Security Design with Flexible Moral Hazard and Limited Liability

Daniel Krähmer*

October 18, 2026

Abstract

I study security design with moral hazard and two-sided limited liability when an entrepreneur can choose any distribution of project returns subject to a cost. I first characterize which return distributions can be induced by a security and when the first-best can be financed. Whenever the first-best can be financed, it is optimal and can be implemented by debt; more generally, any security that implements it pays the investor a constant amount on the support of the first-best distribution. The agency costs that may prevent first-best financing are related to the convexity of the cost function. For a given return distribution, greater convexity of the cost function increases the rent that the entrepreneur has to retain and can therefore make first-best financing more difficult. When the first-best cannot be financed, I derive second-best distributions for cost functions that are monotone in risk or depend on generalized moments. Under suitable conditions, the second-best is first-order stochastically dominated by the first-best. The optimal return distributions I characterize have finite support. Since an optimal security is uniquely determined only on, but not off, the support of the optimal return distribution, optimal securities are not unique.

Keywords: Security Design, Flexible Moral Hazard, Limited Liability

JEL: D21, D82, D86, G32

*Universität Bonn (kraehmer@hcm.uni-bonn.de); I thank Martin Hellwig and Roland Strausz for very valuable discussions. Special thanks go to Mathijs Janssen. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through Germany's Excellence Strategy EXC 2126/2–390838866 and CRC TR 224 (project B03).

1 Introduction

I study a canonical security design problem where an entrepreneur seeks outside funding by issuing a security that promises an investor a portion of the realized project return. The probability distribution of project returns depends on a hidden action by the entrepreneur, that is, there is moral hazard. In Jensen and Meckling's (1976) seminal contribution, security design under moral hazard is shaped by distinct agency problems, depending on whether moral hazard concerns the entrepreneur's effort or risk choices. More recent contributions, including Biais and Casamatta (1999), Hellwig (2009), and Hébert (2018), formalize environments in which the entrepreneur's hidden actions may involve both costly effort that improves project performance and unproductive choices that alter the distribution of returns, such as increasing its risk. In such environments, the question arises how the interaction among the different dimensions of the entrepreneur's hidden action affects agency costs and return distortions.

The present paper studies this question within a flexible moral hazard framework where the entrepreneur can choose any return distribution subject to a cost. Such a flexible approach is also adopted in Hébert (2018). My paper differs in two respects. First, Hébert (2018) is primarily interested in the optimality of certain financial instruments, notably debt, whereas my main focus is on when first-best returns can be implemented and, if not, how second-best returns are distorted. I also characterize the securities that implement the first-best. Second, Hébert (2018) restricts attention to cost functions that are in the class of divergences, whereas I use the smooth flexible moral hazard approach of Georgiadis et al. (2024), which allows for a broad class of convex, increasing, and smooth cost functions.

I study a version of the security design model of Innes (1990) in which the entrepreneur can choose any return distribution subject to a cost. Next to moral hazard, there is two-sided limited liability and a funding constraint: The entrepreneur's liability is limited to the project return, and the investor is not liable for losses incurred by the entrepreneur. Moreover, to participate in the venture, the investor requires that the expected security payout cover the cost of capital he provides.

The main results of my paper are as follows. I first characterize when the first-best can be implemented and is optimal. Any first-best return distribution can be induced by some security under moral hazard and limited liability. Whether it can also be financed depends on whether

the security leaves the investor a sufficiently large share of the project return to cover his capital costs. If this is the case, the first-best is optimal for the entrepreneur and can be implemented by debt. More generally, I show that any security that implements the first-best pays the investor the same amount for all returns in the support of the first-best distribution. This has a familiar residual-claimant interpretation. First-best incentives require that the entrepreneur receive the full marginal return from changing the return distribution. In the flexible approach, where the entrepreneur can alter the entire distribution, this requires the investor's payout to be constant for every return realization in its support, not only that the entrepreneur receives the residual surplus on average. A debt security has exactly this property.

The agency costs that may prevent the first-best from being financed are related to the convexity of the cost function. For a given return distribution, I show that greater convexity increases the rent that the entrepreneur has to retain. Intuitively, the entrepreneur has to be compensated according to her marginal costs of changing the return distribution, while her actual costs reflect average costs. Greater convexity increases the difference between marginal and average costs and therefore increases the entrepreneur's rent.[1]

If the first-best outcome is not optimal, I derive second-best solutions for cost functions which, for a given mean of the distribution, are monotone in its risk (in the sense of a mean-preserving spread). Costs that are decreasing in risk capture situations where breakthrough innovations are easier to achieve than reliable returns (as in VC financing), or where there is a risk-return trade-off (such as portfolio management). Costs that are increasing in risk capture, for example, mature industries where firms' potential for large innovations is relatively small. For these cost functions, fixing the mean identifies the cheapest distribution and reduces the problem to one dimension. I also consider cost functions that depend on a single generalized moment of the distribution. In this case, a similar reduction is obtained by convexification. I derive (fairly mild) conditions such that the second-best is first-order stochastically dominated by the first-best in all these cases.[2]

To the extent that first-order stochastic dominance is not a comparison in pure risk-terms, a conclusion that moral hazard and limited liability lead to excessive (or too little) risk-taking is therefore not warranted in my setting. This may perhaps seem somewhat surprising when

[1] As in Georgiadis et al. (2024) and Krähmer (2026a, 2026b), if costs are linear, this rent is, in fact, zero.

[2] More generally, I also consider "moment-based" costs that depend on finitely many generalized moments of the distribution. Optimal distributions then have finite support, but I do not obtain a general ranking of the first- and second-best.

costs are decreasing in risk, because the entrepreneur then has incentives to add risk to the return distribution, as this reduces her costs. For the same reason, however, risk is also socially desirable. Of course, this normative conclusion might change if risk-taking by the entrepreneur also has wider externalities on society as a whole.

To derive these results, I first ask which return distributions can be implemented by some security and then search for the optimal return distributions among these (in the spirit of Grossman and Hart, 1983).[3] This approach highlights the basic trade-off that underlies the security design problem. In a first-best world, the entrepreneur would like to implement a return distribution that maximizes the total expected return and just compensate the investor for his capital costs. Moral hazard and limited liability require the entrepreneur to retain a rent, and the remaining payout might then be insufficient to cover the investor's capital costs.

A prominent question in the literature is whether the theory can explain real-world securities such as debt or equity as the outcome of an optimal design.[4] In Innes' (1990) parametric and Hébert's (2018) divergence cost based approaches, debt arises as an optimal security when securities are required to be monotone.[5] In the present setting, debt can also implement the first-best whenever the first-best can be financed. At the same time, the security is generally not uniquely determined away from the support of the optimal distribution. Under the cost functions I consider, optimal distributions have finite support within the continuous interval of possible returns, leaving degrees of freedom to specify the security at returns that do not occur at an optimum.[6] Parametric approaches typically consider full-support distributions. Similarly, in Hébert (2018), the cost of a distribution without full support is infinite. This typically pins down the optimal security, and when, in addition, the security is required to be monotone, a debt security arises as optimal.[7]

Related literature

The paper is related to the literatures on flexible moral hazard and security design. The analysis builds on the smooth flexible moral hazard approach of Georgiadis et al. (2024) and applies it to

[3]The literature often imposes the constraint that the security be monotone. In this paper, I do not impose this constraint at any point.

[4]See Allen and Barbalau (2024) for a review.

[5]Hébert (2018) shows that monotonicity is not required in the special case of KL-divergence.

[6]On this point, see also Hellwig (2009).

[7]See also Yang and Zentefis (2024), who identify optimal monotone securities as extreme points of monotone functions.

the security design setting of Innes (1990). In Georgiadis et al. (2024), the principal offers a wage contract to an agent who covertly chooses a return distribution and is protected by limited liability. I use their characterization of the moral hazard constraint in my analysis. The security design problem differs in two aspects: First, it is the party who offers the contract (the entrepreneur) who also chooses the hidden action. Second, both parties are protected by limited liability, and there is a funding constraint that requires the investor to break even. The first point implies a different objective function, while the second implies that not all distributions are implementable by some security.

The smooth flexible moral hazard approach is also used in Krähmer (2026a, 2026b), but in different contracting environments. Krähmer (2026a) studies a hold-up problem in which contract terms are negotiated after investment, while Krähmer (2026b) studies regulation where there is also adverse selection.[8] In the present paper, the entrepreneur instead commits to a security before choosing the return distribution, and the set of implementable distributions is constrained by two-sided limited liability and a funding constraint. These features give rise to results that are specific to security design.

Hébert (2018) also studies security design within a flexible moral hazard framework, but considers a finite return space and focuses on divergence cost functions. These are not included in the set of cost functions that I allow, as divergences are not monotone and are infinite for distributions that do not have full support. Hébert's main focus is on the optimality of certain financial instruments, whereas I characterize the securities that implement the first-best and the distortions caused by moral hazard.[9]

In Hellwig (2009), the entrepreneur can choose the failure risk of the venture as well as its mean return if it succeeds. While the return in case of failure is zero, the return in case of success depends positively on the entrepreneur's costly effort as well as on the chosen failure risk. Hellwig (2009) shows that the second-best exhibits inefficiently large failure risk. In my model with full flexibility, I compare first- and second-best solutions in terms of first-order stochastic dominance.

Biais and Casamatta (1999) consider a security design problem where the entrepreneur has two effort choices and two risk choices with a risk-return trade-off. Only the high-effort, low-risk project is viable, and either the agency costs necessary to prevent excessive risk-taking are

[8] For flexible moral hazard models with adverse selection, see also Castro-Pires et al. (2025), Liu (2025).

[9] Flexible moral hazard problems with finite outcome spaces are also considered in Mattsson and Weibull (2023) in the context of wage design, and in Kocherlakota (2026) in the context of taxation.

sufficiently low, or a complete funding breakdown occurs. In my flexible approach, the trade-off between return and risk (and other moments of the distribution) is captured by the cost function, rather than by a restriction to the action set. This allows me to study how return distortions depend on the cost function.

Hellwig (2025) considers a flexible moral hazard problem where the agent's costs depend on the mean and (minus) the variance of the distribution and are thus decreasing in risk. The key difference to my paper is that the agent is risk averse and effort costs are monetary. Second-best solutions have at most three points in the support, similarly to the finite-support results for moment-based costs in my setting. Hellwig (2025) finds that the agent's remuneration may be non-monotone in returns. In the smooth flexible framework with monotone costs, the agent's/entrepreneur's remuneration is monotone on the support of returns, as already noted in Georgiadis et al. (2024).

