with } -x(\theta)=U'(\theta)=c_F'(\theta), \tag{12}$$

$$\text{If } \underline{\theta}\in supp(F): x(\underline{\theta})\geq -\breve{c}_F'(\underline{\theta}); \quad \text{if } \overline{\theta}\in supp(F): x(\overline{\theta})\leq -\breve{c}_F'(\overline{\theta}). \tag{13}$$

The significance of the lemma is that the joint presence of moral hazard and adverse selection severely restricts the degrees of freedom for a contract to implement a given F. Part (i) implies that the agent's interim utility (up to the constant λ) is pinned down by c_F on the support of F, and part (ii) implies that the allocation x is pinned down by c_F' on the support of F except on the boundary of Θ where the allocation is constrained by $\breve{c}_F'$.

More specifically, part (i) says that to incentivize the agent to choose a distribution F, the principal has to provide her with an interim utility that satisfies (10) and (11).[22] To understand the conditions, it is useful to consider its finite-dimensional analogue. Suppose that there are only finitely many possible types: $\Theta=\{\underline{\theta},\ldots,\theta,\ldots\overline{\theta}\}$. The agent then chooses a probability vector $f=(f_{\underline{\theta}},\ldots,f_\theta,\ldots,f_{\overline{\theta}})$ at cost $C(f)$ so as to maximize $\sum_{\theta\in\Theta}U(\theta)f_\theta-C(f)$ subject to the constraint that f be a probability vector: $\sum_{\theta\in\Theta}f_\theta=1$. The first order condition for f_θ to be optimal and positive—so that θ is in the support—is then that

$$U(\theta)=\frac{\partial C(f)}{\partial f_\theta}+\lambda, \tag{14}$$

where λ is the Lagrange multiplier. This mirrors condition (10) where $c_F(\theta)$ corresponds to the partial derivative of C with respect to f_θ evaluated at f. Condition (11) corresponds to the first order condition for f_θ to be optimal and equal to the corner solution zero.

While part (i) deals with moral hazard, part (ii) of the lemma describes additional restrictions due to adverse selection. Figure 1(a) illustrates the logic behind (12). Due to incentive compatibility, the agent's interim utility U is convex in θ. At the same time, (10) and (11) imply that U is located below $c_F+\lambda$. This implies that when $U(\theta)$ is equal to $c_F(\theta)+\lambda$ at an interior point θ, then U is squeezed in between $c_F+\lambda$ and its tangent. Thus, U and c_F have the same derivatives. Because incentive compatibility also implies that $U'(\theta)=-x(\theta)$, the allocation x is pinned down by $-c_F'$ on $supp(F)\cap(\underline{\theta},\overline{\theta})$.

Condition (13) follows from similar considerations, applied to boundary points where incen-

[21] Unless it leads to confusion, I write $g'(\underline{\theta})$ for the right derivative of a function g at $\underline{\theta}$, and likewise $g'(\overline{\theta})$ for the left derivative at $\overline{\theta}$.

[22] Part (i) are the necessary and sufficient first order conditions for the (concave) maximization problem (8) as established in Georgiadis et al. (2024).

tive constraints are only one-sided. Hence, the allocation on the boundaries is constrained only from below (at $\underline{\theta}$) or from above (at $\overline{\theta}$).

While the lemma characterizes when an IC and IR contract implements F, it leaves open whether and when such a contract exists. To address this question, a necessary condition is readily obtained. Let

$$\breve{\Theta}_F \equiv \{\theta \in \Theta \mid c_F(\theta) = \breve{c}_F(\theta)\} \tag{15}$$

be the set of types where c_F coincides with its lower convex envelope (see Figure 1(b)).[23]

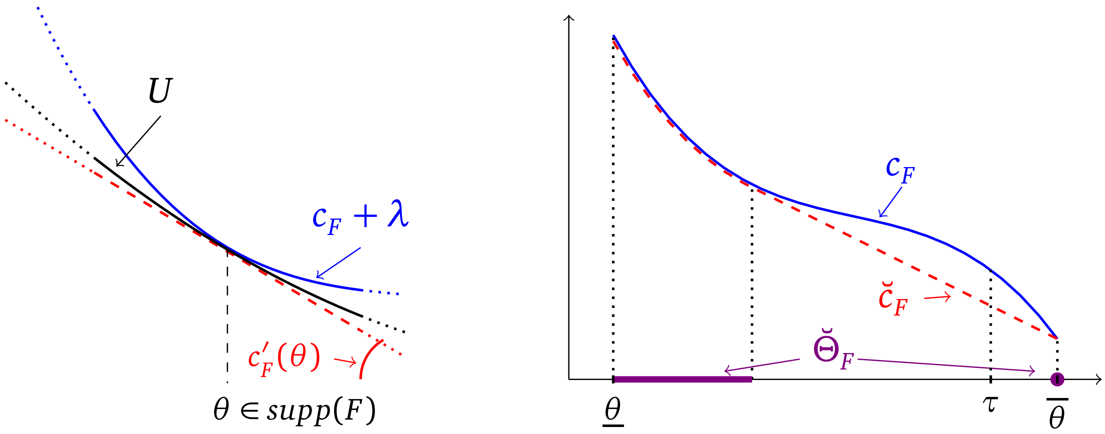

(a) $U'(\theta) = c_F'(\theta)$ (b) The set $\breve{\Theta}_F$

Figure 1: The figure illustrates condition (12) (left panel) and the set $\breve{\Theta}_F$ (right panel)

I now argue that if F is implementable, then

$$supp(F) \subseteq \breve{\Theta}_F. \tag{16}$$

Otherwise, it is impossible to motivate the agent to choose F and report truthfully. Indeed, the conditions (10) and (11) imply that U is located below $c_F + \lambda$. Because U is convex, U is therefore located below the lower convex envelope $\breve{c}_F + \lambda$. Now, consider Figure 1(b) and suppose, contrary to (16), that F had a point τ in its support that is outside of $\breve{\Theta}_F$. Clearly, $U(\tau)$ can then not be equal to $c_F(\tau) + \lambda$ and below $\breve{c}_F(\tau) + \lambda$ at the same .

As it turns out, (16) is also sufficient for F to be implementable. To show this, I next introduce a class of contracts that will implement F if (16) holds.

[23]Because c_F is continuous by assumption, $\breve{\Theta}_F$ contains at least the boundary points $\underline{\theta}$ and $\overline{\theta}$.

Definition 1. *A contract* $(\breve{x}, \breve{U})$ *is called* F*-canonical if*[24]

(i) $\breve{U}(\theta) = \breve{c}_F(\theta) + \breve{\lambda}$ *with* $\breve{\lambda} \geq -\min_{\theta \in \Theta} \breve{c}_F(\theta) = -c_F(\overline{\theta})$.

(ii) $\breve{x}(\theta) = -\breve{c}_F'(\theta)$, $\breve{x}(\underline{\theta}) \geq -\breve{c}_F'(\underline{\theta})$, $\breve{x}(\overline{\theta}) \leq -\breve{c}_F'(\overline{\theta})$.

The basic idea behind an F-canonical contract is to offer the lower convex envelope $\breve{c}_F$ as the utility schedule and its negative derivative $-\breve{c}_F'$ as the allocation schedule. Since $\breve{c}_F$ is convex, the resulting contract is IC. Moreover, the choice of $\breve{\lambda}$ ensures that the contract is IR:

Lemma 2. *An* F*-canonical contract is IC and IR.*

The next result shows that condition (16) is sufficient so that F can be implemented by an F-canonical contract. The reason is that given (16), an F-canonical contract satisfies the implementability conditions from Lemma 1 by construction. Together with the above-mentioned fact that (16) is necessary for F to be implementable, this characterizes the set of implementable distributions.

Theorem 1. *The following are equivalent:*

(i) F *is implementable.*

(ii) $supp(F) \subseteq \breve{\Theta}_F$.

(iii) Any F*-canonical contract implements* F.

Theorem 1 has the following immediate corollary which connects the risk properties of the cost function with the set of implementable distributions.

Corollary 1. *(i) If* C *is strongly decreasing in risk, then only distributions that are supported on the most extreme types* $\{\underline{\theta}, \overline{\theta}\}$ *can be implemented.*

(ii) If C *is increasing in risk, then all distributions can be implemented.*

Recall that C is strongly decreasing in risk if c_F is concave and non-affine for all F, and therefore $\breve{\Theta}_F = \{\underline{\theta}, \overline{\theta}\}$. Thus, the joint restrictions imposed by moral hazard and adverse selection are so severe that only two-point distributions with mass on the lowest and highest possible cost types are implementable. On the other hand, when costs are increasing in risk, then c_F is convex for all F and thus $\Theta_F = \Theta$. Thus, there are no restrictions on what can be implemented.

[24] As is well-known, since c_F is differentiable (by assumption), so is $\breve{c}_F$.

Theorem 1 restricts attention to F-canonical contracts. But are there other contracts that implement F, possibly with different welfare properties? The next result shows that it is indeed sufficient to focus on F-canonical contracts, because they "span" the entire set of allocations and payoff combinations that can be attained when implementing F.

Proposition 1. *Let F be implementable. For any IC and IR contract that implements F, there is an "equivalent" F-canonical contract that implements F, that is, the allocations as well as the payoffs for the principal and the agent are the same under both contracts for all $\theta \in supp(F)$.*

The reason why it is sufficient to focus on F-canonical contracts is that by Lemma 1, implementability of F pins down the allocation and the agent's utility by c_F on the support of F. The remaining degrees of freedom in implementing F can be replicated by some F-canonical contract.

So far, I have focussed on direct contracts. With a view on applications, I conclude this section with a brief discussion on how to implement a distribution with a "simple" indirect contract. As is well known, when the agent's type distribution is exogenous, a direct contract can be implemented by offering the agent an (indirect) menu of allocations and transfers to choose from where the cardinality of the menu can be chosen to be at most the cardinality of the support of the distribution. The next lemma shows that an analogous taxation principle holds in my context when the distribution is endogenous.

Lemma 3. *Let F be implementable. For any IC and IR contract that implements F, there is an "equivalent" menu of allocation-transfer pairs $\{(x(\theta), t(\theta)) \mid \theta \in supp(F)\}$ that implements F and whose cardinality is equal to the cardinality of the support of F, that is, the allocations as well as the payoffs for the principal and the agent are the same under the original contract and the menu for all $\theta \in supp(F)$.*

As with the standard taxation principle, the basic idea is to replace the direct contract with the menu of allocation-transfer pairs that the direct contract implements on the support of F. A potential issue arises when $supp(F)$ does not coincide with the entire set Θ. In this case, the menu might give rise to different terms of trade for types "off" the support of F than the original contract which is defined on the entire set Θ. In principle, this could affect the agent's investment incentives when the menu is offered. The proof shows that this is, however, not the case.

As mentioned in Corollary 1, if C is strongly decreasing in risk, then only distributions that are supported on the boundary of Θ can be implemented. By Lemma 3, all that can be implemented can then also be implemented by a simple menu consisting of two options only. In other words, simple contracts are optimal because effectively no other contracts are feasible.

4 Optimality

In this section, I derive the principal's problem of designing an optimal contract. I identify the agency costs stemming from moral hazard and adverse selection and shed light on the distortions they give rise to. The central insight is that the inefficiencies imposed by an optimal contract are driven by the convexity of the investment cost function. When the investment cost function is linear, an optimal contract will be shown to be first-best.

4.1 Principal's problem

The principal's problem is to choose a feasible combination (F, x, U) that maximizes her profit. I first ask what, for a given implementable F, is the optimal contract (x, U) that implements F. Proposition 1 implies that there is an optimal contract in the class of F-canonical contracts. An F-canonical contract leaves three degrees of freedom: The constant $\breve{\lambda}$ and the boundary allocations $\breve{x}(\underline{\theta})$ and $\breve{x}(\overline{\theta})$.

It is optimal to choose $\breve{\lambda}$ as small as possible so that the individual rationality constraint is binding, that is, $\breve{\lambda} = -c_F(\overline{\theta})$. With this choice, because $c_F(\theta) = \breve{c}_F(\theta)$ on the support of an implementable F, the agent's interim utility is $U(\theta) = c_F(\theta) - c_F(\overline{\theta})$ on the support of F. Thus, her expected interim utility from the contract is[25]

$$\tilde{U}(F) \equiv \int U(\theta)\, dF = \int c_F(\theta)\, d(F - F_0). \tag{17}$$

As to the choice of the boundary allocations, since the principal's (interim) payoff in terms of the agent's interim utility is

$$S(\breve{x}(\theta), \theta) - \alpha C(F) - (1-\alpha) U(\theta) \tag{18}$$

by (2), the principal optimally chooses the allocations $\breve{x}(\underline{\theta})$ and $\breve{x}(\overline{\theta})$ so as to maximize the production surplus $S(x, \underline{\theta})$ and $S(x, \overline{\theta})$ subject to the constraints in Definition 1, (ii). For example, it is optimal to choose $\breve{x}(\underline{\theta})$ as closely to $x^{FB}(\underline{\theta})$ as possible, that is, $\breve{x}(\underline{\theta}) = \max\{x^{FB}(\underline{\theta}), -\breve{c}'_F(\underline{\theta})\}$. The next lemma summarizes these observations.

Lemma 4. *Let F be implementable. An optimal contract that implements F is the F-canonical*

[25] Recall that F_0 places probability 1 on $\overline{\theta}$. Thus $\int U(\theta)\, dF = \int c_F(\theta) - c_F(\overline{\theta})\, dF = \int c_F(\theta)\, d(F - F_0)$.

contract with

$$\breve{\lambda} = -c_F(\overline{\theta}), \quad \breve{x}(\underline{\theta}) = \underline{x}_F \equiv \max\{x^{FB}(\underline{\theta}), -\breve{c}_F'(\underline{\theta})\}, \quad \breve{x}(\overline{\theta}) = \overline{x}_F \equiv \min\{x^{FB}(\overline{\theta}), -\breve{c}_F'(\overline{\theta})\}. \quad (19)$$

The agent's expected interim utility from the contract is $\tilde{U}(F)$, and the principal's profit is

$$V(F) = \int S(x_F(\theta), \theta)\, dF - \alpha C(F) - (1-\alpha)\tilde{U}(F), \quad (20)$$

where[26]

$$x_F(\theta) \quad = \quad \begin{cases} \underline{x}_F & if \quad \theta = \underline{\theta} \\ -c_F'(\theta) & if \quad \theta \in (\underline{\theta}, \overline{\theta}) \\ \overline{x}_F & if \quad \theta = \overline{\theta} \end{cases} . \quad (21)$$

In light of Lemma 4 and Theorem 1, the principal's problem can be written as follows:

$$\mathscr{P}: \qquad \max_F V(F) \quad s.t. \quad supp(F) \subseteq \breve{\Theta}_F. \quad (22)$$

I refer to a cdf as optimal if it is a solution to this problem. I next discuss the principal's trade-offs.[27]

The objective (20) can be interpreted as a *virtual* value of investment: compared to the social value of investment, $\int S^{FB}(\theta)\ dF - C(F)$, the virtual value differs in two respects: first, the principal receives the *virtual* production surplus $S(x_F(\theta), \theta)$ when θ realizes which is smaller than the first-best production surplus because the allocation $x_F(\theta)$ is, in general, distorted. Second, instead of the investment cost $C(F)$, the principal faces the *virtual* investment cost

$$\tilde{C}(F) \equiv \alpha C(F) + (1-\alpha)\tilde{U}(F). \quad (23)$$

As I now argue, the difference between the first-best surplus and the virtual surplus captures the agency costs due to adverse selection whereas the difference between actual and virtual costs captures the agency costs due to moral hazard. To identify the agency costs due to moral hazard, the next lemma describes the outcome of the "pure moral hazard" setting where the agent's type becomes publicly known.