The paper is organized as follows. Section 2 presents the model. Section 3 provides a characterization of implementable distributions. Section 4 derives the security design problem and characterizes when the first-best is implementable and optimal. Section 5 solves the security design problem for various classes of cost functions. Section 6 concludes. All proofs are in the appendix.

2 Model

There are a risk-neutral entrepreneur (she) and a risk-neutral investor (he). The entrepreneur needs funds $K > 0$ to conduct a project that pays a contractible return $x \in X = [\underline{x}, \overline{x}]$, $0 \leq \underline{x} < \overline{x}$. The return x is distributed with a cdf F that is chosen by the entrepreneur. I assume that the entrepreneur's choice is flexible: she can choose any return distribution F subject to the cost $C(F)$.

The entrepreneur owns no capital (for simplicity), and to obtain financing for the project from the investor, she issues a security $S : X \to \mathbb{R}$ that promises to pay the investor the amount $S(x)$ if the project return x realizes.

Both parties are protected by limited liability: The entrepreneur cannot pay out more than the project return, that is, $S(x) \leq x$, and the investor is not liable for any losses incurred by the entrepreneur, that is, $0 \leq S(x)$.

The timing is as follows. The entrepreneur commits to a security. The investor decides whether to invest. If he does not invest, both parties get their outside option of zero. If the investor invests, the entrepreneur covertly chooses a return distribution F, that is, there is moral hazard. Finally, returns realize and are divided as specified by the security.

The entrepreneur's problem is

$$\begin{aligned} P_0: \qquad \max_{S,F} & \int x - S(x)\, dF - C(F) \quad s.t. \\ & F \in \arg\max_{G} \int x - S(x)\, dG - C(G), && (MH) \\ & 0 \leq S(x) \leq x \quad \forall x \in X, && (LL) \\ & \int S(x)\, dF \geq K. && (IR) \end{aligned}$$

Note that limited liability is required for all possible returns x, not only for x in the support of the distribution that the entrepreneur actually chooses. The reason is that if the entrepreneur were to deviate to a distribution with different support, the outcome would still be subject to limited liability.

Next, I state the assumptions on the cost function which parallel those in Georgiadis et al. (2024).

1. C is continuous, monotone, and convex.[10]

2. C is smooth in the sense that C is Gateaux-differentiable with a continuous and differentiable Gateaux derivative $c_F : X \to \mathbb{R}$, that is, for $F, \tilde{F}$, we have[11]

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \big[C(F + \epsilon(\tilde{F} - F)) - C(F) \big] = \int c_F(x)\, d(\tilde{F} - F). \tag{1}$$

3. The cost of the distribution $\delta_{\underline{x}}$, which places mass 1 on $\underline{x}$, is normalized to 0: $C(\delta_{\underline{x}}) = 0$.[12]

Continuity is a technical condition that ensures the existence of various maximizers below. Monotonicity means that higher return distributions are more costly, that is, $C(F) \geq C(G)$ if F

[10] Continuity refers to the weak topology.

[11] It is well known that the Gateaux derivative is unique only up to a constant. The reason is that F and $\tilde{F}$ in (1) are both cdfs. Thus adding a constant to c_F does not affect the right-hand side.

[12] δ_x denotes the cdf that places mass 1 on x.

first-order stochastically dominates G. This is a natural assumption in my context. As pointed out by Georgiadis et al. (2024), convexity is without loss because the entrepreneur can randomize and thus $C(F)$ should be the expected cost of the cheapest randomization that generates it, implying convexity.

Smoothness captures a notion of differentiability which will make the analysis tractable. 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 $c_F(x)$ evaluated at x measures the marginal cost of increasing the probability mass assigned to x given F.[13] The final assumption is a normalization that ensures that "zero effort" has no cost.

I will make use of the fact that for smooth costs, monotonicity is characterized by monotonicity of the Gateaux derivative. That is, monotonicity of C in F is equivalent to $c_F(x)$ being increasing in x for all F (see, e.g., Cerreia-Vioglio et al., 2017).

A first-best distribution F^{fb} maximizes the total expected project value $\int x\, dF - C(F) - K$. Let

$$V^{fb} = \max_F \int x\, dF - C(F) - K$$

be the first-best project value. To make the problem non-trivial, I assume that V^{fb} is positive.

3 Implementability

I start the analysis by studying which distributions are implementable by some security. F is implementable if there is a security so that all the constraints in the entrepreneur's problem are satisfied.

As the first step, it is instructive to isolate the restrictions that arise from moral hazard and limited liability alone which are at the core of the incentive problem.[14] I say that F is incentive-feasible if there is a security S so that the constraints (MH) and (LL) in the entrepreneur's problem are satisfied. The next proposition characterizes incentive-feasibility. To state it, let

$$\underline{\lambda}_F = -c_F(\underline{x}), \quad \overline{\lambda}_F = \min_{x \in supp(F)} (x - c_F(x)). \tag{2}$$

[13] If $F = (F_1, \ldots, F_x, \ldots, F_I)$ is a vector, then $c_F(x)$ is the partial derivative $\frac{\partial C(F)}{\partial F_x}$.

[14] This incentive problem also arises if, for example, the investor has all the bargaining power and there is no participation constraint.

Proposition 1 *F is incentive-feasible if and only if there is $\lambda \in \mathbb{R}$ so that*

$$\underline{\lambda}_F \leq \lambda \leq \overline{\lambda}_F. \tag{LL'}$$

The key to understanding the proposition is a characterization of the moral hazard constraint (MH) in terms of a first-order condition that equates the entrepreneur's marginal benefits and marginal costs. Specifically, following Georgiadis et al. (2024), the constraint (MH) is satisfied if and only if there is λ so that[15]

$$S(x) = x - c_F(x) - \lambda \quad \forall x \in supp(F), \tag{MH_{supp}}$$

$$S(x) \geq x - c_F(x) - \lambda \quad \forall x \in X. \tag{MH_{all}}$$

Intuitively, $c_F(x)$ measures the entrepreneur's cost of marginally increasing the probability mass assigned to the return x, given F, and $x - S(x)$ is the entrepreneur's benefit of marginally increasing the probability mass assigned to return x (which is independent of F). The constant λ corresponds to the marginal shadow costs that come from the constraint that the total probability mass the entrepreneur can allocate is equal to one. A distribution F is therefore optimal if and only if there is no possible return level x at which the entrepreneur can strictly improve by marginally increasing the probability of x (condition (MH_{all})), and for all return levels in the support of F, the benefits and costs of marginally increasing the probability of x are the same (condition (MH_{supp})).

A useful interpretation is that equation (MH_{supp}) defines a security which specifies a "lump sum" λ (possibly negative) that the entrepreneur retains irrespective of the realized return, and when x realizes, the entrepreneur retains the additional amount $c_F(x)$ which is pinned down by the distribution F.

The moral hazard constraints *alone* do not restrict the magnitude of the lump sum, because, as is usual, the magnitude of the lump sum does not affect action incentives.[16] However, the lump sum does affect whether the security satisfies the additional constraints (LL). The lump sum must

[15] Strictly speaking, (MH) has to hold only for F-almost all $x \in supp(F)$. Throughout the paper, I abstract from measure-theoretic subtleties. Moreover, the distributions that arise endogenously in my setting are all discrete and finite.

[16] In particular, the moral hazard constraints alone do not restrict the set of implementable distributions. This is a well-known feature of flexible moral hazard with a risk-neutral agent (see Georgiadis et al., 2024).

not be too large (so that the investor's limited liability is not violated) and not be too small (so that the entrepreneur's limited liability is not violated). Condition (LL') specifies the range of λ's which are consistent with (LL).

It is instructive to illustrate the incentive-feasibility constraints (MH_{supp}), (MH_{all}), and (LL') graphically. The left panel of Figure 1 depicts a distribution F which satisfies these constraints. The support of F is $\{x_1, x_2, x_3\}$. The figure plots the curve $x - c_F(x) - \lambda$ for various values of λ. For a security to induce F with the specific lump sum λ_0, it has to pass through the curve $x - c_F(x) - \lambda_0$ on the support by (MH_{supp}) and lie above the curve elsewhere by (MH_{all}). Due to (LL), the security, moreover, has to be positive and smaller than the 45-degree line. In other words, any security which passes through the dotted points and is otherwise located in the shaded area satisfies the moral hazard and limited liability constraints for the lump sum λ_0. By decreasing (increasing) the lump sum, the curve shifts up (down). The lump sum $\underline{\lambda}_F$ is the smallest lump sum so that the security can be chosen to meet the entrepreneur's limited liability constraint $S(x) \leq x$, and the lump sum $\overline{\lambda}_F$ is the largest lump sum so that the security can be chosen to meet the investor's limited liability constraint $S(x) \geq 0$. Clearly, the depicted constellation has $\underline{\lambda}_F < \overline{\lambda}_F$, and (LL') holds.

In the right panel of Figure 1, the support of the distribution F is $\{x_1, x_2\}$. This distribution is not incentive-feasible. Indeed, the blue curve now plots $x - c_F(x) - \lambda$ for the smallest possible lump sum $\lambda = \underline{\lambda}_F$ that is still consistent with the LL constraint $S(x) \leq x$. Any smaller lump sum would shift the curve upward and imply a violation of the constraint $S(\underline{x}) \leq \underline{x}$. For the security to implement F, it would need to pass through the dotted points in order to satisfy the moral hazard constraint (MH_{supp}). But this is impossible without violating the constraint $S(x_2) \geq 0$. In other words, in the right panel, (LL') cannot be satisfied because $\underline{\lambda}_F > \overline{\lambda}_F$.