[26]To understand the shape of the allocation $x_F(\theta)$, recall (a) that on the interior of Θ the allocation under an F-canonical contract is given by $\breve{x}(\theta) = -\breve{c}_F'(\theta)$ and (b) that $\breve{c}_F'(\theta) = c_F'(\theta)$ on the support of an implementable F.

[27]I discuss the existence of a solution to $\mathscr{P}$ in appendix B.

Lemma 5 (Pure moral hazard, Georgiadis et al. 2024)**.** *Suppose the agent's type is publicly observable (but otherwise nothing is changed). Then any F is implementable. If the principal implements F optimally, the agent's expected interim utility is $\tilde{U}(F)$, and the principal's profit is*

$$V^{MH}(F) = \int S^{FB}(\theta)\, dF - \tilde{C}(F). \tag{24}$$

When θ is publicly observable, the principal faces no incentive compatibility constraint to elicit θ, and her problem is to choose a non-negative interim utility schedule $U(\theta)$ and an allocation $x(\theta)$ subject to the moral hazard constraints (10) and (11) only. For any F, the constraints (10) and (11) can be satisfied by setting $U(\theta) = c_F(\theta) + \lambda$ on the support of F, and $U(\theta) \leq c_F(\theta) + \lambda$ elsewhere. Moreover, it is optimal for the principal to set $\lambda = -c_F(\overline{\theta})$ so that individual rationality binds. Thus, the agent's expected interim utility is $\int U(\theta)\, dF = \tilde{U}(F)$. Finally, since F pins down the agent's expected utility, the principal optimally chooses $x(\theta) = x^{FB}(\theta)$ efficiently so as to maximize his share of the production surplus. Thus, when the principal implements F optimally, he receives $V^{MH}(F)$.

Because the discrepancy between the first-best and the pure moral hazard problem is the difference between virtual and true investment costs, this difference captures the agency costs due to moral hazard. To identify the agency costs due to adverse selection, I compare the cost of implementing a given distribution F when there is only moral hazard with when there is also adverse selection. First, while F is implementable with pure moral hazard, F may not be implementable when there is also adverse selection, as it may violate the implementability condition $supp(F) \subseteq \breve{\Theta}_F$. In this case, the presence of adverse selection makes implementing F infinitely more costly.

Second, suppose that F is implementable with adverse selection and moral hazard. Recall that to implement F in the pure moral hazard case, the principal offers the agent the interim utility $U(\theta) = c_F(\theta) + \lambda$ along with the first-best allocation x^{FB} on the support of F. Now, to implement F with both moral hazard and adverse selection, the principal has to offer the agent the same interim utility (up to a constant) on the support of F, simply because the moral hazard constraints (10) and (11) have not changed. However, it might not be feasible to offer this interim utility and the first-best allocation x^{FB} simultaneously because this might violate incentive compatibility. In this case, to provide the required interim utility in an incentive compatible way, the principal needs to distort the allocation x. The agency cost due to adverse selection is thus the reduction in production surplus caused by this distortion.[28]

4.2 Welfare

I now argue that any inefficiency that arises in my setting can, ultimately, be attributed to the convexity of investment costs. To do so, I first analyze the agent's ex ante profit

$$\Pi(F) = \tilde{U}(F) - C(F) = \int c_F(\theta) d(F - F_0) - C(F). \tag{25}$$

Lemma 6. *Let F be implementable. If the principal implements F optimally, then:*

(i) If C is linear, then $\Pi(F) = 0$.

(ii) If C is strictly convex and $C(F) \neq 0$, then $\Pi(F) > 0$.

The intuition for the lemma becomes apparent when one considers a simplified investment problem where there are only two outcomes, "success" and "failure": The agent chooses a one-dimensional probability $f \in [0,1]$ of success at a smooth, convex cost $C(f)$ with $C(0) = 0$. The principal pays the agent a prize $u \geq 0$ in case of success and a prize of 0 in case of failure. Since the agent maximizes $fu - C(f)$, the principal can implement any probability $\hat{f}$ by choosing the prize so as to satisfy the agent's first order condition $u = C'(\hat{f})$. Plugging the first order condition back into the agent's objective yields the agent's ex ante utility

$$\Pi(\hat{f}) = \hat{f} C'(\hat{f}) - C(\hat{f}). \tag{26}$$

When C is strictly convex (and $\hat{f} > 0$), this expression is strictly positive. The reason is simple: the agent increases her investment f up to the point where marginal costs $C'(\hat{f})$ are equal to marginal benefits u. But since marginal costs are strictly increasing, and marginal benefits are constant, the agent accumulates strictly positive marginal profit for every marginal unit of investment up to $\hat{f}$. When C is linear, on the other hand, marginal benefits are equal to marginal costs for all f up to the optimum, resulting in zero overall profit for the agent. The significance of Lemma 6 is that this basic logic carries over unchanged to the setting with flexible investments.[29]

[28] A different way to measure the agency cost due to adverse selection is to compare the first-best value to the principal's value in the "pure adverse selection" setting where the investment distribution is contractible. This necessitates to solve for an optimal contract for any given distribution as in Hellwig (2010).

[29] Lemma 6 mirrors Proposition 3 in Krähmer (2024).

Lemma 6 has immediate welfare implications. First, the agent's ex ante profit might be increasing in costs (see also Krähmer, 2024). Indeed, consider a cost function $C(F) = \int \ell(\theta)\, dF + \kappa C_0(F)$ that is the sum of a linear part and a scaled strictly convex part C_0. Then Π increases when κ is (locally) increased from 0. Second, part (i) of Lemma 6 readily implies that when investment costs are linear, then the first-best is optimal, provided it is implementable. The next theorem shows this is indeed true. Therefore, whenever costs are linear, there are no agency costs and optimal contracts are efficient.[30]

Theorem 2. *(i) Any first-best distribution F^{FB} is implementable. Moreover, the associated optimal F^{FB}-canonical contract displays first-best allocations $\breve{x}(\theta) = x^{FB}(\theta)$ on the support of F^{FB}.*

(ii) Let C be linear. Then any first-best distribution F^{FB} is optimal, and the principal extracts the full first-best value $\max_F V^{FB}(F)$*, while the agent gets 0.*

To see part (i), note that to implement the first-best outcome, the principal can offer a contract that specifies the first-best allocation $x^{FB}(\theta)$ and interim utility $U(\theta)$ which equals the first-best surplus $S^{FB}(\theta)$ but for a constant. As is well-known, this contract is IC.[31] The contract is also IR if the constant is sufficiently large, and, clearly, induces the agent to choose a first-best distribution. Even though this contract might not be F^{FB}-canonical, it can be replicated by an F^{FB}-canonical contract by Proposition 1.

Part (ii) then follows from Lemma 6: If the principal implements a first-best distribution (along with the first-best allocation) with an optimal contract, the agent's ex ante profit is zero, and thus the principal's receives the full first-best value. Clearly, the principal cannot do better. Hence, any first-best distribution is optimal if investment costs are linear.

Note that since the first-best problem has a solution, the previous proposition in particular implies that a solution to the principal's problem exists if investment costs are linear. In Appendix B, I present sufficient conditions that ensure existence of a solution in the general case. For the cost structures considered in the next section, existence will follow from elementary arguments.

[30]Note that Lemma 6 applies equally to the setting with only moral hazard (see Lemma 5). In this case, any cdf can be implemented and thus, as already noted by Georgiadis et al. (2024), part (i) readily implies that the first-best is optimal in the pure moral hazard setting when costs are linear.

[31]More precisely, because the first-best allocation is decreasing and $\frac{d}{d\theta}S^{FB}(\theta) = -x^{FB}(\theta)$, it follows that U is convex and $U' = -x^{FB}$ which is equivalent to (x, U) being IC.

5 Investment cost structures

The goal of this section is to obtain more detailed insights into the kind of distributions and distortions that arise endogenously under an optimal contract. To do so, I shall distinguish ivnestment cost functions depending on whether they are strongly decreasing in risk or increasing in risk.

As seen in Corollary 1 above, when investment costs are strongly decreasing in risk, only two-point distributions can be implemented, and optimal distributions are therefore necessarily simple and can be implemented by simple contracts consisting of two items (see Lemma 3). In fact, I shall identify conditions under which a fixed-price contract (a single item menu) is optimal.

The case when investment costs are increasing in risk is harder, because then any cdf can be implemented. To make progress, I shall focus on the class of "moment-based" cost functions and show that optimal distributions remain "simple" in the sense that they display only finitely many types. Accordingly, optimal regulatory contracts are simple.

Moreover, I shall provide insights into the allocative and investment distortions that arise in both cases.

5.1 Investment costs that are strongly decreasing in risk

In this section, I consider the case that C is strongly decreasing in risk, that is, c_F is concave and non-affine for all F. As seen above, then only distributions are implementable that are supported on the boundaries. I denote by T_f the distribution that places mass f on $\underline{\theta}$ and mass $1-f$ on $\overline{\theta}$, $f \in [0,1]$. Thus, the principal's problem reduces to the uni-dimensional problem of choosing f.[32] Let $\gamma(f) = C(T_f)$ be the cost of choosing T_f. Since C is convex, so is γ. As shown in the appendix, the boundary allocations (19) that are optimal to implement T_f are

$$\underline{x}_f = \max\left\{x^{FB}(\underline{\theta}), \frac{\gamma'(f)}{\Delta\theta}\right\}, \quad \overline{x}_f = \min\left\{x^{FB}(\overline{\theta}), \frac{\gamma'(f)}{\Delta\theta}\right\}, \tag{27}$$

where $\Delta\theta = \overline{\theta} - \underline{\theta}$. Moreover, the virtual costs of T_f can be computed to

$$\tilde{\gamma}(f) \equiv \tilde{C}(T_f) = \alpha\gamma(f) + (1-\alpha)f\gamma'(f). \tag{28}$$

Theorem 3. *Let C be strongly decreasing in risk.*

[32] Existence of a solution thus follows straightforwardly.

(i) *A first-best distribution is a distribution* $T_{f^{FB}}$ *where* f^{FB} *is a solution to the problem*

$$\max_{f\in[0,1]} f S^{FB}(\underline{\theta}) + (1-f)S^{FB}(\overline{\theta}) - \gamma(f). \tag{29}$$

(ii) *An optimal distribution is a distribution* T_{f^*}. *If* $f^* > 0$, *then the "low cost type"* $\underline{\theta}$ *receives the first-best allocation* $x^{FB}(\underline{\theta})$, *and the "high cost type"* $\overline{\theta}$ *receives the allocation* $\overline{x}_{f^*}$. *Moreover,* f^* *is a solution to the problem*

$$\max_{f\in[0,1]} f S^{FB}(\underline{\theta}) + (1-f)S(\overline{x}_f, \overline{\theta}) - \tilde{\gamma}(f). \tag{30}$$

(iii) *An optimal distribution displays under-investment:* $f^* \leq f^{FB}$.

To shed light on part (i), consider the first-best problem $\max_F \int S^{FB}(\theta)\, dF - C(F)$. Because $S^{FB}(\theta)$ is convex in θ, and C is strongly decreasing in risk, the objective is increasing in risk. A first-best distribution is therefore maximally risky and puts all mass on the smallest and largest cost type and is thus a two-point distribution T_f. Problem (29) is then simply the first-best problem in this class of two-point distributions.

Part (ii) says that an optimal contract features "no distortion at the top" and, possibly, a "downward distortion at the bottom". The reason why there is no distortion at the top is that otherwise it would be an upward distortion given the constraint that $x(\underline{\theta}) = \max\{x^{FB}(\underline{\theta}), \gamma'(f)/\Delta\theta\}$. In light of (27), and since the first-best allocation is decreasing, the allocation $x(\overline{\theta})$ for the high cost type would then be efficient. Intuitively, it would then be an improvement to reduce the mass on $\underline{\theta}$ and so reduce the weight on the inefficiency and save on virtual investment costs.

Finally, part (iii), that the optimal investment is smaller than the first-best investment is not entirely straightforward because there are countervailing effects. On the one hand, marginal virtual investment costs are larger than marginal investment costs.[33] On the other hand, if the allocation is downward distorted at $\overline{\theta}$, that is, $\overline{x}_f < x^{FB}(\overline{\theta})$, then the effect on the virtual production surplus of a marginal increase of f is not only to put more mass on the more efficient type $\underline{\theta}$ but also to make the allocation $\overline{x}_f$ for the high-cost type more efficient. In contrast, the former effect is the only effect a marginal increase of f has on the first-best production surplus. As it turns out, the investment cost effect is the dominant force, implying under-investment.

It is illuminating to compare the solution of the principal's problem to the solution of the pure

[33]Indeed, due to convexity of γ: $\frac{d}{df} f\gamma'(f) = \gamma'(f) + f\gamma''(f) \geq \gamma'(f)$.

moral hazard problem (Lemma 5)

$$\max_F \int S^{FB}(\theta)\,dF - \tilde{C}(F). \tag{31}$$

In contrast to the first-best problem and problem $\mathscr{P}$, the solution to this problem is not necessarily supported on the boundaries $\underline{\theta}$ and $\overline{\theta}$. The reason is that there is no implementability constraint, and virtual investment costs $\tilde{C}(F)$ are not necessarily decreasing in risk. However, for the special case that $\tilde{C}(F)$ is decreasing in risk, the solution is supported on the boundaries for the same reason the first-best is (S^{FB} is convex, so the objective is increasing in risk overall).