It is now straightforward to characterize when a distribution is implementable, that is, when in addition to the moral hazard and limited liability constraints, individual rationality is satisfied. Indeed, observe that by (MH_{supp}), a security that induces the entrepreneur to choose a given F is pinned down on the support of F up to the lump sum λ. Therefore, the investor's expected payoff is $\int S(x)\, dF = \int x - c_F(x)\, dF - \lambda$, and the individual rationality constraint is therefore equivalent to

$$\lambda \leq \int x - c_F(x)\, dF - K \equiv \lambda_F^{IR}. \qquad (IR')$$

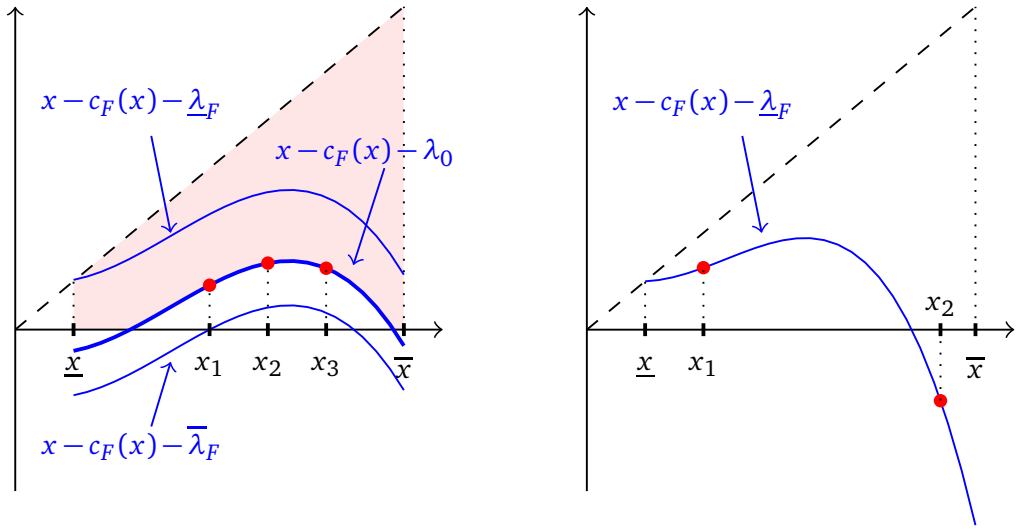


Figure 1: The figure illustrates when the incentive-feasibility conditions (MH_{supp}), (MH_{all}), and (LL') can (left) and cannot (right) be satisfied

4 Optimal security design

I now turn to the problem of finding an optimal security and an optimal distribution. Because a security that implements a given F is pinned down on the support of F up to the lump sum λ by (MH_{supp}), the entrepreneur's expected payoff from implementing F, gross of costs $C(F)$, is $\int x - S(x)\, dF = \int c_F(x)\, dF + \lambda$.

In light of (LL') and (IR'), the entrepreneur's problem P_0 can therefore be re-written as

$$P: \qquad \max_{F,\lambda} \int c_F(x)\, dF + \lambda - C(F) \quad s.t. \quad \underline{\lambda}_F \leq \lambda \leq \overline{\lambda}_F, \quad \lambda \leq \lambda_F^{IR}.$$

Before solving P formally, a reminder of the basic trade-off that underlies the security design problem is useful. Recall that $\lambda = \underline{\lambda}_F = -c_F(\underline{x})$ is the smallest lump sum that allows the entrepreneur to implement F and respect the limited liability constraint $S(x) \leq x$. The combination of moral hazard and this constraint implies that when F is implemented, the entrepreneur obtains *at least* the expected payout

$$\Pi_E(F) = \int c_F(x)\, dF - c_F(\underline{x}) \tag{3}$$

as a gross moral hazard rent (gross of costs). Accordingly, because the remaining portion of the

return is paid out to the investor, the investor obtains *at most* the expected payout

$$\Pi_I(F) = \int x \, dF - \Pi_E(F).$$

Absent any frictions, the entrepreneur would optimally maximize the overall expected project return by committing to choose a first-best distribution F^{fb}. She could then just compensate the investor for his capital costs K and in this way obtain the first-best project value V^{fb}. The key issue, however, is that the entrepreneur's minimal payout $\Pi_E(F^{fb})$ that is required to induce the entrepreneur to choose the first-best under moral hazard and limited liability might be too large for the remaining payout $\Pi_I(F^{fb})$ to allow the investor to recoup his capital costs K.

Conversely, this suggests that if there is a first-best distribution F^{fb} such that the investor can recoup his costs K, that is, $\Pi_I(F^{fb}) \geq K$, and which, in addition, is incentive-feasible, then F^{fb} is a solution to the security design problem. I now make these considerations more precise.

4.1 Implementability of the first-best

To examine when the first-best is implementable, I begin by characterizing first-best distributions.

Lemma 1 *A distribution F^{fb} is first-best if and only if there is μ so that*

$$x - c_{F^{fb}}(x) - \mu = 0 \quad \forall x \in supp(F^{fb}), \tag{FB_{supp}}$$

$$x - c_{F^{fb}}(x) - \mu \leq 0 \quad \forall x \in X. \tag{FB_{all}}$$

The conditions in the lemma are the first-order conditions that equate the marginal (first-best) project return and the marginal costs of taking an action. They are analogous to the moral hazard conditions, but in the first-best the benefit from marginally increasing the probability mass assigned to the return x is equal to x (instead of the entrepreneur's payout $x - S(x)$).

Next, I argue that a first-best distribution is incentive-feasible.

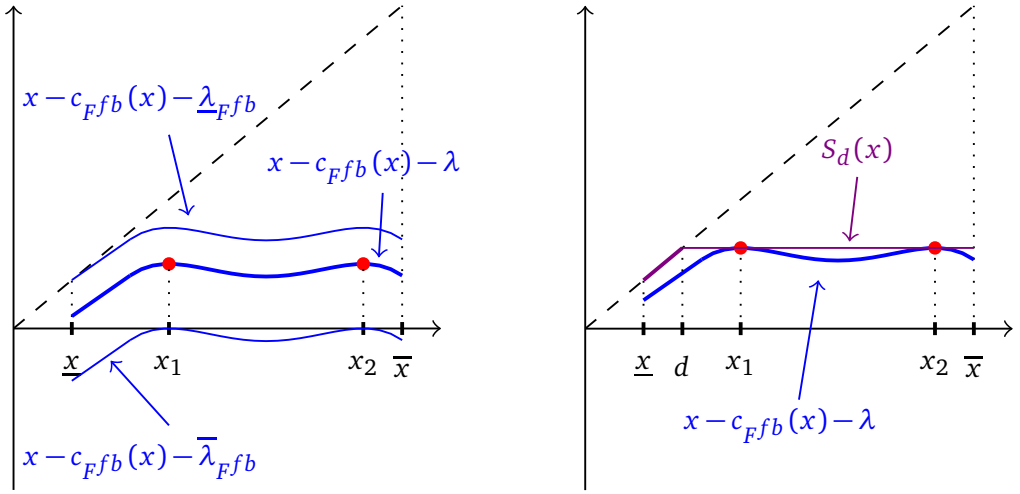
Lemma 2 *A first-best distribution F^{fb} is incentive-feasible, that is,*

$$\underline{\lambda}_{F^{fb}} \leq \overline{\lambda}_{F^{fb}}. \tag{4}$$

An intuitive way to see incentive-feasibility of the first-best is to observe that a constant security,

$S(x) = s$ for all x, makes the entrepreneur the residual claimant of the full surplus and thus induces her to choose a first-best distribution. Moreover, such a security can always be chosen to respect limited liability by setting $0 \leq s \leq \underline{x}$.

Formally, (4) follows directly from the conditions (FB_{supp}) and (FB_{all}) which imply that the function $x - c_{F^{fb}}(x)$ is maximized on the support of F^{fb}.[17] This is illustrated in panel (a) of Figure 2 which displays a first-best distribution with support $\{x_1, x_2\}$. The function $x - c_{F^{fb}}(x)$ is thus maximized on x_1 and x_2. It is therefore evident that the problem that $\underline{\lambda}_{F^{fb}}$ might be larger than $\overline{\lambda}_{F^{fb}}$ that arises in the right panel in Figure 1 does not arise for a first-best distribution, implying (4).



(a) A first-best distribution

(b) Implementation of F^{fb} by a debt security

Figure 2: Incentive-feasibility of the first-best

Moreover, as explained above, a security that satisfies (MH_{supp}), (MH_{all}), and (LL') for F^{fb} has to pass through the dotted points, lie above the curve $x - c_{F^{fb}}(x) - \lambda$, be positive and below the 45-degree line. As shown in panel (b), one such security is a debt security $S_d(x) = \min\{x, d\}$, where the face value d has to satisfy

$$d \in [0, \overline{d}], \quad \overline{d} = \max_{x \in supp(F^{fb})} x - c_{F^{fb}}(x) - \underline{\lambda}_{F^{fb}}$$

[17] Because $x - c_{F^{fb}}(x)$ is maximized by any point in the support of F^{fb} by (FB_{supp}), it follows

$$\underline{\lambda}_{F^{fb}} = -c_{F^{fb}}(\underline{x}) \leq \underline{x} - c_{F^{fb}}(\underline{x}) \leq \min_{x \in supp(F^{fb})} (x - c_{F^{fb}}(x)) = \overline{\lambda}_{F^{fb}}.$$

in order to respect (LL') and (MH_{all}).

Debt is not the only security that makes the first-best incentive feasible, because such a security is not pinned down off the support. Interestingly, however, any security that makes the first-best incentive feasible has the structure of "on-path" debt. Indeed, the first-best conditions (FB_{supp}) and (FB_{all}) imply that the *only* way to induce the entrepreneur to choose a first-best distribution is to pay the investor the *same* amount for *any* return realization in the support of the first-best distribution; it is not sufficient that the entrepreneur receives the residual surplus on average. Indeed, because $x - c_{F^{fb}}(x)$ is maximized and in particular constant on the support of F^{fb}, the moral hazard condition (MH_{supp}) can only be satisfied if the security $S(x) = x - c_{F^{fb}}(x) - \lambda$ is constant on the support of F^{fb}. This is a distinctive feature of the flexible moral hazard framework. Intuitively, if the security is not constant F^{fb}-almost surely, the entrepreneur's incentives are not aligned with first-best incentives across first-best support points. Because she can flexibly reallocate probability mass across returns, she then has an incentive to deviate from the first-best distribution.

Because a first-best distribution F^{fb} is incentive-feasible, it can be implemented if, in addition, the investor's individual rationality constraint can be satisfied. By (4), this is the case if and only if $\underline{\lambda}_{F^{fb}} \leq \lambda^{IR}_{F^{fb}}$. This inequality says that the smallest lump sum $\underline{\lambda}_{F^{fb}}$ that the entrepreneur has to retain to satisfy her limited liability constraint $S(x) \leq x$, is sufficiently small so that the investor's IR constraint can be satisfied.

Equivalently, the IR constraint can be satisfied if the face value of the implementing debt security covers the investor's capital costs, which is the case if and only if $\overline{d} \geq K$.

Finally observe that if F^{fb} is implementable, then F^{fb} is optimal for the entrepreneur. The reason is that she can then choose the lump sum so as to just compensate the investor for his capital costs K and in this way extract the entire first-best project value. The next proposition summarizes.

Proposition 2 *1. A first-best distribution F^{fb} is implementable if and only if $\underline{\lambda}_{F^{fb}} \leq \lambda^{IR}_{F^{fb}}$, or, equivalently, $\Pi_I(F^{fb}) - K \geq 0$, or, equivalently, $\overline{d} \geq K$.*

2. If a first-best distribution F^{fb} is implementable, then $F^* = F^{fb}$ and $\lambda^* = \Pi_I(F^{fb}) - K - c_{F^{fb}}(\underline{x})$ is a solution to P, and the entrepreneur's ex ante profit is the first-best project value V^{fb}.

3. If a first-best distribution F^{fb} is implementable, it can be implemented by a debt security with

face value $d \in [K, \overline{d}]$, and the optimal debt for the entrepreneur is $d^* = K$.

Note from part (iii) that a necessary condition for first-best implementation is that the minimum of the support of the first-best distribution exceed the implementation cost K. Thus, when implementable, a first-best debt security avoids default. This stands in contrast to the parametric approach in Innes (1990). There, low realizations cannot be avoided, and under an optimal debt security, default arises on path, and first-best is not implementable.[18]

Next, I show that the degree of convexity of the cost function drives whether the first-best is implementable or not. To do so, I denote by

$$\Pi_E^{net}(F) = \Pi_E(F) - C(F)$$

the entrepreneur's *net moral hazard rent*. The condition $\Pi_I(F^{fb}) \geq K$ for the implementability of the first-best can then be expressed as saying that the difference between the first-best surplus and the entrepreneur's net moral hazard rent is larger than K:

$$\int x \, dF^{fb} - C(F^{fb}) - \Pi_E^{net}(F^{fb}) \geq K. \tag{5}$$

Since the first-best surplus is larger than K by assumption, the first-best is therefore implementable if the net moral hazard rent is not too large.

Next, I show that the first-best is indeed implementable in the benchmark case that costs are linear in F. C is linear (in F) if $C(F) = \int \varphi(x) \, dF$ for some function φ.[19]

Proposition 3 *Let C be linear. Then the net moral hazard rent $\Pi_E^{net}(F)$ is zero for all F. In particular, the first-best is implementable.*

The result follows from the fact that for linear costs, the entrepreneur's minimal payout, $\Pi_E(F)$, is equal to costs $C(F)$ for all distributions F.[20] In other words, any implementable distribution can be implemented without affording the entrepreneur a net moral hazard rent.

[18] I thank the editor for pointing this out to me.

[19] Costs being linear means that the entrepreneur can generate a stochastic return distribution F only through a "mixed strategy" that randomizes over deterministic returns x, each costing $\varphi(x)$, according to F, and the costs of doing so are "expected costs". (When costs C are convex, such a mixed strategy is more costly than choosing F directly.)

[20] A similar observation appears in slightly different contexts in Georgiadis et al. (2024) and Krähmer (2026a, 2026b).

The previous reasoning suggests that the more convex the cost function, the higher the net moral hazard rent the entrepreneur accumulates. The next result confirms this. I say that $\tilde{C}$ is more convex than C if the difference $\tilde{C} - C$ is convex.

Proposition 4 *Let $\tilde{C}$ be more convex than C. Then $\tilde{\Pi}_E^{net}(F) \geq \Pi_E^{net}(F)$ for all F.*

The intuition is similar to standard marginal versus average cost reasoning. Fix a distribution F and consider moving from the no-investment distribution $\delta_{\underline{x}}$ towards F. The entrepreneur's payout required to induce F is determined by marginal costs at F, whereas her actual investment costs reflect average costs along this change. Greater convexity increases the difference between marginal and average costs and therefore increases the entrepreneur's net moral hazard rent.

Note that in the statement, F is arbitrary but fixed. Because a first-best distribution F^{fb} depends on the cost function, the result does not say that the net moral hazard rent, evaluated at the first-best, increases in the degree of convexity of the cost function. Therefore, Proposition 4 identifies merely one force that contributes to the severity of the agency problem. The following example illustrates this point and provides comparative statics for first-best implementability in a simple case.

Example. Normalize the range of possible returns such that $\overline{x} - \underline{x} = 1$, and let

$$\Phi_F = \int (x - \underline{x})\, dF \tag{6}$$

be the expected return in excess of the lowest possible return. Consider the family of cost functions

$$C_p(F) = \Gamma_p(\Phi_F), \qquad \Gamma_p(\Phi) = \frac{1}{2p}\left[(1+\Phi)^p - 1\right], \qquad p \geq 1. \tag{7}$$

For $p = 1$, costs are linear. For $p > 1$, costs are strictly convex. Moreover, if $\tilde{p} > p$, then $C_{\tilde{p}}$ is more convex than C_p in the sense of Proposition 4.[21]

Since costs depend only on the mean return, the first-best excess mean Φ^{fb} maximizes

$$\int x\, dF - C_p(F) = \underline{x} + \Phi - \Gamma_p(\Phi)$$

[21] Indeed, $\Gamma_p''(\Phi) = \frac{p-1}{2}(1+\Phi)^{p-2}$, and hence $\Gamma_{\tilde{p}}''(\Phi) \geq \Gamma_p''(\Phi)$ for all $\Phi \in [0,1]$.

over $\Phi \in [0,1]$. The first-order condition $1 = \Gamma_p'(\Phi)$ yields

$$\Phi^{fb} = \begin{cases} 1 & \text{if} \quad p \leq 2 \\ 2^{1/(p-1)} - 1 & \text{if} \quad p > 2. \end{cases}$$

For $p \leq 2$, the first-best distribution is therefore the degenerate distribution which puts mass 1 on the highest return $\overline{x}$: $F^{fb} = \delta_{\overline{x}}$. For $p > 2$, the first-best excess mean is interior and decreases with p. Any distribution with mean $\underline{x} + \Phi^{fb}$ is first-best.

For the cost function in (7), the entrepreneur's net moral hazard rent is[22]

$$\Pi_E^{net}(F) = \Gamma_p'(\Phi_F)\Phi_F - \Gamma_p(\Phi_F).$$

As long as $p \leq 2$, the first-best remains fixed at $\overline{x}$. Proposition 4 therefore implies that the net moral hazard rent evaluated at the first-best increases with p. Once $p > 2$, the first-best changes with p. Using the first-order condition $\Gamma_p'(\Phi^{fb}) = 1$, (9) becomes

$$\Pi_E^{net}(F^{fb}) = \frac{p-1}{p} 2^{1/(p-1)} - 1 + \frac{1}{2p}.$$

A tedious but straightforward calculation shows that $\Pi_E^{net}(F^{fb})$ strictly decreases for every $p > 2$. Thus, the net moral hazard rent evaluated at the first-best first increases and then decreases with p (see Figure 3). This illustrates why Proposition 4, which holds for a fixed distribution, does not imply monotone comparative statics for the net moral hazard rent evaluated at the first-best.

The example also gives a simple characterization of first-best implementability. Let $V(p)$ denote the first-best surplus before subtracting the investor's capital costs:

$$V(p) = \int x \, dF^{fb} - C_p(F^{fb}) = \begin{cases} \overline{x} - \dfrac{2^p - 1}{2p} & \text{if} \quad p \leq 2 \\ \underline{x} - 1 + \dfrac{p-1}{p} 2^{1/(p-1)} + \dfrac{1}{2p} & \text{if} \quad p > 2. \end{cases}$$

Thus, the first-best project value is $V^{fb} = V(p) - K$, and the first-best is viable if $K < V(p)$.

[22] Note that the Gateaux derivative is $c_F(x) = \Gamma_p'(\Phi_F)(x - \underline{x})$, and hence

$$\Pi_E^{net}(F) = \int c_F(x) \, dF - c_F(\underline{x}) - C_p(F) = \Gamma_p'(\Phi_F)\Phi_F - \Gamma_p(\Phi_F). \tag{8}$$

By definition, the maximal expected payout to the investor is the difference between $V(p)$ and the net moral hazard rent:

$$\Pi_I(F^{fb}) = V(p) - \Pi_E^{net}(F^{fb}) = \begin{cases} \overline{x} - 2^{p-2} & \text{if } p \leq 2 \\ \underline{x} & \text{if } p > 2. \end{cases}$$

Recall from Proposition 2 that the first-best is implementable if and only if $K \leq \Pi_I(F^{fb})$.

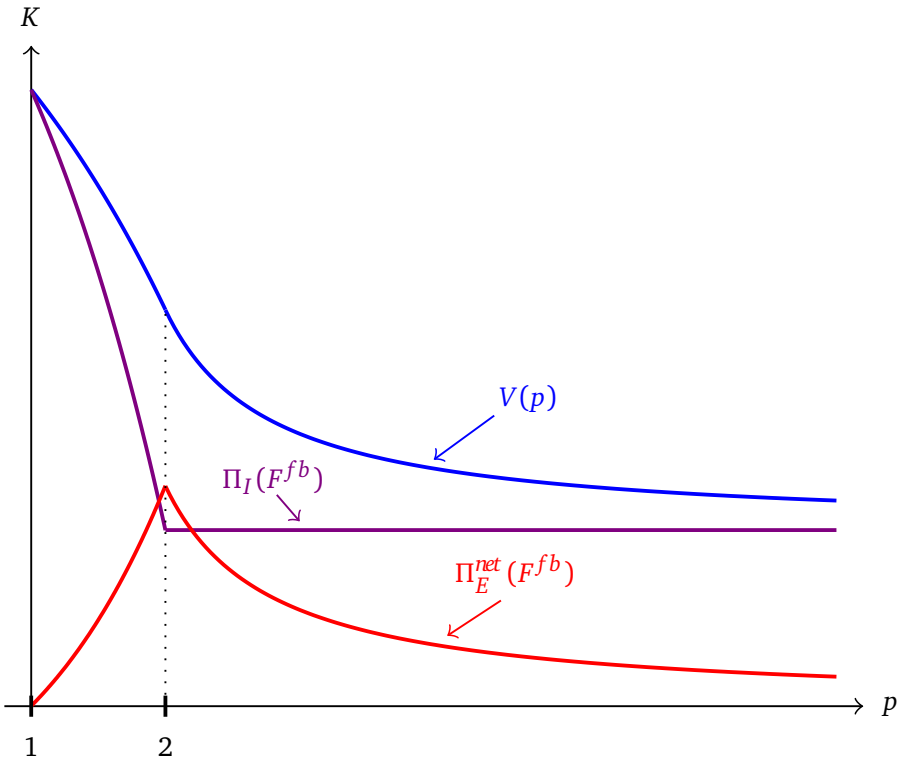


Figure 3: First-best viability, maximal expected payout to the investor, and net moral hazard rent as functions of p for $\underline{x} = 0.2$ and $\overline{x} = 1.2$.