Hence, suppose that $\tilde{C}(F)$ is decreasing in risk.[34] In this case, (31) boils down to choosing a probability f that maximizes

$$f S^{FB}(\underline{\theta}) + (1-f)S^{FB}(\overline{\theta}) - \tilde{\gamma}(f). \tag{32}$$

Denote by f^m a solution to (32). Figure 2 plots the marginal virtual production surplus of problem (30), denoted $E\tilde{S}'(f)$,[35] as well as the marginal production surplus in the first-best problem and the pure moral hazard problem (33), given by $\Delta S^{FB} = S^{FB}(\underline{\theta}) - S^{FB}(\overline{\theta})$, as a function of f. The curve $E\tilde{S}'(f)$ has a kink at the value f_0 where the allocation $\overline{x}_f$ for the high cost type becomes equal to the first-best allocation so that, from that point onward, the virtual and the first-best surplus coincide. The intersection of the marginal virtual and the marginal first-best surplus, respectively, with the marginal virtual investment cost curve $\tilde{\gamma}'$ delivers the solutions f^* and f^m, respectively.

The left panel depicts the case where the marginal virtual investment cost curve $\tilde{\gamma}'$ intersects the marginal virtual surplus curve at a point smaller than f_0, and therefore $f^m < f^* < f_0$. The allocation $\overline{x}_{f^*}$ of the high cost type $\overline{\theta}$ is therefore downward distorted. The right panel depicts the opposite case where the marginal virtual investment cost curve intersects the marginal virtual surplus curve at a point larger than f_0, and therefore $f_0 < f^m = f^*$. The allocation $\overline{x}_{f^*}$ of the high cost type is therefore efficient, and the additional adverse selection problem causes no additional agency costs. The next lemma summarizes.[36]

[34] An example is a cost function $C(F) = \Gamma(\int \varphi(\theta)\,dF)$ with a convex function Γ and a concave function φ.

[35] Formally,

$$E\tilde{S}'(f) = \frac{d}{df}\left[f S^{FB}(\underline{\theta}) + (1-f)S(\overline{x}_f, \overline{\theta})\right]. \tag{33}$$

[36] Formally, f_0 is defined by $\Delta\theta x^{FB}(\overline{\theta}) = \gamma'(f_0)$, and efficient allocations obtain if and only if $f_0 \le f^*$ in which case

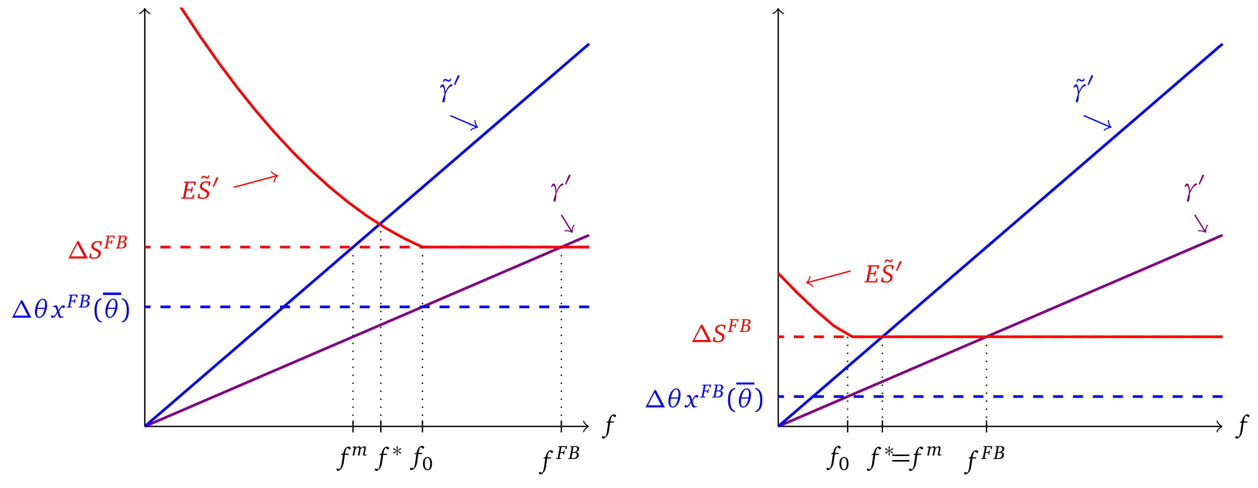

(a) $\Delta\theta x^{FB}(\overline{\theta}) > \gamma'(f^m)$ (b) $\Delta\theta x^{FB}(\overline{\theta}) < \gamma'(f^m)$

Figure 2: The figure plots the marginal virtual surplus ($E\tilde{S}'$) and the marginal surplus of the pure moral hazard and the first-best problems (ΔS^{FB}) as well as marginal costs (γ') and marginal virtual costs ($\tilde{\gamma}'$). In the left panel, the solutions to the pure moral hazard problem and the problem with adverse selection on top differ; in the right panel, they coincide. The plot is for the model specification: $\underline{\theta} = 1$, $\overline{\theta} = 2$, $\beta(x) = \sigma x - 1/2 \cdot x^2$, $C(F) = (\int \varphi \, dF)^2$, φ strictly concave with $\varphi(\underline{\theta}) = 1$, $\varphi(\overline{\theta}) = 0$, and $\sigma = 3$ (left), $\sigma = 9/4$ (right)

Lemma 7. *Let C be strongly decreasing in risk, and let $\tilde{C}$ be decreasing in risk. Then the solution to the principal's problem (30) displays efficient allocations ($\underline{x}_{f^*} = x^{FB}(\underline{\theta})$, $\overline{x}_{f^*} = x^{FB}(\overline{\theta})$) if and only if*

$$\Delta\theta x^{FB}(\overline{\theta}) \leq \gamma'(f^m). \tag{34}$$

Therefore, the case in the right panel of Figure 2 occurs if and only if condition (34) holds. In this case, the interim utility schedule that is necessary to motivate the agent to choose f^m in the pure moral hazard case, is already "steep enough" to render it incentive compatible when combined with the first-best allocation x^{FB}. Therefore, the pure moral hazard outcome can be implemented also when there is adverse selection.

More precisely, under pure moral hazard, to induce the agent to choose f^m, the principal optimally offers the agent utility $U^m = \gamma'(f^m)$ if the low cost type $\underline{\theta}$ realizes and utility zero if the high cost type $\overline{\theta}$ realizes. Thus the low cost type receives a rent of $\gamma'(f^m)$ relative to the high cost type. If condition (34) is violated, this rent is not sufficient for the low cost type to not mimic the hight cost type if the agent's type is her private information. In fact, there is then no way

$f^* = f^m$. Monotonicity of γ' implies that $f_0 \leq f^m$ if and only if (34).

to device a contract that makes the pure moral hazard outcome incentive compatible. To attain incentive compatibility, the principal can distort the allocation for the high cost type downwards and thus reduce the low cost type's utility from mimicking the high cost type, or offer the low cost type a higher utility and thus increase investment incentives. At an optimum, the principal chooses a combination of both leading to a downward distortion for the high cost type and higher investment.[37]

I conclude this section with a discussion on how to implement an optimal contract. Note first that, in general, a two-item menu $\{(x(\underline{\theta}), t(\underline{\theta})), (x(\overline{\theta}), t(\overline{\theta}))\}$ can be interpreted as a fixed-price-award-fee contract where the principal pays the "base fee" $t(\overline{\theta})$ when the agent supplies the "default" allocation $x(\overline{\theta})$ and pays the "award fee" $t(\underline{\theta})$ if the agent supplies the larger allocation $x(\underline{\theta})$.[38]

In the special case that the two-item menu specifies allocation $x(\overline{\theta}) = 0$ for the high cost type, the menu amounts to a fixed-price contract which offers the agent to produce $x(\underline{\theta})$ for the fixed price $t(\underline{\theta})$. In this case, the high cost type "shuts down". Notice that the optimal contract from Theorem 3 thus corresponds to a fixed-price contract if zero production is efficient at the high cost type: $x^{FB}(\overline{\theta}) = 0$. In other words, if some investment is needed for production to be efficient, a fixed-price contract is optimal.

Corollary 2. *Let costs be strongly decreasing in risk and suppose $x^{FB}(\overline{\theta}) = 0$. Then the optimal contract can be implemented by a fixed-price contract that pays the agent the price $\tilde{\gamma}(f^m)$ if she produces $x^{FB}(\underline{\theta})$.*[39]

5.2 Investment costs that are increasing in risk and moment-based

I now consider the case that investment costs are moment-based and have a convex Gateaux-derivative. C is moment-based if it only depends on finitely many generalized moments of F. Formally, let $\varphi = (\varphi_1, \ldots, \varphi_K)$ be a vector of K (integrable) functions $\varphi_k : \mathbb{R} \to \mathbb{R}$ and define the

[37] Note also that the agent's ex ante utility $\Pi = \tilde{U}(F) - C(F) = f\gamma'(f) - \gamma(f)$ is increasing in f. Thus, the agent is weakly better of when there is adverse selection on top of moral hazard.

[38] FPAF contracts are frequently used in infrastructure and defense procurement. In practice, the contractor is paid an award fee if she attains certain objective pre-determined performance criteria. These can include time to completion, and various reliability and performance targets (e.g. fuel consumption in a military vehicle).

[39] If $x^{FB}(\overline{\theta}) = 0$, then condition (34) holds, and as explained in the text, transfers are such that the high cost type obtains 0 utility and the low cost type obtains $\gamma'(f^m)$: $\overline{t} = \overline{\theta} x^{FB}(\overline{\theta}) = 0$, $\underline{t} = \gamma'(f^m) + \underline{\theta} x^{FB}(\overline{\theta}) = \gamma'(f^m)$.

K-dimensional moment vector of F by

$$\Phi_F = (\Phi_{1,F}, \ldots, \Phi_{K,F}), \quad \Phi_{k,F} = \int \varphi_k(\theta) dF. \tag{35}$$

I call C moment-based if there is a convex, increasing, and differentiable function $\Gamma : \mathbb{R}^K \to \mathbb{R}$ so that $C(F) = \Gamma(\Phi_F)$. In this case, the Gateaux derivative is the dot-product of the gradient Γ' and $\varphi(\theta)$, that is, $c_F(\theta) = \Gamma'(\Phi_F) \cdot \varphi(\theta)$. In this section I assume that

$$\varphi_k \text{ is decreasing, differentiable, and convex for all } k. \tag{36}$$

This implies that the Gateaux derivative is decreasing (ensuring monotonicity of C) and convex (ensuring that C is increasing in risk).

The main result of this section says that also in this setting, optimal contracts remain simple.

Theorem 4. *Let C be moment-based with (36). Then there is a solution F^* to the principal's problem $\mathscr{P}$ that has at most $K+1$ points in its support. In particular, there is an optimal menu with at most $K+1$ items that implements the solution (by Lemma 3).*

The reason behind the result is that when investment costs are moment-based, the principal's problem becomes a linear problem (in F) that has a simple solution. To see this, note that also virtual investment costs are moment-based and have the form

$$\tilde{C}(F) = \tilde{\Gamma}(\Phi_F) \equiv \alpha \Gamma(\Phi_F) + (1-\alpha)\Gamma'(\Phi_F) \cdot (\Phi_F - \varphi(\overline{\theta})). \tag{37}$$

Moreover, the virtual production surplus $S(x_F(\theta), \theta)$ depends only on the moment vector because c_F' and hence x_F depend on F only through the moment vector. With the notation $\tilde{S}(\Phi_F, \theta) = S(x_F(\theta), \theta)$, the principal's problem can thus be written as

$$\max_F \int \tilde{S}(\Phi_F, \theta)\, dF - \tilde{\Gamma}(\Phi_F). \tag{38}$$

Note that this is now an unconstrained problem because if investment costs are increasing in risk, any F is implementable by Corollary 1. To see that this problem has a "simple" solution, let $\Phi_k \in [\varphi_k(\overline{\theta}), \varphi_k(\underline{\theta})]$, and consider the constrained problem where each moment is kept fix: $\Phi_{k,F} = \Phi_k$ for $1 \leq k \leq K$. Inserting the constraint into the objective in (38), the constrained problem becomes:

$$\max_F \int \tilde{S}(\Phi_1, \dots, \Phi_K, \theta)\, dF - \tilde{\Gamma}(\Phi_1, \dots, \Phi_K) \quad s.t. \quad \int \varphi_k\, dF = \Phi_k, 1 \le k \le K. \tag{39}$$

This is a linear problem (in F) with K linear constraints. Thus, there is a solution which is an extreme point of the constraint set.[40] It is well-known (Winkler, 1988) that the set of extreme points is the set of discrete distributions with at most $K+1$ points in their support.

5.2.1 Mean-based investment costs

In the special case that investment costs are mean-based, the solution can be fully characterized. C is mean-based if $C(F) = \Gamma(M_F)$, $M_F = \int \overline{\theta} - \theta \; dF$. It turns out that with mean-based costs, both a first-best and an optimal distribution is in the class of distributions T_f that are supported on $\{\underline{\theta}, \overline{\theta}\}$ very much like in the case treated in Theorem 3. Once this is established, Theorem 3 carries over verbatim to the case with mean-based investment costs. The only difference is that now, the function γ is defined as the cost of T_f when investment costs are mean-based:

$$\gamma(f) = \Gamma(M_{T_f}) = \Gamma(f\,\Delta\theta). \tag{40}$$

Theorem 5. *Let C be mean-based.*

(i) A first-best distribution is a distribution $T_{f^{FB}}$. Moreover, there is an optimal distribution T_{f^*}.

(ii) Points (i) to (iii) from Theorem 3 carry over unchanged with γ given by (40).

(iii) Corollary 2 carries over unchanged.

With mean-based investment costs, the first-best objective $\int S^{FB}(\theta)\, dF - \Gamma(M_F)$ is increasing in risk. ($S^{FB}(\theta)$ is convex in θ and Γ is "constant" in risk.) A first-best distribution therefore puts all mass on the smallest and largest cost type and is thus a two-point distribution T_f.