Figure 3 illustrates the comparative statics. The region between $\Pi_I(F^{fb})$ and $V(p)$ is the set of capital costs K for which the first-best is viable but cannot be implemented, and the region below $\Pi_I(F^{fb})$ is the set of capital costs K for which the first-best can be implemented (and is viable).

For $p = 1$, costs are linear and the net moral hazard rent is zero (Proposition 3). As p increases and costs become more convex, the wedge between viability and implementability first increases. It is largest at $p = 2$, where the first-best becomes interior, and decreases thereafter. For $p > 2$, $\Pi_I(F^{fb}) = \underline{x}$, while $V(p)$ converges to $\underline{x}$ as p increases.

5 Second-Best

I now turn to the entrepreneur's problem when the first-best cannot be implemented and identify the resulting distortions. The main take-away is that under certain cost structures specified below, a second-best solution is first-order stochastically dominated by the first-best. Moreover, the second-best solutions I identify have finite support. This will imply that optimal securities that implement the second-best are not unique.

The fact that the first-best is incentive-feasible by Lemma 2 implies that if the first-best is not implementable, it is because the investor's IR constraint is violated even when the lump sum takes its smallest possible value $\underline{\lambda}_{Ffb}$ that is consistent with the entrepreneur's LL constraint $S(x) \leq x$. This suggests considering the relaxed problem where the lump sum λ is only required to satisfy the entrepreneur's limited liability constraint $S(x) \leq x$ and the IR constraint, that is,

$$\underline{\lambda}_F \leq \lambda \leq \lambda_F^{IR}. \tag{9}$$

Since the entrepreneur's objective in P is increasing in λ, it is optimal to set $\lambda = \lambda_F^{IR}$. If one plugs this into the objective and observes that the constraint (9) is equivalent to the condition $\int x\, dF - \Pi_E(F) \geq K$, one arrives at the relaxed problem

$$R: \qquad \max_F \int x\, dF - C(F) - K \quad s.t. \quad \int x\, dF - \Pi_E(F) - K \geq 0. \tag{10}$$

The relaxed problem has the intuitive interpretation that the entrepreneur maximizes the total project value subject to the constraint that the minimal payout to the investor covers his capital costs.

If K is large, the feasible set in problem R is empty, and the project is not funded. To rule out this rather uninteresting case, I assume from now on:

$$\sup_F \int x\, dF - \Pi_E(F) > K.$$

I shall solve problem R for various cost structures. Costs are increasing (resp. decreasing) in risk if $C(F) \geq C(G)$ (resp. $\leq$) whenever F is a mean-preserving spread of G. I say C is mean-based if it depends only on the mean of the distribution, in which case it is both increasing and decreasing

in risk.

Economically, the shape of the cost function depends on the application in question. For example, costs that are decreasing in risk capture features of VC financing of young, high-growth firms for whom a breakthrough innovation is easier to achieve than reliable returns. Likewise, the costs of a portfolio manager tend to be decreasing in risk, as it is easier to attain a given average return by investing in riskier asset classes. On the other hand, costs are increasing in risk for entrepreneurs in mature industries where firms' potential for large innovations is relatively small. Costs are mean-based when the entrepreneur has free access to fair gambles as in Barron et al. (2020).

I also consider costs that depend only on one generalized moment

$$\Phi_F = \int \varphi(x)\, dF$$

where $\varphi : X \to \mathbb{R}$ is an increasing function with $\varphi(\underline{x}) = 0$. Then C is moment-based if $C(F) = \Gamma(\Phi_F)$ for a strictly convex, increasing, and twice differentiable function $\Gamma : \mathbb{R} \to \mathbb{R}$.[23,24]

The general idea to solve for the second-best (and first-best) when costs have the specific structure described above is to reduce the problem to a one-dimensional problem parameterized by the mean $M = \int x\, dF$ of the distribution. This is formalized in the following lemma.

Lemma 3 *For $M \in [\underline{x}, \overline{x}]$, let $\mathscr{F}_M = \{F \mid \int x\, dF = M\}$ be the set of all distributions with mean M. Suppose that for all M, there is a $G_M \in \mathscr{F}_M$ so that for all $F \in \mathscr{F}_M$:*

$$C(G_M) \leq C(F) \quad \text{and} \quad \Pi_E(G_M) \leq \Pi_E(F). \tag{11}$$

Then there are solutions $F^{fb} = G_{M^{fb}}$ and $F^* = G_{M^*}$ to the first-best problem and problem R, respectively, where M^{fb} and M^* are the solutions to the one-dimensional maximization problems

$$M^{fb} \in \arg\max_M M - C(G_M), \quad M^* \in \arg\max_M M - C(G_M) \quad s.t. \quad M - \Pi_E(G_M) \geq K.$$

The lemma is straightforward: under the conditions of the lemma, the first-best problem and

[23] Setting $\varphi(\underline{x}) = 0$ is without loss. Because $C(\delta_{\underline{x}}) = 0$ means that $\Gamma(\varphi(\underline{x})) = 0$, one obtains an equivalent model with functions $\tilde{\Gamma}$ and $\tilde{\varphi}$ and with the property that $\tilde{\varphi}(\underline{x}) = 0$ by setting: $\tilde{\varphi}(x) = \varphi(x) - \varphi(\underline{x})$ and $\tilde{\Gamma}(\Phi) = \Gamma(\Phi + \varphi(\underline{x}))$.

[24] Convexity of Γ ensures convexity of C. Assuming strict convexity and differentiability simplifies some arguments, but is not required for my results. Note that when Γ is linear, costs C are linear so that Proposition 3 applies.

problem R can each be solved in two steps. One first solves for the optimal distributions among the set $\mathscr{F}_M$ of distributions with given mean M:

$$\begin{aligned} P_M^{fb}: &\qquad \max_{F\in\mathscr{F}_M} M - C(F), \\ R_M: &\qquad \max_{F\in\mathscr{F}_M} M - C(F) \quad s.t. \quad M - \Pi_E(F) \geq K. \end{aligned}$$

When condition (11) holds, G_M is a solution to both P_M^{fb} and, provided that G_M is feasible, also to R_M. In the second step, one maximizes over M to find the optimal distribution among all distributions.

5.1 C is increasing in risk

In this section, I assume that C is increasing in risk. I begin by showing that the first-best and second-best distributions are degenerate in this case.

Proposition 5 *Let C be increasing in risk.*

1. *The degenerate distribution $F^{fb} = \delta_{x^{fb}}$ where x^{fb} maximizes $x - C(\delta_x)$ is a first-best distribution. Moreover, the first-best is implementable if and only if $x^{fb} - \Pi_E(\delta_{x^{fb}}) - K \geq 0$.*

2. *If, in addition, Π_E is increasing in risk, then a solution to the relaxed problem R is the degenerate distribution δ_{x^*} where x^* maximizes*

$$x - C(\delta_x) \quad s.t. \quad x - \Pi_E(\delta_x) - K \geq 0. \tag{12}$$

Moreover, this solution is also a solution to the original problem P.

The result follows from Lemma 3 together with the observation that when C and Π_E are increasing in risk, then the degenerate distribution $G_M = \delta_M$ minimizes C and Π_E within the set $\mathscr{F}_M$ and, hence, satisfies (11). The implementability condition in part 1. is immediate from Proposition 2. The claim in part 2. that the solution to R also solves P can be easily verified.

Part 2. requires Π_E to be increasing in risk. Next, I show that this is the case if C is moment-based. In this case, the first- and the second-best can also be ranked.

Proposition 6 *Let C be increasing in risk and moment-based. Then Π_E is increasing in risk. Moreover, $x^* \leq x^{fb}$. In particular, the first-best first-order stochastically dominates the second-best.*

To see that Π_E is increasing in risk, I use the well-known fact that a Gateaux-differentiable function is increasing in risk if and only if its Gateaux derivative is convex in x for all F.[25] To see why Π_E has a convex Gateaux derivative, note that when C is moment-based, its Gateaux derivative is $c_F(x) = \Gamma'(\Phi_F)\varphi(x)$. Since C is increasing in risk, c_F is convex and since, in addition, Γ is increasing, it follows that φ is convex. Moreover, if C is moment-based, then

$$\Pi_E(F) = \int c_F(x)\, d(F - \delta_{\underline{x}}) = \Gamma'(\Phi_F)\Phi_F, \tag{13}$$

and the Gateaux derivative of Π_E is

$$(\Pi_E)_F(x) = [\Gamma''(\Phi_F)\Phi_F + \Gamma'(\Phi_F)]\varphi(x).$$

Since Γ is increasing and convex, the term in the square bracket is non-negative so that convexity of $(\Pi_E)_F(x)$ for all F follows from convexity of φ.

The fact that $x^* \leq x^{fb}$ is driven by two features of moment-based costs. First, $C(\delta_x)$ is convex in x. Second, the maximizer of the constraint in problem R is smaller than the maximizer of the objective. These two features do not need to hold in general, however. For example, it does not need to be the case that $C(\delta_x)$ is convex in x even though C is convex on the whole domain of cdfs.

5.2 C is decreasing in risk

In this section, I assume that C is decreasing in risk. This is to some extent the mirror image of the case discussed in the previous section, and an important role will now be played by distributions that are maximally risky in the sense that they put the entire mass on the most extreme return realizations.