[40] Existence of a solution to (39) is always guaranteed. To see this, observe that the virtual production surplus $\tilde{S}(\Phi, \theta)$ is upper semicontinuous in θ: It is continuous on $(\underline{\theta}, \overline{\theta})$ (because $x_F(\theta) = -c_F'(\theta)$ is continuous), and may have a downward jump at $\underline{\theta}$ and an upward jump at $\overline{\theta}$. Thus, the objective $\int \tilde{S}(\Phi, \theta)\; dF - \tilde{\Gamma}(\Phi)$ is upper semicontinuous in F with respect to weak convergence. Moreover, the constraint set is compact. The reason is that φ_k is continuous for all k by assumption. Thus, if a sequence (F_n) from the constraint set converges weakly to F, then $\int \varphi_k\, dF = \lim_n \int \varphi_k\, dF_n = \Phi_k$ and thus F is in the constraint set. As the objective is upper semicontinuous and the constraint set is compact, a solution exists by Weierstraß' extreme value theorem. That there is a solution that is an extreme point follows from Bauer's maximum principle for upper semicontinuous and convex (especially linear) objectives. See, e.g., Ok, 2007, Chapter J.5.

As to the principal's problem, unlike when costs are strongly decreasing risk, any distribution is now implementable. The basic intuition for why the optimal distribution still puts all mass on the smallest and largest cost typ is that the allocation $x_F(\theta)$ is now equal to $c_F'(\theta) = -\Gamma'(M_F)$ and thus constant on $(\underline{\theta}, \overline{\theta})$. Accordingly, the virtual production surplus $\tilde{S}(M_F, \theta) = \beta(\Gamma'(M_F)) - \theta\Gamma'(M_F)$ is convex in θ.[41] Therefore, the principal's objective $\int \tilde{S}(M_F, \theta)\, dF - \tilde{\Gamma}(M_F)$ is increasing in risk. Thus, the optimal distribution puts all mass on the smallest and largest cost type.

5.2.2 Distortions

In this section, I shed light on the distortions implied by an optimal contract. For tractability reasons, I restrict attention to the case where investment costs depend on one moment ($K = 1$).

I first argue that investment distortions do not have a clear cut direction. When investment cost are mean-based (Theorem 5), there is under-investment in the sense that an optimal distribution is first and second order stochastically dominated by the first-best distribution. I now show by example that this is not generally the case.

In the example, the first-best distribution is degenerate and places all mass on an interior point $\theta^{FB} \in (\underline{\theta}, \overline{\theta})$. In contrast, an optimal distribution puts positive mass on the lowest possible cost type $\underline{\theta}$. Thus, the optimal distribution is not first or second order stochastically dominated by the first-best distribution.

Lemma 8. *Let $\overline{\theta} = \underline{\theta} + 1$, $\alpha = 0$, $\beta(x) = \overline{\theta}x - 1/2 \cdot x^2$, and*

$$C(F) = \kappa \Phi_F^2, \quad \Phi_F = \int \varphi \, dF, \quad \varphi(\theta) = \left(\overline{\theta} - \theta\right)^{5/2}, \quad \kappa > 0. \tag{41}$$

There are $0 < \kappa_0 < \kappa_1$ so that for $\kappa \in (\kappa_0, \kappa_1)$, the (unique) first-best distribution places mass 1 on an interior point $\theta^{FB} \in (\underline{\theta}, \overline{\theta})$, while an optimal distribution places positive mass on the lowest possible cost type $\underline{\theta}$.[42]

The comparison of the optimal distribution and the first-best is driven by two forces. On the one hand, the principal's marginal virtual investment costs are larger than marginal investment costs. This force weakens the investment incentives relative to the first-best. On the other hand, the marginal virtual surplus might be larger than the marginal surplus. As discussed after Theorem 3, a change in F not only affects the expected value of the surplus but also affects the

[41] In fact, $S(M_F, \theta)$ is linear on $(\underline{\theta}, \overline{\theta})$ and may have a downward jump at $\underline{\theta}$ and an upward jump at $\overline{\theta}$ (because the first-best allocation may be implemented on the boundaries).

[42] The values can be explicitly calculated: $\kappa_0 = 1/(2\rho) = .2$ and $\kappa_1 = (100/81)^2 \cdot (5/27) \approx .82$. Moreover, $\theta^{FB} = \sigma - (\frac{1}{2\rho\kappa})^{\frac{1}{2(\rho-1)}} \in (\underline{\theta}, \overline{\theta})$.

allocation directly and possibly moves it closer to the first-best allocation. This force strengthens the investment incentives relative to the first-best and is the driver behind the example.

Next, I turn to the question how the allocative distortions look like under an optimal contract. Under Theorem 3 and 5, the implemented allocations $x_{F^*}(\theta)$ are never distorted upwards on the support of F^*, and sometimes there are no distortions at all on the support of F^*. The next proposition generalizes this insight in the sense that when $K = 1$, an optimal contract does not impose upward distortions.

Proposition 2. *Let costs be given by*

$$C(F) = \Gamma(\Phi_F), \quad \Phi_F = \int \varphi(\theta)\, dF \tag{42}$$

with Γ and φ twice differentiable, strictly decreasing, and strictly convex. Then there is an optimal contract with no upward distortions.

To illustrate the intuition, recall that when investment costs depend on one moment only, then there is an optimal distribution with two points θ_1 and θ_2 in its support. Assume that both points are interior so that the allocations are given by $x_1 = -\Gamma'(\Phi)\varphi'(\theta_1)$ and $x_2 = -\Gamma'(\Phi)\varphi'(\theta_2)$. (If one of the support points is not interior, similar, yet more tedious arguments apply.) Now imagine that both x_1 and x_2 were upward distorted and consider the effect of decreasing θ_2 marginally. Since this decreases production costs, there is a direct positive effect on the virtual surplus. On the other hand, decreasing θ_2 decreases the moment Φ and thus increases virtual investment costs. Moreover, both allocations x_1 and x_2 decrease, and, since they are upward distorted, they become more efficient, increasing the virtual surplus. As it turns out, the positive effects dominate the negative effect, making the modification profitable.

If one allocation, say x_1, were upward distorted, and the other, x_2, were downward distorted, then the principal could decrease θ_1 and increase θ_2 so as to keep the moment, Φ, constant. This modification leaves virtual investment costs unchanged, decreases x_1 and increases x_2, hence making the allocations more efficient, and thus increasing the virtual surplus.

The latter argument becomes more tricky when costs depend on more than one moment. In this case, an optimal distribution has three points in its support, and if some allocations are upward distorted and some downward, the question is whether a modification can be found that leaves both moments constant and makes each allocation more efficient. I leave this for future research.

6 Conclusion

In this paper, I study an optimal regulation model with moral hazard and adverse selection where the monopolist can flexibly invest to reduce her production costs. I show that inefficiencies are driven by the degree of convexity of the investment cost function. Moreover, I show that optimal contracts and the distribution of agent types that they induce endogenously are often simple. This provides a novel rationale why real world regulatory schemes are often simple in contrast to predictions based on models with exogenous type distributions. For specific cost structures, I show that optimal contracts may induce investment distortions only but display efficient production levels.

My approach can be extended to various other applications. A case in point is optimal auctions with flexible investments. While my implementability result can be extended to characterize feasible outcomes in terms of interim allocations ("reduced form auctions"), the key question is then which interim allocations can be obtained from feasible mechanisms (as in, for example, Gershkov et al. 2021), taking into account that agents are now engaged in an investment game. Moreover, my approach is applicable to other single agent design problems where incentive compatibility can be characterized in terms of a convex indirect utility function. This includes optimal delegation problems without money (Kleiner, 2022) or multi-dimensional screening problems. Finally, it is an interesting question what happens when the agent first receives private information and then chooses a flexible investment.

A Proofs

Proof of Lemma 1 (i) (x, U) implements F if and only if (8). By Proposition 1 in Georgiadis et al. (2024), (8) is equivalent to

$$\int U(\theta) - c_F(\theta)\, dF(\theta) \geq \int U(\theta) - c_F(\theta)\, dG(\theta) \tag{43}$$

for all cdfs G. Let λ be equal to the left hand side. Then the inequality is equivalent to

$$\int U(\theta) - c_F(\theta) - \lambda\, dF(\theta) = 0 \quad and \quad \int U(\theta) - c_F(\theta) - \lambda\, dG(\theta) \leq 0 \quad \forall G. \tag{44}$$

Because the right inequality holds for all G, it is equivalent to $U(\theta) - c_F(\theta) - \lambda \leq 0$ for all θ. Since U is continuous by IC, and since c_F is continuous by assumption, this implies that the left equality is equivalent to $U(\theta)-c_F(\theta)-\lambda=0$ for all $\theta \in supp(F)$. This shows that (8) is equivalent to (10) and (11).

(ii) Let (x,U) be a direct IC and IR contract that implements F. As to (12), IC implies that U is convex and therefore right and left differentiable with

$$\partial^- U(\theta) \le \partial^+ U(\theta). \tag{45}$$

Moreover, by part(i), (x,U) satisfies (10) and (11). Together with the differentiability of c_F, this implies for $\theta \in supp(F) \cap (\underline{\theta}, \overline{\theta})$:

$$\partial^- U(\theta) = \lim_{\theta^- \uparrow \theta} \frac{U(\theta^-)-U(\theta)}{\theta^- - \theta} \ge \lim_{\theta^- \uparrow \theta} \frac{c_F(\theta^-)-c_F(\theta)}{\theta^- - \theta} = c_F'(\theta); \tag{46}$$

$$\partial^+ U(\theta) = \lim_{\theta^+ \downarrow \theta} \frac{U(\theta^+)-U(\theta)}{\theta^+ - \theta} \le \lim_{\theta^+ \downarrow \theta} \frac{c_F(\theta^+)-c_F(\theta)}{\theta^+ - \theta} = c_F'(\theta). \tag{47}$$

Putting the inequalities (45) to (47) together implies that $\partial^- U(\theta) = \partial^+ U(\theta) = c_F'(\theta)$. In particular, U is differentiable at θ with $U'(\theta) = c_F'(\theta)$.

Finally, since IC implies that at all points of differentiability of U, one has $U'(\theta) = -x(\theta)$. Thus, $x(\theta) = -c_F'(\theta)$, proving (12).

As to (13), let $\underline{\theta} \in supp(F)$. Then $U(\underline{\theta}) = c_F(\underline{\theta}) + \lambda$ by (10). Moreover, since U is convex and $U \le c_F + \lambda$ by (11), we have that $U \le \breve{c}_F + \lambda$. Since $\breve{c}_F \le c_F$, this implies that $U(\underline{\theta}) = \breve{c}_F(\underline{\theta}) + \lambda$. Therefore, $U \le c_F + \lambda$ implies that $\partial^+ U(\underline{\theta}) \le \partial^+ \breve{c}_F(\underline{\theta})$. Further, by IC, we have $-x(\underline{\theta}) \le \partial^+ U(\underline{\theta})$. Thus, $x(\underline{\theta}) \ge -\partial^+ \breve{c}_F(\underline{\theta})$, as claimed. The claim for $\underline{\theta} \in supp(F)$ follows from analogous arguments. This completes the proof. QED

Proof of Lemma 2 Since $\breve{c}_F$ is convex, it is absolutely continuous. Thus, "payoff equivalence" holds, since by construction, $\breve{U}(\theta)-\breve{U}(\overline{\theta}) = \int_\theta^{\overline{\theta}} -\breve{c}_F'(t)\,dt = \int_\theta^{\overline{\theta}} \breve{x}(t)\,dt$. Moreover, $\breve{x}$ is decreasing since c_F, and thus $\breve{c}_F$ is decreasing. It is well known that payoff equivalence and monotonicity implies IC. Moreover, $\breve{U}(\overline{\theta}) \ge 0$ by construction, and $\breve{U}$ is decreasing. Thus IR follows. QED

Proof of Theorem 1 (i) $\Rightarrow$ (ii): Let F be implementable. Let (x,U) be an IC and IR contract that implements F. Thus, U is convex. Moreover, $U \le c_F + \lambda$ by (11) in Lemma 1. Thus,

$$U(\theta) \le \breve{c}_F(\theta) + \lambda \quad \forall \theta. \tag{48}$$

Therefore, the fact that $\breve{c}_F \le c_F$ and (10) imply that $c_F(\theta)+\lambda = \breve{c}_F(\theta)+\lambda$ for all $\theta \in supp(F)$. Thus, $supp(F) \subseteq \breve{\Theta}_F$, as desired.

(ii) $\Rightarrow$ (iii): Let $supp(F) \subseteq \breve{\Theta}_F$. Let (x,U) be F-canonical. It is straightforward to verify that

(x, U) satisfies (10) and (11). Thus, (x, U) implements F by part (i) of Lemma 1.

(ii) $\Rightarrow$ (iii): trivial. QED

Proof of Proposition 1 Let (x, U) implement F with the associated λ from part (i) of Lemma 1. Let $(\breve{x}, \breve{U})$ be the F-canonical contract with $\breve{\lambda} = \lambda$ and

$$\breve{x}(\underline{\theta}) = x(\underline{\theta}), \quad \breve{x}(\overline{\theta}) = x(\overline{\theta}). \tag{49}$$

By part (i) of Lemma 1, under (x, U), the agent obtains for every $\theta \in supp(F)$ the utility

$$U(\theta) = c_F(\theta) + \lambda = \breve{c}_F(\theta) + \lambda, \tag{50}$$

where the second equality follows, since $supp(F) \subseteq \breve{\Theta}_F$ by Theorem 1. Under $(\breve{x}, \breve{U})$, we have $\breve{U}(\theta) = \breve{c}_F(\theta) + \lambda$ by definition of an F-canonical contract and since $\breve{\lambda} = \lambda$. Thus, the two contracts are payoff-equivalent for the agent.

Turning to the principal, note that both contracts induce the same allocation $x = \breve{x}$ on $supp(F)$ by part (ii) of Lemma 1, the definition of an F-canonical contract, and (49). Thus, since the principal's interim payoff in terms of the agent's interim payoff is $S(x(\theta), \theta) - \alpha C(F) - (1-\alpha)U(\theta)$ by (2), the two contracts are payoff-equivalent for the principal. QED.

Proof of Lemma 3 Let F be implementable and $(\hat{x}, \hat{U})$ be an IC and IR contract that implements F. Define for all $\theta \in supp(F)$ the transfer-allocation pair

$$x(\theta) = \hat{x}(\theta), \quad t(\theta) = \theta\hat{x}(\theta) + \hat{U}(\theta). \tag{51}$$

Consider the menu $M = \{(x(\theta), t(\theta)) \mid \theta \in supp(F)\}$. It follows from a standard taxation principle argument that if the agent chooses F and type $\theta \in supp(F)$ realizes, the agent selects $(x(\theta), t(\theta))$ from M. Thus, given F, the same allocations and payoffs are implemented under M as under the contract $(\hat{x}, \hat{U})$ on the support of F.