I denote by T_f the distribution whose support contains at most the boundary points $\underline{x}$ and $\overline{x}$ and which places mass $f \in [0,1]$ on $\overline{x}$ (and mass $1-f$ on $\underline{x}$). Define $\kappa(f) = C(T_f)$ and $\pi_E(f) = \Pi_E(T_f)$. Moreover, I impose throughout this section the assumption that implementing "zero effort" is not valuable: $\underline{x} < K$. The main result is as follows.

[25] This characterization is due to Cerreia-Vioglio et al. (2017). The research program to characterize the risk properties of non-expected utility functionals in terms of their "local utility functions" goes back to Machina (1982) who considered Frechet-differentiable functionals. See also Hong and Nishimura (1992) for a partial characterization for Gateaux-differentiable functionals.

Proposition 7 *Let C be decreasing in risk.*

1. *A first-best distribution is a distribution $T_{f^{fb}}$ where f^{fb} maximizes $(1-f)\underline{x} + f\overline{x} - \kappa(f)$. Moreover, the first-best $T_{f^{fb}}$ is implementable if and only if it places mass 1 on the point $x = \overline{x}$ (that is, $f^{fb} = 1$) and*

$$\overline{x} - \pi_E(1) - K \geq 0. \tag{14}$$

2. *If, in addition, Π_E is decreasing in risk, then the solution to the relaxed problem R is the distribution T_{f^*} where f^* maximizes*

$$(1-f)\underline{x} + f\overline{x} - \kappa(f) \quad s.t. \quad (1-f)\underline{x} + f\overline{x} - \pi_E(f) - K \geq 0. \tag{15}$$

   Moreover, this solution is also a solution to the original problem P.

3. *It holds: $f^* \leq f^{fb}$. In particular, the first-best first-order stochastically dominates the second-best.*

The observations in part 1. and 2. that the first-best solution and the solution to R are in the class of distributions T_f follow from Lemma 3. Indeed, when C and Π_E are decreasing in risk, then the distribution $G_M = T_f$, where $f = f(M) = \frac{M-\underline{x}}{\overline{x}-\underline{x}}$ is such that T_f has mean M, minimizes C and Π_E within the set $\mathscr{F}_M$. Thus, G_M satisfies (11).

The implementability condition in part 1. says that the first-best can be implemented only if it assigns probability 1 to the largest return realization $\overline{x}$. Otherwise, the investor's IR constraint can never be satisfied. To see this intuitively, recall from Lemma 2 and Figure 2 that if the first-best is implementable, then the security $S(x)$ that implements it is constant on the support. Therefore, because $S(x) \leq x$ by LL, it follows that if a first-best distribution puts positive probability on the lowest return realization $\underline{x}$, then the security pays at most $\underline{x}$ for all return realizations in the support. Because $\underline{x} < K$ by assumption, the investor cannot recoup his capital costs under a security that pays less than K.

The claim in part 2. that the solution to R also solves P can be easily verified. Finally, the ranking in part 3. now (unlike the analogous ranking in Proposition 6) follows without additional assumptions. The reason is that $C(T_f)$ is now convex in f, which can be used to show that $f^* \leq f^{fb}$.

I conclude with the remark that, analogously to Proposition 6, Π_E is decreasing in risk if C is moment-based. The result by Cerreia-Vioglio et al. (2017) quoted above now says that a Gateaux-differentiable function is decreasing in risk if its Gateaux-derivative is concave in x for all F. Therefore, the arguments to show Proposition 6 equally apply.

5.3 Moment-based costs

In this section, I assume that costs are moment-based, but not necessarily increasing or decreasing in risk. Recall that in this case, $C(F) = \Gamma(\Phi_F)$ with $\Phi_F = \int \varphi(x)\, dF$. Let

$$\breve{\varphi}(x) = \sup\{h(x) \mid h \text{ convex}, h(y) \leq \varphi(y) \text{ for all } y \in X\}$$

be the lower convex envelope of φ. The next proposition states the result for moment-based costs.

Proposition 8 *Let C be moment-based.*

1. *There is a first-best distribution with at most two points in its support. The first-best value V^{fb} is the value of the maximization problem $\max_{M \in [\underline{x}, \overline{x}]} M - \Gamma(\breve{\varphi}(M)) - K$, and the maximizer is the first-best mean $M^{fb} = \int x\, dF^{fb}$.*

2. *There is a solution F^* to problem R with at most two points in its support, and the value of R is the value of the maximization problem*

$$\max_M M - \Gamma(\breve{\varphi}(M)) - K \quad s.t. \quad M - \Gamma'(\breve{\varphi}(M))\breve{\varphi}(M) - K \geq 0. \tag{16}$$

   and its maximizer M^* is the mean of F^*.

3. *Moreover, $M^* \leq M^{fb}$, and there exist a first-best distribution F^{fb} and a solution F^* to R, each with at most two points in their support, such that F^{fb} first-order stochastically dominates F^*.*

4. *If F^* is a degenerate distribution, then it also solves the original problem P. Otherwise, F^* solves the original problem if and only if for both points x_1, x_2 in the support, it holds: $x_i - \Gamma'(\breve{\varphi}(M^*))\varphi(x_i) \geq 0$. A sufficient primitive condition for this is that the slope of φ is bounded by $1/\Gamma'(\varphi(\overline{x}))$.*

Parts 1. and 2. follow again from Lemma 3. To see this, note that $C(F) = \Gamma(\Phi_F)$, and Γ is increasing. Likewise, recall from (13) that with moment-based costs, one has

$$\Pi_E(F) = \Gamma'(\Phi_F)\Phi_F.$$

Due to convexity of Γ, the function $\Gamma'(\Phi)\Phi$ is increasing in Φ. Therefore, minimizing both C and Π_E among all distributions with mean M now amounts to minimizing Φ_F in that class: $\min_{F:\int x\,dF=M} \Phi_F$.

This is a linear problem with one linear constraint, and it is well-known that there is a solution G_M to this problem that has at most two points in its support (see Winkler, 1988), and thus satisfies (11). Moreover, standard convexification arguments imply that the moment of the solution G_M is $\Phi_{G_M} = \breve{\varphi}(M)$, which explains the form of the first-best and second-best objectives in the statement.[26]

Part 3. of the proposition can best be seen for the case that F^{fb} and F^* are unique. Recall that F^{fb} and F^* are obtained by convexifying φ. The interval $[\underline{x}, \overline{x}]$ can be divided into disjoint segments in each of which the convexification $\breve{\varphi}$ corresponds to an affine function (at a point where $\breve{\varphi}$ is strictly convex, the segment is just this point). The support of F^i consists of the lower and upper points of the segment in which M^i is located, $i \in \{fb, *\}$. Now, because $M^* \leq M^{fb}$, there are only two possibilities. Either M^* and M^{fb} are located in the same segment. In this case, F^{fb} and F^* have the same support. Or, M^* is located in a "lower" segment than M^{fb}. In this case, all points in the support of F^* are smaller than all points in the support of F^{fb}. In either case, F^{fb} first-order stochastically dominates F^*. (Part 4. is technical and shows that the constraints that were neglected in the relaxed problem are met under the stated conditions.)

I conclude this section with a remark on the case with general moment-based cost functions that depend not on one but on an arbitrary finite number of moments. In this case, the one-dimensional reduction of Lemma 3 does not work. But the problem can still be reduced to a finite-dimensional problem. The next proposition characterizes a second-best solution.

Proposition 9 *Let $\varphi = (\varphi_1, \ldots, \varphi_N)$ be a vector of N increasing functions $\varphi_n : X \to \mathbb{R}$ with $\varphi_n(\underline{x}) =$*

[26] I am grateful to Siwen Liu for pointing this out to me.

0.[27] *Define the moment vector*

$$\Phi_F = (\Phi_{1,F}, \ldots, \Phi_{N,F}), \quad \Phi_{n,F} = \int \varphi_n(x) dF.$$

Suppose that costs depend only on the moment vector, that is, $C(F) = \Gamma(\Phi_F)$*, where* $\Gamma : \mathbb{R}^N \to \mathbb{R}$ *is convex, (coordinate-wise) increasing and differentiable.*

Then there is a discrete distribution F^* *with at most* $N+1$ *mass points that is a solution to R.* F^* *is also a solution to the original problem if and only if for all points* $x \in supp(F^*)$*, it holds:* $x - \Gamma'(\Phi_{F^*}) \cdot \varphi(x) \geq 0$*.*

The result rests on the fact that the Gateaux derivative $c_F(x) = \Gamma'(\Phi_F) \cdot \varphi(x)$ is the dot-product of the gradient Γ' and the vector $\varphi(x)$. Therefore, the entrepreneur's minimal payout also depends only on the moment vector,

$$\Pi_E(F) = \int c_F(x)\, d(F - \delta_{\underline{x}}) = \Gamma'(\Phi_F) \cdot \Phi_F,$$

and problem *R* writes:

$$\max_F \int x\, dF - \Gamma(\Phi_F) - K \quad s.t. \quad \int x\, dF - \Gamma'(\Phi_F) \cdot \Phi_F - K \geq 0.$$

For given $\Phi_n \in [\varphi_n(\underline{x}), \varphi_n(\overline{x})]$, consider now the subproblem where the moments are fixed: $\Phi_{n,F} = \Phi_n, n = 1, \ldots, N$. Both the objective and the constraint then depend only on the mean $\int x\, dF$ of the distribution, and both are increasing in the mean. Therefore, the problem is equivalent to maximizing the mean subject to the N moment constraints $\Phi_{n,F} = \Phi_n, n = 1, \ldots, N$. Because this is a linear problem with N linear constraints, there is a solution in the set of extreme points of the constraint set. It is well known (Winkler, 1988) that the set of extreme points is the set of discrete distributions with at most $N+1$ mass points.

The comparison of the second-best solution with the first-best is difficult. While similar arguments as in the previous paragraph imply that there is a discrete first-best distribution with finitely many mass points, an argument in the spirit of the convexification argument that I used in the single-moment case is not available here. I leave this question for future research.

[27] As in the case with a single moment ($N = 1$), this is without loss.

6 Conclusion

In this paper, I study security design with flexible moral hazard and two-sided limited liability. I characterize when the first-best can be implemented and the securities that implement it. Whenever the first-best can be financed, it can be implemented by debt. More generally, any security that implements the first-best must pay the investor the same amount for all return realizations in its support. The agency costs that may prevent the first-best from being financed are related to the convexity of the cost function: for a given return distribution, greater convexity increases the rent that the entrepreneur has to retain. When the first-best cannot be implemented, I derive second-best solutions for various cost structures. Under the conditions stated in the paper, the first-best first-order stochastically dominates the second-best. Thus, even when the entrepreneur has incentives to increase risk, the main distortion need not take the form of excessive risk-taking.

A prominent question in the literature is whether the theory can explain real-world securities such as debt or equity as the outcome of an optimal design. In the present setting, the security is sharply pinned down on the support of the optimal return distribution, while there are generally degrees of freedom away from the support. The optimal distributions I characterize have finite support within an interval of possible returns, so several securities can implement the same optimal return distribution. The approach therefore gives sharper predictions about security payouts at returns that occur at an optimum than about the complete shape of the security. Introducing additional restrictions on the entrepreneur's choice of distributions, for example through noise, may further restrict the security away from the support and seems an interesting avenue for future research.

7 Proofs

Proof of Proposition 1 By Lemma 1 in Georgiadis et al. (2024), and as explained in the main text, the moral hazard constraint (MH) is equivalent to (MH_{supp}) and (MH_{all}).

Now, let F be incentive-feasible. The inequality $\underline{\lambda}_F \leq \lambda$ in (LL') follows from (MH_{all}) and the inequality $S(x) \leq x$ in (LL). The inequality $\lambda \leq \overline{\lambda}_F$ in (LL') follows from (MH_{supp}) and the inequality $0 \leq S(x)$ in (LL).

Conversely, let (LL') hold. Define the security $S(x) = x - c_F(x) - \lambda$ for all $x \in supp(F)$, and $S(x) = x$ for all $x \notin supp(F)$. Then S satisfies (MH_{supp}) by construction. Condition (MH_{all}) holds because $\lambda \geq \underline{\lambda}_F$ by (LL') and because c_F is monotone in x. Hence, for $x \notin supp(F)$:

$$S(x) = x = x - c_F(\underline{x}) - \underline{\lambda}_F \geq x - c_F(x) - \lambda.$$

Moreover, using a similar chain of inequalities, it follows that S satisfies LL for $x \in supp(F)$ because $\underline{\lambda}_F \leq \lambda \leq \overline{\lambda}_F$ by (LL') and because c_F is monotone in x. Finally, S trivially satisfies LL for $x \notin supp(F)$. QED

Proof of Lemma 1 The claim follows from Lemma 1 in Georgiadis et al. (2024). QED

Proof of Lemma 2 The argument is in footnote 17. QED

Proof of Proposition 2 As to 1. Let F^{fb} be a first-best distribution. By (FB_{supp}), it follows that

$$\overline{\lambda}_{F^{fb}} = \mu, \quad \lambda^{IR}_{F^{fb}} = \mu - K, \tag{17}$$

and thus $\lambda^{IR}_{F^{fb}} < \overline{\lambda}_{F^{fb}}$. By Proposition 1, F^{fb} is therefore implementable if and only if $\underline{\lambda}_{F^{fb}} \leq \lambda^{IR}_{F^{fb}}$. Moreover, using the definition of $\underline{\lambda}_F$ and λ^{IR}_F, it is straightforward to verify that this inequality is equivalent to $\Pi_I(F^{fb}) \geq K$.

As to 2. Let F^{fb} be implementable. Part 1. thus implies that $\underline{\lambda}_{F^{fb}} \leq \lambda^{IR}_{F^{fb}}$. Accordingly, F^{fb} can be implemented by choosing any $\lambda \in [\underline{\lambda}_{F^{fb}}, \lambda^{IR}_{F^{fb}}]$. For the particular choice

$$\lambda = \lambda^{IR}_{F^{fb}} = \int x - c_{F^{fb}}(x)\, dF^{fb} - K = \Pi_I(F^{fb}) - K - c_{F^{fb}}(\underline{x}), \tag{18}$$

the entrepreneur's expected payoff is

$$\int c_{F^{fb}}(x)\, dF^{fb} - C(F^{fb}) + \lambda^{IR}_{F^{fb}} = \int x\, dF^{fb} - C(F^{fb}) - K. \tag{19}$$

Thus, the entrepreneur gets the first-best project value V^{fb}, and she can clearly not do better than implementing F^{fb} in this way.

As to 3. The argument follows from the main text. QED

Proof of Proposition 3 If C is linear, the Gateaux derivative is $c_F(x) = \varphi(x)$. Moreover, $C(\delta_{\underline{x}}) = \varphi(\underline{x}) = 0$ by assumption. The claim is thus immediate from the definition of $\Pi_E(F)$. QED

Proof of Proposition 4 Let $B(F) = \tilde{C}(F) - C(F)$. Then

$$\widetilde{\Pi}^{net}_E(F) - \Pi^{net}_E(F) = \int b_F(x)\, d(F - \delta_{\underline{x}}) - B(F). \tag{20}$$

Since B is convex by assumption, it follows from Proposition 2 in Krähmer (2026a) that the expression on the right-hand side is non-negative. QED

Proof of Lemma 3 The argument is in the main text. QED

Proof of Proposition 5 All arguments are in the main text except for the claim that the solution to R is a solution to the original problem P. To see this claim, recall that when deriving problem R, λ was optimally set equal to λ^{IR}_F. Because the constraint $\lambda \leq \overline{\lambda}_F$ is dropped in the relaxed problem, it follows that a solution F^* to the relaxed problem is also a solution to the original problem if and only if $\lambda^{IR}_{F^*} \leq \overline{\lambda}_{F^*}$, that is,

$$\int x - c_{F^*}(x)\, dF^* - K \leq \min_{x \in supp(F^*)} (x - c_{F^*}(x)). \tag{21}$$

Now observe that for the degenerate distribution $F^* = \delta_{x^*}$, condition (21) writes

$$x^* - c_{F^*}(x^*) - K \leq x^* - c_{F^*}(x^*), \tag{22}$$

and is thus satisfied. QED

Proof of Proposition 6 It remains to show that $x^* \leq x^{fb}$. This is trivial if $x^{fb} = \overline{x}$. Hence, let

$x^{fb} < \overline{x}$. Denote by $\alpha(x)$ the objective function and by $\beta(x)$ the constraint in problem (12). Since costs are moment-based,

$$\alpha(x) = x - \Gamma(\varphi(x)), \quad \beta(x) = x - \Gamma'(\varphi(x))\varphi(x) - K. \tag{23}$$

Note that because Γ is strictly convex and φ is convex (because costs are increasing in risk), x^{fb} uniquely maximizes α, and it follows that $\alpha'(x) < 0$ for all $x > x^{fb}$. Moreover, $\alpha'(x) \geq \beta'(x)$, since Γ is strictly convex. Therefore, it follows that also $\beta'(x) < 0$ for all $x > x^{fb}$. Now, suppose, contrary to the claim, that $x^{fb} < x^*$. Then lowering x^* would increase the objective and relax the constraint, a contradiction to the optimality of x^*. QED

Proof of Proposition 7 Part 1.: The argument for why a distribution T_f is optimal is given in the main text. The problem stated in the proposition is the first-best problem in the class of these distributions.

For the rest of the proof, I first establish the following properties of the functions κ and π_E:

$$\kappa'(f) = c_{T_f}(\overline{x}) - c_{T_f}(\underline{x}), \quad \textit{and} \quad \pi_E(f) = f\,\kappa'(f). \tag{24}$$

To see this, let $f \in [0,1)$ and note that $T_{f+\epsilon} = T_f + \frac{\epsilon}{1-f}(T_1 - T_f)$, which implies

$$\begin{aligned}
\kappa'(f) &= \frac{dC(T_f)}{df} && (25)\\
&= \lim_{\epsilon\to 0} \frac{1}{\epsilon}\big[C(T_{f+\epsilon}) - C(T_f)\big] && (26)\\
&= \lim_{\epsilon\to 0} \frac{1}{\epsilon}\left[C(T_f + \frac{\epsilon}{1-f}(T_1 - T_f)) - C(T_f)\right] && (27)\\
&= \frac{1}{1-f}\lim_{\epsilon\to 0} \frac{1}{\frac{\epsilon}{1-f}}\left[C(T_f + \frac{\epsilon}{1-f}(T_1 - T_f)) - C(T_f)\right] && (28)\\
&= \frac{1}{1-f}\int c_{T_f}(x)\, d(T_1 - T_f) && (29)\\
&= c_{T_f}(\overline{x}) - c_{T_f}(\underline{x}). && (30)
\end{aligned}$$

Moreover,

$$\pi_E(f) = \Pi_E(T_f) = \int c_{T_f}(x)\, dT_f - c_{T_f}(\underline{x}) = f\,[c_{T_f}(\overline{x}) - c_{T_f}(\underline{x})] = f\,\kappa'(f). \tag{31}$$

The argument to calculate $\kappa'(1)$ is analogous.

I next prove the "only if"-part of the implementability statement. Let T_{ffb} be implementable. I show first that $f^{fb} = 1$. To the contrary, assume T_{ffb} had positive mass $1 - f^{fb} > 0$ on $x = \underline{x}$. I derive a contradiction to the first-best implementability condition $\Pi_I(T_{ffb}) - K \geq 0$ stated in Proposition 2. Indeed, if there is positive mass on $\underline{x}$, then (FB_{supp}) and (FB_{all}) imply

$$\underline{x} - c_{T_{ffb}}(\underline{x}) = \mu \quad and \quad \overline{x} - c_{T_{ffb}}(\overline{x}) \leq \mu. \tag{32}$$

Thus,

$$\begin{aligned}
\Pi_I(T_{ffb}) - K &= \int x - c_{T_{ffb}}(x)\, dT_{ffb} + c_{T_{ffb}}(\underline{x}) - K && (33)\\
&= (1 - f^{fb})(\underline{x} - c_{T_{ffb}}(\underline{x})) + f^{fb}(\overline{x} - c_{T_{ffb}}(\overline{x})) + c_{T_{ffb}}(\underline{x}) - K && (34)\\
&\leq \mu + c_{T_{ffb}}(\underline{x}) - K && (35)\\
&= \underline{x} - K. && (36)
\end{aligned}$$

Because $\underline{x} - K < 0$ by assumption, this is the desired contradiction.