What remains to be shown is that it is indeed optimal for the agent to choose F when M is offered. This is trivial if the support of F is the entire set Θ, because then M and $(\hat{x}, \hat{U})$ implement the same allocations and payoffs for all $\theta \in \Theta$. However, when $supp(F)$ is a strict subset of Θ, there can be types $\theta' \notin supp(F)$ who get a different allocation and payoff under M than under $(\hat{x}, \hat{U})$, and this, in principle, could induce the agent to choose a distribution different from F.

To show that this is not the case, I show that the indirect utility induced by M satisfies the sufficient conditions (10) and (11) to implement F in Lemma 1. To see this, let

$$U_M(\theta') = \max\{\max_{\theta \in supp(F)} t(\theta) - \theta' x(\theta), 0\} \tag{52}$$

be the utility for agent type $\theta' \in \Theta$ when choosing from the menu M. (The first "max" operator accounts for the possibility that some types might prefer to reject the menu.)

To see that (10) holds, note that, by construction, type $\theta \in supp(F)$ chooses $(x(\theta), t(\theta))$ from the menu because the original mechanism $(\hat{x}, \hat{U})$ from which the menu is derived is IC. Thus,

$$U_M(\theta) = t(\theta) - \theta x(\theta) = \theta\hat{x}(\theta) + \hat{U}(\theta) - \theta\hat{x}(\theta) - \theta\hat{x}(\theta) = \hat{U}(\theta) = c_F(\theta) + \lambda, \tag{53}$$

where the last equality follows because $(\hat{x}, \hat{U})$ implements F and thus satisfies (10) by assumption.

As to (11), observe that for all $\theta' \in \Theta$:

$$\begin{aligned} U_M(\theta') &= \max\{\max_{\theta \in supp(F)} t(\theta) - \theta' x(\theta), 0\} && (54)\\ &= \max\{\max_{\theta \in supp(F)} \theta\hat{x}(\theta) + \hat{U}(\theta) - \theta'\hat{x}(\theta), 0\} && (55)\\ &\leq \max\{\max_{\theta \in \Theta} \theta\hat{x}(\theta) + \hat{U}(\theta) - \theta'\hat{x}(\theta), 0\} && (56)\\ &= \max\{\hat{U}(\theta'), 0\} && (57)\\ &= \hat{U}(\theta'). && (58) \end{aligned}$$

Here, the first two lines follow by definition of (x, t) in (51), the third line because $supp(F) \subseteq \Theta$, the fourth line because $(\hat{x}, \hat{U})$ is IC, and the last line because $(\hat{x}, \hat{U})$ is IR. Hence, U_M satisfies (11), because $\hat{U}$ satisfies (11) by assumption. QED.

Proof of Lemma 4 In the text. QED

Proof of Lemma 5 In the text. QED

Proof of Lemma 6 Part (i) is immediate from the definition of linear costs. To see part (ii), define for $\tau \in [0, 1]$:

$$F^\tau = \tau F + (1 - \tau)F_0, \qquad \phi(\tau) = C(F^\tau). \tag{59}$$

Below, I show that

$$\phi'(\tau) = \int c_{F^\tau}(\theta)\, d(F - F_0). \tag{60}$$

Moreover, because C is convex, so is ϕ, and hence $\phi'(\tau) \le \phi'(1)$. Thus, since $C(F_0) = 0$,

$$C(F) = C(F) - C(F_0) = \phi(1) - \phi(0) = \int_0^1 \phi'(\tau)\, d\tau \le \phi'(1) = \tilde{U}(F), \tag{61}$$

which is (ii). It remains to show (60).[43] Note that for $h > 0$:

$$F^{\tau+h} = F^\tau + h(F - F_0) = F^\tau + \frac{h}{1-\tau}(F - F^\tau). \tag{62}$$

Thus, by definition of the Gateaux derivative:

$$\begin{aligned}
\phi'(\tau) &= \lim_{h\downarrow 0} \frac{1}{h}\left[C(F^{\tau+h}) - C(F^\tau)\right] && (63)\\
&= \lim_{h\downarrow 0} \frac{1}{h}\left[C\left(F^\tau + \frac{h}{1-\tau}(F - F^\tau)\right) - C(F^\tau)\right] && (64)\\
&= \frac{1}{1-\tau} \lim_{h\downarrow 0} \frac{1}{\frac{h}{1-\tau}}\left[C\left(F^\tau + \frac{h}{1-\tau}(F - F^\tau)\right) - C(F^\tau)\right] && (65)\\
&= \frac{1}{1-\tau}\int c_{F^\tau}(\theta)\, d(F - F^\tau) && (66)\\
&= \int c_{F^\tau}(\theta)\, d(F - F_0), && (67)
\end{aligned}$$

and this completes the proof. QED

Proof of Theorem 2 The argument is given in the text. QED

Proof of Theorem 3 I first derive (27). To see this, note that since c_{T_f} is concave, its lower convex envelope is simply the line that connects $c_{T_f}(\underline{\theta})$ and $c_{T_f}(\overline{\theta})$ and has constant slope

$$\check{c}'_{T_f}(\theta) = \frac{c_{T_f}(\underline{\theta}) - c_{T_f}(\overline{\theta})}{\Delta\theta} = \frac{\gamma'(f)}{\Delta\theta}. \tag{68}$$

The second equality follows from the fact that $T_{f+\epsilon} = T_f + \epsilon(T_1 - T_0)$, which implies

$$\begin{aligned}
\gamma'(f) = \frac{dC(T_f)}{df} &= \lim_{\epsilon\to 0} \frac{1}{\epsilon}\left[C(T_f + \epsilon(T_1 - T_0)) - C(T_f)\right] && (69)\\
&= \int c_{T_f}(\theta)\, d(T_1 - T_0) = c_{T_f}(\underline{\theta}) - c_{T_f}(\overline{\theta}). && (70)
\end{aligned}$$

Inserting (68) into (19) yields (27).

[43] A similar argument is used to prove equation (5.8) in Huber (1981).

To see expression (28) for the virtual costs of T_f, recall from (17) and (23) that

$$\tilde{\gamma}(f) = \tilde{C}(T_f) = \alpha C(T_f) + (1-\alpha)\int c_{T_f}(\theta)\, d(T_f - F_0) \quad (71)$$

$$= \alpha\gamma(f) + (1-\alpha)\left[f c_{T_f}(\underline{\theta}) + (1-f)c_{T_f}(\overline{\theta}) - c_{T_f}(\overline{\theta})\right] \quad (72)$$

$$= \alpha\gamma(f) + (1-\alpha) f\gamma'(f), \quad (73)$$

as desired.

I now prove (i). By Georgiadis et al. (2024), a first-best distribution is characterized by the condition

$$supp(F^{FB}) \subseteq \arg\max_{\theta}\left(S^{FB}(\theta) - c_{F^{FB}}(\theta)\right). \quad (74)$$

Since c_F is concave and non-affine and S^{FB} is convex, $S^{FB} - c_{F^{FB}}$ can attain a maximum only on the boundary points, and thus $supp(F^{FB}) \subseteq \{\underline{\theta}, \overline{\theta}\}$. The first-best problem therefore reduces to (29), and this establishes (i).

As to (ii). With (28) the principal's problem (22) becomes

$$\max_{f\in[0,1]} fS(\underline{x}_f, \underline{\theta}) + (1-f)S(\overline{x}_f, \overline{\theta}) - \alpha\gamma(f) - (1-\alpha)f\gamma'(f). \quad (75)$$

I now argue that if $f > 0$ is an optimum, we have $\underline{x}_f = x^{FB}(\underline{\theta})$. Indeed, if the contrary was true, then (27) and the fact that x^{FB} decreases would imply that $\underline{x}_f = \gamma'(f)/\Delta\theta$ and $\overline{x}_f = x^{FB}(\overline{\theta})$. Thus, the principal's expected utility would be

$$V(f) = fS\left(\frac{\gamma'(f)}{\Delta\theta}, \underline{\theta}\right) + (1-f)S^{FB}(\overline{\theta}) - f\gamma'(f) - \alpha[\gamma(f) - f\gamma'(f)] \quad (76)$$

$$= f\left[S\left(\frac{\gamma'(f)}{\Delta\theta}, \overline{\theta}\right) - S^{FB}(\overline{\theta})\right] + S^{FB}(\overline{\theta}) - \alpha[\gamma(f) - f\gamma'(f)], \quad (77)$$

where the equality uses the fact that $S(x,\underline{\theta}) = \beta(x) - \underline{\theta}x - = \beta(x) - \overline{\theta}x + x\Delta\theta = S(x,\overline{\theta}) + x\Delta\theta$. The derivative is

$$V'(f) = \left[S\left(\frac{\gamma'(f)}{\Delta\theta}, \overline{\theta}\right) - S^{FB}(\overline{\theta})\right] + f\frac{\partial S\left(\frac{\gamma'(f)}{\Delta\theta}, \overline{\theta}\right)}{\partial x}\frac{\gamma''(f)}{\Delta\theta} - \alpha f\gamma''(f). \quad (78)$$

I now argue that this expression is negative so that the principal would benefit from setting $f = 0$, a contradiction to the optimality of $f > 0$. Indeed, because $\gamma'(f)/\Delta\theta > x^{FB}(\overline{\theta})$ by assumption, both the term in the square brackets and $\partial S\left(\frac{\gamma'(f)}{\Delta\theta}, \overline{\theta}\right)/\partial x$ are strictly negative since the total surplus $S(x, \overline{\theta})$ is concave in x and maximized at $x^{FB}(\overline{\theta})$. Since $\gamma'' > 0$, this implies that $V'(f) \leq 0$, as desired. Finally, having established that $\overline{x}_f = x^{FB}(\overline{\theta})$ at an optimum, (75) becomes (30), as desired.

As to (iii). Note first that by definition of the first-best, we have $dS^{FB}(\theta)/d\theta = -x^{FB}(\theta)$. Thus, since S^{FB} is convex, it follows:

$$S^{FB}(\underline{\theta}) - S^{FB}(\overline{\theta}) \geq x^{FB}(\overline{\theta})\Delta\theta. \tag{79}$$

Let $f_0 \in [0, 1]$ be the smallest solution to

$$x^{FB}(\overline{\theta})\Delta\theta = \gamma'(f), \tag{80}$$

whenever a solution exists. If no solution exists and if $x^{FB}(\overline{\theta})\Delta\theta < \gamma'(f)$ for all f, let $f_0 = 0$, and if $x^{FB}(\overline{\theta})\Delta\theta > \gamma'(f)$ for all f, let $f_0 = 1$.

The first-order condition for the first-best problem (29) together with (79) implies

$$f_0 \leq f^{FB}. \tag{81}$$

I now show the claim that $f^* \leq f^{FB}$. Indeed, if $f^* \leq f_0$, the claim is immediate from the previous inequality. Suppose next that $f^* > f_0$. By definition of f_0 and because γ is convex and thus γ' is increasing, it follows that $x^{FB}(\overline{\theta}) \leq \gamma'(f^*)/\Delta\theta$, and thus $\overline{x}_{f^*} = x^{FB}(\overline{\theta})$. By (30), f^* is thus a solution to

$$\max_{f \geq f_0} \; f S^{FB}(\underline{\theta}) + (1-f)S^{FB}(\overline{\theta}) - \alpha\gamma(f) - (1-\alpha) f \gamma'(f). \tag{82}$$

This problem has the same marginal benefits, $S^{FB}(\overline{\theta}) - S^{FB}(\underline{\theta})$, from increasing f as the first-best problem (29) but higher marginal costs, because

$$\frac{d}{df} f\gamma'(f) = \gamma'(f) + f\gamma''(f) \geq \gamma'(f), \tag{83}$$

where the inequality follows from the convexity of γ. Therefore, $f^* \leq f^{FB}$, and this completes the proof. QED

Proof of Theorem 4: In the text. QED

Proof of Theorem 5 As to (i). To see that there is a first-best distribution with support in $\{\underline{\theta}, \overline{\theta}\}$, consider the first-best problem subject to the constraint that the mean is kept fix:

$$\max_F \int S^{FB}(\theta)\, dF - \Gamma(M_F) \quad s.t. \quad M_F = M. \tag{84}$$

Note that $S^{FB}(\theta)$ is convex in θ. Thus, the (unique) distribution that has mean M and whose support is in $\{\underline{\theta}, \overline{\theta}\}$ maximizes the objective.

To see that there is an optimal distribution with support in $\{\underline{\theta}, \overline{\theta}\}$, consider the principal's problem subject to the constraint that the mean is kept fix:

$$\max_F \int \tilde{S}(M_F, \theta)\, dF - \tilde{\Gamma}(M_F) \quad s.t. \quad M_F = M. \tag{85}$$

As explained in the main text, $\tilde{S}(M_F, \theta)$ is convex in θ. Thus, the (unique) distribution that has mean M and whose support is in $\{\underline{\theta}, \overline{\theta}\}$ maximizes the objective.

The proofs of part (ii) and (iii) are identical to the proofs of Theorem 3 and Lemma 7. QED

Proof of Lemma 8 Note first that $x^{FB}(\theta) = \overline{\theta} - \theta$, and $S^{FB}(\theta) = (\overline{\theta} - \theta)^2)/2$. The Gateaux derivative is $c_F(\theta) = 2\kappa\Phi_F\varphi(\theta)$, and (since $\alpha = 0$), virtual costs are $\tilde{C}(F) = 2\kappa\Phi_F^2$.

I proceed in various steps: (i) I first derive the first-best distribution. By Georgiadis et al. (2024), the first-best distribution is characterized by

$$supp(F^{FB}) \subseteq \arg\max_{\theta\in[\underline{\theta},\overline{\theta}]} S^{FB}(\theta) - c_{F^{FB}}(\theta) \tag{86}$$

$$\Longleftrightarrow \qquad supp(F^{FB}) \subseteq \arg\max_{\theta\in[\underline{\theta},\overline{\theta}]} \frac{(\overline{\theta}-\theta)^2}{2} - 2\kappa\Phi_{F^{FB}}(\overline{\theta}-\theta)^{5/2}. \tag{87}$$

Since the objective function is concave, there is a unique maximizer θ^{FB}. If interior, the maximizer is given by the first order condition

$$(\overline{\theta} - \theta^{FB}) - 5\kappa\Phi_{F^{FB}}(\overline{\theta} - \theta^{FB})^{3/2} = 0. \tag{88}$$

Since the maximizer is unique, only degenerate distributions can be a solution, and, in particular $\Phi_{F^{FB}} = (\overline{\theta} - \theta^{FB})^{5/2}$. Thus, by the first order condition: $\theta^{FB} = \overline{\theta} - \left(\frac{1}{5\kappa}\right)^{\frac{1}{3}}$. Because $\underline{\theta} = \overline{\theta} - 1$, it follows that θ^{FB} is interior if and only if

$$\kappa > \frac{1}{5} = \kappa_0. \tag{89}$$

(ii) Next, I show that an optimal distribution is either degenerate or has a mass point at $\underline{\theta}$. Indeed, by Theorem 4, there is an optimal distribution with at most two points θ_1, θ_2, $\theta_1 < \theta_2$, in its support. I first show that there is no solution that has two points in its support with $\theta_1, \theta_2 \in (\underline{\theta}, \overline{\theta}]$. Indeed, I can infer from (19) that

$$\overline{x}_F = 0 = -c_F'(\overline{\theta}). \tag{90}$$

Therefore, since virtual costs are $\tilde{\Gamma}(\Phi) = 2\kappa\Phi^2$, the principal's profit from a two-point distribution with support points $\theta_1, \theta_2 \in (\underline{\theta}, \overline{\theta}]$ and $f = Pr(\theta_1)$ is

$$V(f, \theta_1, \theta_2) \;=\; fS(-2\kappa\Phi\varphi'(\theta_1), \theta_1) + (1-f)S(-2\kappa\Phi\varphi'(\theta_2), \theta_2) - 2\kappa\Phi^2. \tag{91}$$

I now show that the principal's profit can be improved by increasing θ_1 and decreasing θ_2 slightly while keeping f and Φ fixed. Indeed, let $\theta_2(\theta_1)$ be such that $f\varphi(\theta_1) + (1-f)\varphi(\theta_2(\theta_1)) = \Phi$. Hence

$$\theta_2'(\theta_1) = -\frac{f\varphi'(\theta_1)}{(1-f)\varphi'(\theta_2)}. \tag{92}$$

Denoting $x_i = -\kappa\Phi\varphi'(\theta_i)$, the derivative of V with respect to θ_1 when Φ is kept fix is

$$\begin{aligned}
\frac{dV(f, \theta_1, \theta_2(\theta_1))}{d\theta_1} &= f\left[\frac{\partial S(x_1, \theta_1)}{\partial x}(-2\kappa\Phi\varphi''(\theta_1)) + \frac{\partial S(x_1, \theta_1)}{\partial \theta}\right] \\
&\quad + (1-f)\left[\frac{\partial S(x_2, \theta_2)}{\partial x}(-2\kappa\Phi\varphi''(\theta_2))\theta_2'(\theta_1) + \frac{\partial S(x_2, \theta_2)}{\partial \theta}\theta_2'(\theta_1)\right] &(93)\\
&= f\left[\frac{\partial S(x_1, \theta_1)}{\partial x}(-2\kappa\Phi\varphi''(\theta_1)) + \frac{\partial S(x_1, \theta_1)}{\partial \theta}\right. \\
&\quad \left. + \frac{\partial S(x_2, \theta_2)}{\partial x}2\kappa\Phi\varphi''(\theta_2)\frac{\varphi'(\theta_1)}{\varphi'(\theta_2)} - \frac{\partial S(x_2, \theta_2)}{\partial \theta}\frac{\varphi'(\theta_1)}{\varphi'(\theta_2)}\right]. &(94)
\end{aligned}$$

Because $\partial S/\partial\theta = -x$, the second and the fourth term in the square bracket cancel. Moreover, with $\partial S/\partial x = \overline{\theta} - x - \theta$, the previous expression becomes

$$-2\kappa\Phi f\varphi'(\theta_1)\left[(\overline{\theta} - \theta_1)\frac{\varphi''(\theta_1)}{\varphi'(\theta_1)} - x_1\frac{\varphi''(\theta_1)}{\varphi'(\theta_1)} - (\overline{\theta} - \theta_2)\frac{\varphi''(\theta_2)}{\varphi'(\theta_2)} + x_2\frac{\varphi''(\theta_2)}{\varphi'(\theta_2)}\right]. \tag{95}$$

Since $\varphi(\theta) = (\overline{\theta} - \theta)^{5/2}$, the first and the third term in the bracket cancel, and with the definition of x_i, I am left with

$$\frac{dV(f,\theta_1,\theta_2(\theta_1))}{d\theta_1} = -4\kappa^2\Phi^2 f\varphi'(\theta_1)\big[\varphi''(\theta_1)-\varphi''(\theta_2)\big]. \tag{96}$$

Since $\theta_1 < \theta_2$, the term in the square brackets is strictly positive, and since $\varphi'(\theta_1) < 0$, the expression is strictly positive overall. Hence, it is not optimal to have $\theta_1 > 0, \theta_2 \in (\underline{\theta}, \overline{\theta}]$

Therefore, either a degenerate distribution with a single point in its support is optimal, or a distribution which places positive mass on $\underline{\theta}$ is optimal. I now derive a sufficient condition for the latter to be the case.

(iii) I first determine the best degenerate distribution $F_{\hat{\theta}}$ for the principal which puts all mass on one point $\hat{\theta}$. Note,

$$\Phi_{F_{\hat{\theta}}} = \varphi(\hat{\theta}) = (\overline{\theta}-\hat{\theta})^{5/2}, \quad c_{F_{\hat{\theta}}}(\theta) = 2\kappa\varphi(\hat{\theta})\varphi(\theta), \quad c'_{F_{\hat{\theta}}}(\theta) = 2\kappa\varphi(\hat{\theta})\varphi'(\theta). \tag{97}$$

For $\hat{\theta} > \underline{\theta}$, the principal's profit in (38) writes

$$\begin{aligned} V(F_{\hat{\theta}}) &= S(-c'_{F_{\hat{\theta}}}(\hat{\theta}),\hat{\theta}) - 2\kappa\Phi^2_{F_{\hat{\theta}}} & (98)\\ &= -(\overline{\theta}-\hat{\theta})2\kappa\varphi(\hat{\theta})\varphi'(\hat{\theta}) - \frac{1}{2}4\kappa^2(\varphi(\hat{\theta})\varphi'(\hat{\theta}))^2 - 2\kappa\varphi(\hat{\theta})^2. & (99)\end{aligned}$$

Plugging in φ and simplifying yields:

$$V(F_{\hat{\theta}}) = 2\kappa\Big[3/2\cdot(\overline{\theta}-\hat{\theta})^5 - \kappa\cdot 25/4\cdot(\overline{\theta}-\hat{\theta})^8\Big]. \tag{100}$$

This expression is maximized at $\hat{\theta}^* = \overline{\theta} - (3/20\kappa)^{1/3}$ which is larger than $\underline{\theta} = \overline{\theta} - 1$ if $\kappa > 3/20$. In this case, the principal's profit is

$$V(F_{\hat{\theta}^*}) = \frac{9}{8}\left(\frac{3}{20}\right)^{5/3}\left(\frac{1}{\kappa}\right)^{2/3}. \tag{101}$$

(iv) I next derive the optimal distribution T_f that is supported on the boundaries $\underline{\theta}$ and $\overline{\theta}$ and places mass f on $\underline{\theta}$. I derive the solution under the assumption that $x^{FB}(\underline{\theta}) \geq -c'_{T_f}(\underline{\theta}) = -2\kappa\Phi_{T_f}\varphi'(\underline{\theta})$ so that $\underline{x}_{T_f} = x^{FB}(\underline{\theta}) = \overline{\theta} - \underline{\theta}$, and then verify later that this is indeed the case.

Under this assumption, the principal's profit is

$$V(T_f) = fS^{FB}(\underline{\theta}) + (1-f)(S(\overline{x}_{T_f}, \overline{\theta})) - 2\kappa(\Phi_{T_f})^2 \tag{102}$$

$$= f \cdot \frac{(\overline{\theta} - \underline{\theta})^2}{2} + (1-f)\cdot 0 - 2\kappa f^2 = 1/2 \cdot f - 2\kappa f^2. \tag{103}$$

where I have used that $\overline{\theta} - \underline{\theta} = 1$ and $\Phi_{T_f} = f\varphi(\underline{\theta}) + (1-f)\varphi(\overline{\theta}) = f$. This expression is maximized at $f^* = 1/(8\kappa)$, and thus $f^* \leq 1$ if $\kappa \geq 1/8$. Moreover, note that the solution satisfies the imposed assumption that $x^{FB}(\underline{\theta}) \geq -c'_{T_f}(\underline{\theta})$, since $-c'_{T_f}(\underline{\theta}) = -2\kappa\Phi_{T_{f^*}}\varphi'(\underline{\theta}) = 2\kappa f^* \cdot 5/2 = 5/8$ which is indeed smaller than $x^{FB}(\underline{\theta}) = \overline{\theta} - \underline{\theta} = 1$. Hence

$$V(T_{f^*}) = \frac{1}{32\kappa}. \tag{104}$$

(v) Finally, I compare profits. We have

$$V(T_{f^*}) > V(F_{\hat{\theta}^*}) \quad \Leftrightarrow \quad \kappa^{1/3} < \frac{1}{32} \cdot \frac{8}{9}\left(\frac{20}{3}\right)\Big)^{1/5} = \left(\frac{100}{81}\right)^2 \frac{5}{27} \approx 0.282 = \kappa_1. \tag{105}$$

Thus, for $\kappa < \kappa_1$, an optimal distribution places positive mass on $\underline{\theta}$. Together with (89), this implies the claim. QED

Proof of Proposition 2 By Theorem 4, there is an optimal distribution in the class of two-point distributions with support $\{\theta_1, \theta_2\}$, $\theta_1 \leq \theta_2$ and $f = Pr(\theta_1)$. Recall that $c'_F(\theta) = \Gamma'(\Phi_F)\varphi'(\theta)$. In what follows, I omit the subindex F on Φ. Lemma 4 implies that the optimal allocations $x_i = x(\theta_i)$ are given by

$$x_i = \begin{cases} \max\{x^{FB}(\underline{\theta}), -\Gamma'(\Phi)\varphi'(\underline{\theta})\} & if \quad \theta_i = \underline{\theta} \\ -\Gamma'(\Phi)\varphi'(\theta_i) & if \quad \theta_i \in (\underline{\theta}, \overline{\theta}) \\ \min\{x^{FB}(\overline{\theta}), -\Gamma'(\Phi)\varphi'(\overline{\theta})\} & if \quad \theta_i = \overline{\theta} \end{cases}. \tag{106}$$

To show that there are no upward distortions, it is sufficient to show that at an optimal distribution,

$$\frac{\partial S(x_i, \theta_i)}{\partial x} \geq 0, \quad i = 1, 2. \tag{107}$$

By (38), the principal's profit from a distribution in the above class is

$$V(f, \theta_1, \theta_2) = fS(x_1, \theta_1) + (1-f)S(x_2, \theta_2) - \tilde{\Gamma}(\Phi), \tag{108}$$

where $\tilde{\Gamma}(\Phi) = \alpha\Gamma(\Phi) + (1-\alpha)\Gamma'(\Phi)(\Phi - \varphi(\overline{\theta}))$. I first consider the case that $\theta_1 < \theta_2$ and distinguish various subcases.

Case 1: $x_2 = -\Gamma'(\Phi)\varphi'(\theta_2)$.

By (106), this means that θ_2 is interior, or if $\theta_2 = \overline{\theta}$, then x_2 is smaller than $x^{FB}(\overline{\theta})$.

Case 1(a): $x_1 = -\Gamma'(\Phi)\varphi'(\theta_1)$.

By (106), this means that θ_1 is interior, or if $\theta_1 = \underline{\theta}$, then x_1 is larger than $x^{FB}(\underline{\theta})$.

For a given value of Φ, define the function $\theta_1^\Phi(\theta_2)$ so that Φ is kept constant when θ_2 is varied, and let $V^\Phi(f, \theta_2)$ be the associated profit:

$$f\varphi(\theta_1^\Phi(\theta_2)) + (1-f)\varphi(\theta_2) = \Phi, \quad V^\Phi(f, \theta_2) = V(f, \theta_1^\Phi(\theta_2), \theta_2). \tag{109}$$

Then the change in profits when θ_2 is increased and Φ kept constant is given by

$$\begin{aligned}
\frac{\partial V^\Phi(f,\theta_2)}{\partial \theta_2} &= f\left[\frac{\partial S(x_1,\theta_1^\Phi)}{\partial x}\frac{\partial x_1}{\partial \theta_1}\frac{d\theta_1^\Phi}{d\theta_2} + \frac{\partial S(x_1,\theta_1^\Phi)}{\partial \theta}\frac{d\theta_1^\Phi}{d\theta_2}\right] \\
&\quad +(1-f)\left[\frac{\partial S(x_2,\theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_2} + \frac{\partial S(x_2,\theta_2)}{\partial \theta}\right] && (110)\\
&= f\frac{\partial S(x_1,\theta_1^\Phi)}{\partial x}\frac{\partial x_1}{\partial \theta_1}\frac{d\theta_1^\Phi}{d\theta_2} + (1-f)\frac{\partial S(x_2,\theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_2} - f x_1\frac{d\theta_1^\Phi}{d\theta_2} - (1-f)x_2 && (111)\\
&= f\frac{\partial S(x_1,\theta_1^\Phi)}{\partial x}\frac{\partial x_1}{\partial \theta_1}\frac{d\theta_1^\Phi}{d\theta_2} + (1-f)\frac{\partial S(x_2,\theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_2} && (112)\\
&= -f\frac{\partial S(x_1,\theta_1^\Phi)}{\partial x}\Gamma'(\Phi)\varphi''(\theta_1^\Phi)\frac{d\theta_1^\Phi}{d\theta_2} - (1-f)\frac{\partial S(x_2,\theta_2)}{\partial x}\Gamma'(\Phi)\varphi''(\theta_2). && (113)
\end{aligned}$$

Here, I have used that $\partial S/\partial\theta = -x$ in the second equality. In the third equality, I have used the definition of θ_2^Φ which implies that $f\varphi'(\theta_1^\Phi)\frac{d\theta_1^\Phi}{d\theta_2} = -(1-f)\varphi'(\theta_2)$ and the fact that $x_i = -\Gamma'(\Phi)\varphi'(\theta_i)$, and in the fourth equality I have inserted the derivatives of x_1 and x_2, keeping Φ fixed.

Moreover, the change in profit when (only) θ_2 is increased is

$$
\begin{aligned}
\frac{\partial V(f,\theta_1,\theta_2)}{\partial \theta_2} &= f\frac{\partial S(x_1,\theta_1)}{\partial x}\frac{\partial x_1}{\partial \theta_2}+(1-f)\Big[\frac{\partial S(x_2,\theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_2}+\frac{\partial S(x_2,\theta_2)}{\partial \theta}\Big] \\
&\quad -\alpha\Gamma'(\Phi)(1-f)\varphi'(\theta_2) \\
&\quad -(1-\alpha)\Gamma''(\Phi)(1-f)\varphi'(\theta_2)(\Phi-\varphi(\overline{\theta}))-(1-\alpha)\Gamma'(\Phi)(1-f)\varphi'(\theta_2) \quad (114)\\
&= f\frac{\partial S(x_1,\theta_1)}{\partial x}\frac{\partial x_1}{\partial \theta_2}+(1-f)\frac{\partial S(x_2,\theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_2}-(1-f)x_2 \\
&\quad -\Gamma'(\Phi)(1-f)\varphi'(\theta_2) \\
&\quad -(1-\alpha)\Gamma''(\Phi)(1-f)\varphi'(\theta_2)(\Phi-\varphi(\overline{\theta})) \quad (115)\\
&= f\frac{\partial S(x_1,\theta_1)}{\partial x}\frac{\partial x_1}{\partial \theta_2}+(1-f)\frac{\partial S(x_2,\theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_2} \\
&\quad -(1-\alpha)\Gamma''(\Phi)(1-f)\varphi'(\theta_2)(\Phi-\varphi(\overline{\theta})), \quad (116)
\end{aligned}
$$

where in the second line, I have used that $\partial S/\partial\theta = x$, and in third equality that $x_2 = -\Gamma'(\Phi)\varphi'(\theta_2)$.

Case 1(a, i): Suppose θ_1 and θ_2 are interior at an optimum. Then, optimality implies that expressions (113) and (116) are each equal to zero. Assume by contradiction to (107) that $\frac{\partial S(x_1,\theta_1)}{\partial x} < 0$. Then, (116) being equal to zero implies

$$
(1-f)\frac{\partial S(x_2,\theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_2} = -f\frac{\partial S(x_1,\theta_1)}{\partial x}\frac{\partial x_1}{\partial \theta_2}+(1-\alpha)\Gamma''(\Phi)(1-f)\varphi'(\theta_2)(\Phi-\varphi(\overline{\theta})). \quad (117)
$$

Since $\partial x_1/\partial\theta_2 \leq 0$, and since $\varphi'(\theta_2) < 0$, the right hand side is negative. Thus, since $\partial x_2/\partial\theta_2 \leq 0$, it follows that $\frac{\partial S(x_2,\theta_2)}{\partial x} > 0$. But now note that Γ' and φ'' are strictly positive by assumption and $d\theta_2^{\Phi}/d\theta_1 < 0$ because $\varphi'(\theta_i) < 0$ by assumption. Therefore, the two inequalities $\frac{\partial S(x_1,\theta_1)}{\partial x} < 0$ and $\frac{\partial S(x_2,\theta_2)}{\partial x} > 0$ imply that (113) is not zero, a contradiction. The symmetric argument shows that also $\frac{\partial S(x_2,\theta_2)}{\partial x} \geq 0$.

Case 1(a, ii): Suppose that $\theta_1 = \underline{\theta}$ and θ_2 is interior at an optimum. Then expression (116) is zero. Assume by contradiction to (107) that $\frac{\partial S(x_2,\theta_2)}{\partial x} < 0$. Then, because (116) is equal to zero, $\frac{\partial S(x_1,\theta_1)}{\partial x} > 0$ for the same reason as in the previous paragraph. But this contradicts the fact that in Case 1(a), when $\theta_1 = \underline{\theta}$, then $x_1 \geq x^{FB}(\underline{\theta})$ by (106).

Case 1(a, iii): Suppose that $\theta_2 = \overline{\theta}$ at an optimum. As explained above, we have $x_2 \leq x^{FB}(\overline{\theta})$ in this case. Hence, $\frac{\partial S(x_2,\theta_2)}{\partial x} \geq 0$. Moreover, expression (113) is (weakly) positive when $\theta_2 = \overline{\theta}$ at an optimum. For the same reason as in the paragraph before the previous one, this implies that $\frac{\partial S(x_1,\theta_1)}{\partial x} \geq 0$, thus establishing (107).

Case 1(b): $x_1 \neq -\Gamma'(\Phi)\varphi'(\theta_1)$.

By (106), this is only possible if $\theta_1 = \underline{\theta}$ and $-\Gamma'(\Phi)\varphi'(\theta_1) > x^{FB}(\underline{\theta})$ so that $x_1 = x^{FB}(\underline{\theta})$. In particular, x_1 is not upward distorted.

As to x_2, suppose that $\theta_2 = \overline{\theta}$ at an optimum. By (106), x_2 is then smaller than the first-best and also not upward distorted.

Suppose next that θ_2 is interior at an optimum. Since x_1 is first-best, $\partial S(x_1, \theta_1)/\partial x = 0$, and (116) becomes

$$\frac{\partial V(f, \theta_1, \theta_2)}{\partial \theta_2} = (1-f)\frac{\partial S(x_2, \theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_2} - (1-\alpha)\Gamma''(\Phi)(1-f)\varphi'(\theta_2)(\Phi - \varphi(\underline{\theta})). \quad (118)$$

Since θ_2 is interior, this expression is zero. Since $\partial x_2/\partial \theta_2 \leq 0$ and $\varphi'(\theta_2) < 0$, it follows that $\frac{\partial S(x_2,\theta_2)}{\partial x} \geq 0$, thus establishing (107).

Case 2: $x_2 \neq \Gamma'(\Phi)\varphi'(\theta_2)$.

By (106), this is only possible if $\theta_2 = \overline{\theta}$ and $x_2 = x^{FB}(\overline{\theta})$.

Case 2(a): $x_1 = \Gamma'(\Phi)\varphi'(\theta_1)$.

Consider the derivative of the principal's profit with respect to θ_1:

$$\begin{aligned}
\frac{V(f, \theta_1, \theta_2)}{\partial \theta_1} &= f\left[\frac{\partial S(x_1, \theta_1)}{\partial x}\frac{\partial x_1}{\partial \theta_1} + \frac{\partial S(x_1, \theta_1)}{\partial \theta}\right] + (1-f)\frac{\partial S(x_2, \theta_2)}{\partial x}\frac{\partial x_2}{\partial \theta_1} \\
&\quad -\alpha\Gamma'(\Phi)f\varphi'(\theta_1) \\
&\quad -(1-\alpha)\Gamma''(\Phi)f\varphi'(\theta_1)(\Phi - \varphi(\overline{\theta})) - (1-\alpha)\Gamma'(\Phi)f\varphi'(\theta_1) && (119) \\
&= f\frac{\partial S(x_1, \theta_1)}{\partial x}\frac{\partial x_1}{\partial \theta_1} - fx_1 - \Gamma'(\Phi)f\varphi'(\theta_1) && (120) \\
&\quad -(1-\alpha)\Gamma''(\Phi)f\varphi'(\theta_1)(\Phi - \varphi(\overline{\theta})) \\
&= f\frac{\partial S(x_1, \theta_1)}{\partial x}\frac{\partial x_1}{\partial \theta_1} - (1-\alpha)\Gamma''(\Phi)f\varphi'(\theta_1)(\Phi - \varphi(\overline{\theta})), && (121)
\end{aligned}$$

where in the second equality, I have used that $\frac{\partial S(x_2,\theta_2)}{\partial x} = 0$ since x_2 is first-best, and that $\partial S/\partial \theta = -x$, and in third equality that $x_1 = -\Gamma'(\Phi)\varphi'(\theta_1)$. At an optimum, (121) is (weakly) negative. Therefore, since $\partial x_1/\partial \theta_1 \leq 0$ and $\varphi'(\theta_1) < 0$, it follows that $\frac{\partial S(x_1,\theta_1)}{\partial x} \geq 0$, establishing (107).

Case 2(b): $x_1 \neq -\Gamma'(\Phi)\varphi'(\theta_1)$.

As in Case 1(b), x_1 is then first-best, and since also x_2 is first-best, (107) follows. This completes the proof for the case $\theta_1 < \theta_2$.

It remains to consider the case that $\theta_1 = \theta_2$, that is, the optimal distribution is degenerate with a single interior support point θ. Suppose first that $\theta = \overline{\theta}$. Then there is no upward distortion, because the optimal allocation $x(\overline{\theta})$ is always (weakly) smaller than the first-best by (106).

Suppose next that $\theta < \overline{\theta}$ and $x(\theta) \neq -\Gamma'(\Phi)\varphi'(\theta)$. By (106), this is only possible if $\theta = \underline{\theta}$ and $x(\underline{\theta}) = x^{FB}(\underline{\theta})$. Hence, there is no upward distortion.

Finally, suppose that $\theta < \overline{\theta}$ and $x(\theta) = -\Gamma'(\Phi)\varphi'(\theta)$. The principal's profit is then

$$V(\theta) = S(x(\theta), \theta) - \alpha\Gamma(\Phi) - (1-\alpha)\Gamma'(\Phi)(\Phi - \varphi(\overline{\theta})). \tag{122}$$

Hence, with $\partial S/\partial\theta = -x$, we have:

$$\frac{\partial V(\theta)}{\partial \theta} = \frac{\partial S(x(\theta),\theta)}{\partial x}\frac{\partial x(\theta)}{\partial \theta} - x(\theta) - \Gamma'(\Phi)\varphi'(\theta) - (1-\alpha)\Gamma''(\Phi)\varphi'(\theta)(\Phi - \varphi(\overline{\theta})) \tag{123}$$
$$= \frac{\partial S(x(\theta),\theta)}{\partial x}\frac{\partial x(\theta)}{\partial \theta} - (1-\alpha)\Gamma''(\Phi)\varphi'(\theta)(\Phi - \varphi(\overline{\theta})). \tag{124}$$

At an optimum, this expression is (weakly) negative. Since $\partial x/\partial\theta$ is negative and $\varphi'(\theta) < 0$, it follows that $\frac{\partial S(x(\theta),\theta)}{\partial x}$ is positive. This establishes (107) and completes the proof. QED

B Existence

In this appendix, I state sufficient conditions for the existence of a solution to the principal's problem. While these conditions might appear strong, given that the objective and the constraint depend in complicated ways on F (in particular through the allocation $x_F(\cdot)$), it is not surprising that in general, somewhat strong assumptions are needed to ensure continuity of the objective and compactness of the constraint set.

Proposition 3. *Suppose that for any sequence $(F_n)_n$ that converges weakly to F, we have that c_{F_n} and c'_{F_n} converge pointwise to c_F and c'_F respectively, and that c'_F and c''_F are bounded uniformly for all F.*[44] *Then the principal's problem $\mathscr{P}$ has a solution.*

Proof of Proposition 3 Let $(F_n)_n$ be a sequence converging weakly to F. I show: (i) The objective function V is upper semi-continuous, that is,

$$\limsup_n V(F_n) \le V(F). \tag{125}$$

(ii) The constraint set is compact, that is,

$$supp(F_n) \subseteq \breve{\Theta}_{F_n} \quad \forall n \quad \Rightarrow \quad supp(F) \subseteq \breve{\Theta}_F. \tag{126}$$

[44] That is, there is $B > 0$ so that $|c'_F(\theta)| < B$ and $|c''_F(\theta)| < B$ for all F and θ.

The claim then follows from Weierstraß' extreme value theorem applied to upper semi-continuous functions.

As to (i). I first rewrite the principal's objective as

$$V(G) = \int v(G,\theta)\,dG(\theta), \quad \text{with } v(G,\theta) = S(x_G(\theta),\theta) - c_G(\theta) + c_G(\underline{\theta}). \tag{127}$$

Moreover, let

$$\hat{v}(G,\theta) \;=\; S(-c_G'(\theta),\theta) - c_G(\theta) + c_G(\underline{\theta}). \tag{128}$$

Below, I show that for $n \to \infty$:

$$\hat{v}(F_n,\cdot) \to \hat{v}(F,\cdot) \quad \text{uniformly,} \tag{129}$$

$$\limsup_n v(F_n,\theta) \le v(F,\theta) \quad \text{for } \theta \in \{\underline{\theta},\overline{\theta}\}. \tag{130}$$

Note that $\hat{v}(F,\theta)$ coincides with $v(F,\theta)$ for all $\theta \in (\underline{\theta},\overline{\theta})$. Therefore:

$$\int v(F_n,\theta)\,dF_n \;=\; \int v(F_n,\theta) - v(F,\theta)\,dF_n + \int v(F,\theta)\,dF_n \tag{131}$$

$$= \big([v(F_n,\underline{\theta}) - v(F,\underline{\theta})] - [\hat{v}(F_n,\underline{\theta}) - \hat{v}(F,\underline{\theta})]\big)F(\underline{\theta}) \tag{132}$$

$$+ \int \hat{v}(F_n,\theta) - \hat{v}(F,\theta)\,dF_n \tag{133}$$

$$+ \big([v(F_n,\overline{\theta}) - v(F,\overline{\theta})] - [\hat{v}(F_n,\overline{\theta}) - \hat{v}(F,\overline{\theta})]\big)(F_n(\overline{\theta}) - F_n(\overline{\theta}^-)) \tag{134}$$

$$+ \int v(F,\theta)\,dF_n. \tag{135}$$

I now first argue that the terms (132)-(134) become arbitrarily small as n gets large. To do so, let $\epsilon > 0$. Due to (129) and (130), the lines (132) and (134) each get smaller than $\epsilon/4$ for large enough n.

Moreover, consider the integral in (133). In absolute value, this integral is smaller than $\sup_\theta |\hat{v}(F_n,\theta) - \hat{v}(F,\theta)|$, and thus gets smaller than $\epsilon/4$ when n gets large since $\hat{v}(F_n,\cdot)$ converges uniformly to $\hat{v}(F_n,\cdot)$ by (129).

Next, I provide a bound for (135) for large n. Observe that $v(F,\cdot)$ is continuous on $(\underline{\theta},\overline{\theta})$ (because $x_F(\theta)$ and $c_F(\theta)$ are continuous on $(\underline{\theta},\overline{\theta})$) and may have a downward jump at $\underline{\theta}$ and an upward jump at $\overline{\theta}$ (because at these points the allocation $x_F(\theta)$ may change discontinuously).

Thus, $v(F,\cdot)$ is upper semi-continuous. Now, it is well known that weak convergence of F_n to F and upper semi-continuity of $v(F,\cdot)$ implies that $\limsup_n \int v(F,\theta)\, dF_n \leq \int v(F,\theta)\, dF$. Consequently $\int v(F,\theta)\, dF_n < \epsilon/4 + \int v(F,\theta)\, dF$ for large enough n.

Taken together, these observations imply that one can find an N so that for all $n > N$, one has

$$V(F_n) = \int v(F_n,\theta)\, dF_n \leq \epsilon + \int v(F,\theta)\, dF. \tag{136}$$

Because ϵ was arbitrary, it follows that

$$\limsup_n V(F_n) \leq \int v(F,\theta)\, dF = V(F), \tag{137}$$

and this is (125).

To complete the proof of part (i), I have to show (129) and (130). As to (129), note first that $\hat{v}(F_n,\theta)$ converges to $\hat{v}(F,\theta)$ pointwise, because c_{F_n} and c'_{F_n}, respectively, converges pointwise to c_F and c'_F by assumption, and because S is continuous in x. It is well-known that pointwise convergence implies uniform convergence if the derivatives of the elements of the sequence are bounded uniformly for all elements of the sequence. Thus it is sufficient to show that the derivative of $\hat{v}(F_n,\theta)$ is bounded uniformly for all n. Indeed,

$$\begin{aligned}
\frac{d}{d\theta}\hat{v}(F_n,\theta) &= -\frac{\partial S(-c'_{F_n}(\theta),\theta)}{\partial x} c''_{F_n}(\theta) + \frac{\partial S(-c'_{F_n}(\theta),\theta)}{\partial \theta} - c'_{F_n}(\theta) && (138)\\
&= -\frac{\partial S(-c'_{F_n}(\theta),\theta)}{\partial x} c''_{F_n}(\theta) && (139)\\
&= -\Big(\beta'(-c'_{F_n}(\theta)) - \theta\Big) c''_{F_n}(\theta), && (140)
\end{aligned}$$

where the second line follows from the fact that $\partial S(x,\theta)/\partial\theta = -x$. Because c'_{F_n} and c''_{F_n} are bounded uniformly for all n by assumption, and since β' is continuous, $\frac{d}{d\theta}\hat{v}(F_n,\theta)$ is bounded uniformly for all n. This establishes (129).

As to (130). Below, I show that

$$\liminf_n -\breve{c}'_{F_n}(\underline{\theta}) \geq -\breve{c}'_F(\underline{\theta}) \quad and \quad \limsup_n -\breve{c}'_{F_n}(\overline{\theta}) \leq -\breve{c}'_F(\overline{\theta}). \tag{141}$$

Now recall that $\underline{x}_F = \max\{x^{FB}(\underline{\theta}), -\breve{c}'_F(\underline{\theta})\}$. If $\underline{x}_F = x^{FB}(\underline{\theta})$, then clearly

$$S(\underline{x}_{F_n},\underline{\theta}) \leq S(\underline{x}_F,\underline{\theta}). \tag{142}$$

If $\underline{x}_F = -\breve{c}'_F(\underline{\theta}) > x^{FB}(\underline{\theta})$, then the left part of (141) implies that at most finitely many elements of the sequence x_{F_n} are more efficient than $\underline{x}_F$, and hence,

$$\limsup_n S(\underline{x}_{F_n}, \underline{\theta}) \leq S(\underline{x}_F, \underline{\theta}). \tag{143}$$

A similar argument using the right part of (141) shows that

$$\limsup_n S(\overline{x}_{F_n}, \overline{\theta}) \leq S(\overline{x}_F, \overline{\theta}). \tag{144}$$

Together with the fact that c_{F_n} converges pointwise to c_F by assumption, this implies (130).

I next prove the left part of (141). The right part follows from symmetric consideration. To simplify notation, I write c_n for c_{F_n}, and c for c_F for the rest of the proof. Let $\epsilon > 0$, and define

$$\kappa_\epsilon \quad = \quad \sup\{k \mid k(\theta - \underline{\theta}) - \epsilon \leq c(\theta) - c(\underline{\theta}) + \epsilon \quad \forall \theta \in \Theta\}. \tag{145}$$

Note that κ_ϵ is decreasing in ϵ with $\lim_{\epsilon \to 0} \kappa_\epsilon = \breve{c}'(\underline{\theta})$. Recall that c_n converges uniformly to c by assumption, and thus for all $\epsilon > 0$ there is N_ϵ so that for all $n > N_\epsilon$,

$$-2\epsilon < c_n(\theta) - c_n(\underline{\theta}) - [c(\theta) - c(\underline{\theta})] < 2\epsilon \quad \forall \theta \in \Theta. \tag{146}$$

Suppose now, contrary to left part of (141), that $\liminf -\breve{c}'_n(\underline{\theta}) < -\breve{c}'(\underline{\theta})$. Then because $\lim_{\epsilon \to 0} \kappa_\epsilon = \breve{c}'(\underline{\theta})$, we can find $\hat{\epsilon} > 0$ so that for infinitely many $n > N_{\hat{\epsilon}}$, we have $\breve{c}'_n(\underline{\theta}) > \kappa_{\hat{\epsilon}}$. Hence, for all θ:

$$\kappa_{\hat{\epsilon}}(\theta - \underline{\theta}) \quad < \quad \breve{c}'_n(\underline{\theta})(\theta - \underline{\theta}) \leq c_n(\theta) - c_n(\underline{\theta}) \leq c(\theta) - c(\underline{\theta}) + 2\hat{\epsilon}, \tag{147}$$

where the second inequality follows by definition of $\breve{c}'_n(\underline{\theta})$ and the third inequality from (146). But this is a contradiction to the definition of $\kappa_{\hat{\epsilon}}$. This shows (141) and completes the proof of part (i).

As to (ii). I use the following well-known fact about weak convergence. For all $\theta \in supp(F)$ and for all $\epsilon > 0$ there is N so that for all $n > N$ we have that $supp(F_n) \cap (\theta - \epsilon, \theta + \epsilon) \neq \emptyset$.

Contrary to (126) suppose that there is $\theta \in supp(F)$ with $\theta \notin \breve{\Theta}_F$, that is,

$$\Delta = c_F(\theta) - \breve{c}_F(\theta) > 0. \tag{148}$$

Since c_F is continuous, the inequality holds in a neighbourhood around θ. Thus, there is $\epsilon > 0$ so that

$$c_F(\tau) - \breve{c}_F(\tau) > \frac{\Delta}{2} \quad \text{for all } \tau \in (\theta - \epsilon, \theta + \epsilon). \tag{149}$$

Thus, it follows from the above-mentioned fact that there is N so that for all $n > N$ there is $\theta_n \in supp(F_n) \cap (\theta - \epsilon, \theta + \epsilon)$ so that

$$c_F(\theta_n) - \breve{c}_F(\theta_n) > \frac{\Delta}{2}. \tag{150}$$

As mentioned above, since c_{F_n} converges pointwise to c_F, and the derivatives c'_{F_n} are bounded uniformly for all n by assumption, the convergence is, in fact, uniform. Below, I show that this implies that $\breve{c}_{F_n}$ also converges uniformly to $\breve{c}_F$. Thus, there is $\hat{N}$ so that for all $n > \hat{N}$ and all $\tau \in (\theta - \epsilon, \theta + \epsilon)$

$$|c_{F_n}(\tau) - c_F(\tau)| < \frac{\Delta}{8}, \quad |\breve{c}_{F_n}(\tau) - \breve{c}_F(\tau)| < \frac{\Delta}{8}. \tag{151}$$

Therefore, for $n > \max\{N, \hat{N}\}$, we have (recall that for $n > N$, we have $\theta_n \in (\theta - \epsilon, \theta + \epsilon)$):

$$\begin{aligned} c_{F_n}(\theta_n) - \breve{c}_{F_n}(\theta_n) &= c_{F_n}(\theta_n) - c_F(\theta_n) + c_F(\theta_n) - \breve{c}_F(\theta_n) + \breve{c}_F(\theta_n) - \breve{c}_{F_n}(\theta_n) && (152) \\ &> -\frac{\Delta}{8} + c_F(\theta_n) - \breve{c}_F(\theta_n) - \frac{\Delta}{8} > \frac{\Delta}{4}, && (153) \end{aligned}$$

where the final inequality follows from (150). But because $\theta_n \in supp(F_n)$, this is a contradiction to the assumption that $supp(F_n) \subseteq \breve{\Theta}_{F_n}$.

To complete the proof, I have to show that uniform convergence of c_{F_n} to c_F implies uniform convergence of $\breve{c}_{F_n}$ to $\breve{c}_F$. Indeed, because of uniform convergence of c_{F_n} to c_F, there is N so that for all $n > N$ and all τ:

$$c_F(\tau) + \epsilon \geq c_{F_n}(\tau) \geq c_F(\tau) - \epsilon. \tag{154}$$

Now, because $c_{F_n}(\tau) \geq \breve{c}_{F_n}(\tau)$ for all τ, the first inequality implies that $c_F(\tau) + \epsilon \geq \breve{c}_{F_n}(\tau)$ for all τ. Since $\breve{c}_{F_n}$ is convex, it follows that lower convex envelope of $c_F + \epsilon$, which equals $\breve{c}_F + \epsilon$, is larger than $\breve{c}_{F_n}$, that is,

$$\breve{c}_F(\tau) + \epsilon \geq \breve{c}_{F_n}(\tau). \tag{155}$$

Likewise, the second inequality implies that

$$\breve{c}_{F_n}(\tau) \geq \breve{c}_F(\tau) - \epsilon. \tag{156}$$

But these two inequalities imply the uniform convergence of $\breve{c}_{F_n}$ to $\breve{c}_F$, and this completes the proof of part (ii). QED

References

Anderson, Eric T., and James D. Dana Jr. When is price discrimination profitable?. Management Science 55.6 (2009): 980-989.

Arozamena, Leandro, and Estelle Cantillon. Investment incentives in procurement auctions. The Review of Economic Studies 71.1 (2004): 1-18.

Barron, Daniel, George Georgiadis, and Jeroen Swinkels. Optimal contracts with a risk-taking agent. Theoretical Economics 15.2 (2020): 715-761.

Baron, David P., and Roger B. Myerson. Regulating a monopolist with unknown costs. Econometrica: Journal of the Econometric Society (1982): 911-930.

Castro-Pires, Henrique, Hector Chade, and Jeroen Swinkels. Disentangling moral hazard and adverse selection. American Economic Review 114.1 (2024): 1-37.

Cerreia-Vioglio, Simone, Fabio Maccheroni, and Massimo Marinacci. Stochastic dominance analysis without the independence axiom. Management Science 63.4 (2017): 1097-1109.

Condorelli, Daniele, and Balázs Szentes. Information design in the holdup problem. Journal of Political Economy 128.2 (2020): 681-709.

Crémer, Jacques, and Fahad Khalil. Gathering information before signing a contract. The American Economic Review (1992): 566-578.

Crémer, Jacques, Fahad Khalil, and Jean-Charles Rochet. Strategic information gathering before a contract is offered. Journal of Economic Theory 81.1 (1998a): 163-200.

Crémer, Jacques, Fahad Khalil, and Jean-Charles Rochet. Contracts and productive information gathering. Games and Economic Behavior 25.2 (1998b): 174-193.

Diamond, Peter. Managerial incentives: On the near linearity of optimal compensation. Journal of Political Economy 106.5 (1998): 931-57.

Gasmi, Farid, Jean-Jacques Laffont, and William W. Sharkey. Empirical evaluation of regulatory regimes in local telecommunications markets. Journal of Economics & Management Strategy 8.1 (1999): 61-93.

Georgiadis, George, Doron Ravid, and Balazs Szentes. Flexible moral hazard problems. Econometrica 92.2 (2024): 387-409.

Gershkov, Alex, Benny Moldovanu, Philipp Strack, and Mengxi Zhang. A theory of auctions with endogenous valuations. Journal of Political Economy 129.4 (2021): 1011-1051.

Gottlieb, Daniel, and Humberto Moreira. Simple contracts with adverse selection and moral hazard. Theoretical Economics 17.3 (2022): 1357-1401.

Hébert, Benjamin. Moral hazard and the optimality of debt. The Review of Economic Studies 85.4 (2018): 2214-2252.

Hellwig, Martin F. Incentive problems with unidimensional hidden characteristics: A unified approach. Econometrica 78.4 (2010): 1201-1237.

Huber, Peter J. *Robust Statistics*. John Wiley & Sons, 1981.

Jewitt, Ian. Justifying the first-order approach to principal-agent problems. Econometrica (1988): 1177-1190.

Kleiner, Andreas. Optimal delegation in a multi-dimensional world. Mimeo. 2022.

Krähmer, Daniel. The hold-up problem with flexible unobservable investments. No. 278. ECONtribute Discussion Paper, 2024.

Krähmer, Daniel, and Roland Strausz. VCG mechanisms and efficient ex ante investments with externalities. Economics Letters 94.2 (2007): 192-196.

Laffont, Jean-Jacques, and David Martimort. *The theory of incentives: the principal-agent model*. Princeton University Press, 2002.

Laffont, Jean-Jacques, and Jean Tirole. Using cost observation to regulate firms. Journal of political Economy 94.3, Part 1 (1986): 614-641.

Liu, Siwen. Flexible Moral Hazard Problems with Adverse Selection. Mimeo. 2025.

Martimort, David, Jean-Christophe Poudou, and Lionel Thomas. Optimal contracts under moral hazard, adverse selection and limited liability. Mimeo. 2025.

Mensch, Jeffrey. Screening inattentive buyers. American Economic Review 112.6 (2022): 1949-1984.

Mensch, Jeffrey, and Doron Ravid. Monopoly, product quality, and flexible learning. Mimeo. 2024.

Mussa, Michael, and Sherwin Rosen. Monopoly and product quality. Journal of Economic Theory 18.2 (1978): 301-317.

Ok, Efe A. *Real analysis with economic applications*. Vol. 10. Princeton University Press, 2007.

Ollier, Sandrine, and Lionel Thomas. Ex post participation constraint in a principal–agent model with adverse selection and moral hazard. Journal of Economic Theory 148.6 (2013): 2383-2403.

Piccione, Michele, and Guofu Tan. Cost-reducing investment, optimal procurement and implementation by auctions. International Economic Review (1996): 663-685.

Rogerson, William P. Simple menus of contracts in cost-based procurement and regulation. American Economic Review 93.3 (2003): 919-926.

Rogerson, William P. Contractual solutions to the hold-up problem. The Review of Economic Studies 59.4 (1992): 777-793.

Sandmann, Christopher. When are sparce menus profit-maximizing?. Mimeo. 2024.

Szalay, Dezsö. Contracts with endogenous information. Games and Economic Behavior 65.2 (2009): 586-625.

Thereze, João. Screening costly information. Mimeo. 2024.

Winkler, Gerhard (1988). Extreme Points of Moment Sets. Mathematics of Operations Research 13(4), 581-587.