To complete the proof of the "only if"-part, I show that $\overline{x} - \pi_E(1) - K \geq 0$. Indeed, because T_{ffb} is implementable, the inequality $\Pi_I(T_{ffb}) - K \geq 0$ holds by Proposition 2. Since $f^{fb} = 1$ and because $\pi_E(f) = f[c_{T_f}(\overline{x}) - c_{T_f}(\underline{x})]$ by (24), the said inequality simplifies to $\overline{x} - \pi_E(1) - K \geq 0$, as desired.

To see the "if"-part, let $f^{fb} = 1$ and $\overline{x} - \pi_E(1) - K \geq 0$. I show that $T_{ffb} = T_1$ is implementable by verifying the condition $\Pi_I(T_1) - K \geq 0$ from Proposition 2. Indeed, since $\pi_E(f) = f[c_{T_f}(\overline{x}) - c_{T_f}(\underline{x})]$ by (24), one has

$$\begin{aligned}
\Pi_I(T_1) - K &= \int x - c_{T_1}(x)\, dT_1 + c_{T_1}(\underline{x}) - K && (37)\\
&= \overline{x} - c_{T_1}(\overline{x}) + c_{T_1}(\underline{x}) - K && (38)\\
&= \overline{x} - \pi_E(1) - K, && (39)
\end{aligned}$$

which is non-negative by assumption. Thus, $T_{ffb} = T_1$ is implementable by Proposition 2.

Part 2.: All arguments are in the main text except for the claim that the solution to R is a solution to the original problem P. With the same arguments as in the proof of Proposition 5 (part 2.),

it follows that the solution F^* to the relaxed problem is also a solution to the original problem if and only if

$$\int x - c_{F^*}(x)\, dF^* - K \leq \min_{x \in supp(F^*)} (x - c_{F^*}(x)). \tag{40}$$

Recall that $F^* = T_{f^*}$ for some $f^* \in [0,1]$. If $f^* \in \{0,1\}$, that is, the support of F^* is a singleton, then the inequality is obvious. For $f^* \in (0,1)$, condition (40) writes

$$(1-f^*)[\underline{x} - c_{T_{f^*}}(\underline{x})] + f^*[\overline{x} - c_{T_{f^*}}(\overline{x})] - K \leq \min\{\underline{x} - c_{T_{f^*}}(\underline{x}),\ \overline{x} - c_{T_{f^*}}(\overline{x})\}. \tag{41}$$

I now distinguish two cases. First, suppose T_{f^*} coincides with a first-best distribution. Recall from (FB_{supp}) and (FB_{all}) that in this case, $x - c_{T_{f^*}}(x)$ is maximized—and in particular takes on the same value—on the support of T_{f^*}. Thus, (41) is equivalent to $-K \leq 0$ and therefore holds.

Second, suppose T_{f^*} is not a first-best distribution. Then the constraint in (15) is binding, and, by (24), reads

$$(1-f^*)\underline{x} + f^*\overline{x} - f^*[c_{T_{f^*}}(\overline{x}) - c_{T_{f^*}}(\underline{x})] - K = 0. \tag{42}$$

Therefore, the left-hand side of (41) is equal to $-c_{T_{f^*}}(\underline{x})$. To determine the right-hand side, (42) can be rearranged to

$$\underline{x} - c_{T_{f^*}}(\underline{x}) + \frac{K - \underline{x}}{f^*} = \overline{x} - c_{T_{f^*}}(\overline{x}). \tag{43}$$

Since $\underline{x} < K$ by assumption, this implies that

$$\min\{\underline{x} - c_{T_{f^*}}(\underline{x})), \overline{x} - c_{T_{f^*}}(\overline{x})\} = \underline{x} - c_{T_{f^*}}(\underline{x}). \tag{44}$$

Therefore, the right-hand side of (41) is equal to $\underline{x} - c_{T_{f^*}}(\underline{x})$. Because the left-hand side is equal to $-c_{T_{f^*}}(\underline{x})$, and $0 \leq \underline{x}$, (41) is shown, and this establishes Part 2.

Part 3.: To see that $f^* \leq f^{fb}$, I show below that κ is convex and that $\kappa'(f) \leq \pi_E'(f)$ for all f. The claim then follows from the same formal arguments that establish that $x^* \leq x^{fb}$ in the proof of Proposition 6, where now $\alpha(f) = (1-f)\underline{x} + f\overline{x} - \kappa(f)$, and $\beta(f) = (1-f)\underline{x} + f\overline{x} - \pi_E(f) - K$.

To see that κ is convex, recall that $\kappa(f) = C(T_f)$. Now, observe that for a convex combination

$\tau f + (1-\tau)g$, $\tau \in [0,1]$, it holds:

$$T_{\tau f+(1-\tau)g} = \tau T_f + (1-\tau)T_g. \tag{45}$$

Therefore, the convexity of κ follows from the convexity of C. To see that $\kappa'(f) \leq \pi'_E(f)$ for all f, recall that $\pi_E(f) = f\kappa'(f)$, and thus

$$\pi'_E(f) = \kappa'(f) + \kappa''(f)f \geq \kappa'(f), \tag{46}$$

as desired. QED

Proof of Proposition 8 The arguments for Part 1. and Part 2. of the statement are in the main text.

As to Part 3., note first that problem (16) is structurally identical to problem (12), because $\breve{\varphi}$ is convex by definition. The proof that $M^* \leq M^{fb}$ is thus identical to the proof that $x^* \leq x^{fb}$ in Proposition 6.

To see the second statement of Part 3., I construct, for each mean M, a distribution G_M that minimizes $\int \varphi(x)\, dF$ among all distributions with mean M, and show that G_M is increasing in M in the sense of first-order stochastic dominance.

Indeed, recall that $\breve{\varphi}$ denotes the lower convex envelope of φ. For any M define the interval

$$I(M) = \begin{cases} [M,M] & \text{if} \quad \varphi(M) = \breve{\varphi}(M) \\ [a,b] & \text{if} \quad \varphi(M) > \breve{\varphi}(M) \end{cases}, \tag{47}$$

where in the second case, (a,b), $a < b$, is the connected component of $\{\varphi > \breve{\varphi}\}$ containing M.

Moreover, let G_M be the distribution supported on the boundary points of $I(M)$ with mean M. It follows from construction that $\int \varphi(x)\, dG_M = \breve{\varphi}(M)$. Hence, G_M minimizes the moment among all distributions F with mean M, because

$$\int \varphi(x)\, dG_M = \breve{\varphi}(M) \leq \int \breve{\varphi}(x)\, dF \leq \int \varphi(x)\, dF, \tag{48}$$

where the first inequality follows from Jensen's inequality.

It remains to observe that this family of distributions can be chosen to be ordered by first-order

stochastic dominance. Consider $M_1 < M_2$. If $I(M_1) = I(M_2) = [a, b]$, then because $M_1 < M_2$, the interval is non-degenerate, and G_{M_1} and G_{M_2} both have support $\{a, b\}$. The probability assigned to the upper point b is $\frac{M-a}{b-a}$, which is increasing in M. Hence G_{M_2} first-order stochastically dominates G_{M_1}.

If $I(M_1) = [a_1, b_1] \neq I(M_2) = [a_2, b_2]$, then $b_1 \leq a_2$. This follows, because the non-degenerate intervals $I(M)$ are the closures of distinct connected components of $\{\varphi > \breve{\varphi}\}$, which are disjoint and ordered, while a degenerate interval $I(M) = [M, M]$ has $\varphi(M) = \breve{\varphi}(M)$ and therefore cannot lie in the interior of any such component. Therefore, every point in the support of G_{M_1} is weakly smaller than every point in the support of G_{M_2}. Hence, also in this case, G_{M_2} first-order stochastically dominates G_{M_1}.

Now take $F^* = G_{M^*}$ and $F^{fb} = G_{M^{fb}}$. By Parts 1. and 2., these distributions are, respectively, a solution to R and a first-best distribution. Since $M^* \leq M^{fb}$, the argument above implies that F^{fb} first-order stochastically dominates F^*. Both distributions have at most two points in their support, as desired.

As to Part 4., I show that the solution to R is a solution to the original problem P under the stated sufficient conditions. Indeed, the same arguments as in the proof of Proposition 5 apply to show that a solution F^* to the relaxed problem is also a solution to the original problem if and only if

$$\int x - c_{F^*}(x)\, dF^* - K \leq \min_{x \in supp(F^*)} (x - c_{F^*}(x)). \tag{49}$$

If F^* is a degenerate distribution, the inequality is evidently always true. For the case that F^* coincides with a first-best distribution, recall from (FB_{supp}) and (FB_{all}) that $x - c_{F^*}(x)$ is maximized—and in particular takes on the same value—on the support of F^*. Thus, (49) is equivalent to $-K \leq 0$ and therefore holds.

Thus, suppose that F^* has two points $\{x_1, x_2\}$ in its support with $f^* = Prob(x_2)$, and is not first-best. Condition (49) writes

$$\begin{aligned} & (1 - f^*)[x_1 - \Gamma'(\Phi_{F^*})\varphi(x_1)] + f^*[x_2 - \Gamma'(\Phi_{F^*})\varphi(x_2)] - K \\ \leq \;& \min\{x_1 - \Gamma'(\Phi_{F^*})\varphi(x_1),\; x_2 - \Gamma'(\Phi_{F^*})\varphi(x_2)\}, \end{aligned} \tag{50}$$

where I have used that for single-moment-based costs, $c_F(x) = \Gamma'(\Phi_F)\varphi(x)$ and $\Phi_{F^*} = \breve{\varphi}(M^*)$ at

an optimum.

Because F^* is not first-best, the constraint in (16) is binding (otherwise, moving slightly toward a first-best distribution would be feasible by continuity and increase the objective):

$$M^* - \Gamma'(\Phi_{F^*})\Phi_{F^*} - K = \int x - \Gamma'(\Phi_{F^*})\varphi(x)\, dF^* - K = 0. \tag{51}$$

Therefore, the left hand side of inequality (50) is equal to zero. Accordingly, inequality (50) holds if and only if

$$x_i - \Gamma'(\breve{\varphi}(M^*))\varphi(x_i) \geq 0, \quad i = 1, 2. \tag{52}$$

To complete the proof, I show that a sufficient condition for (52) is

$$\varphi'(x) \leq \frac{1}{\Gamma'(\varphi(\overline{x}))} \quad \forall x \in X. \tag{53}$$

Indeed, recall that $\varphi(\underline{x}) = 0$. Thus a sufficient condition for (52) is that $x - \Gamma'(\breve{\varphi}(M^*))\varphi(x)$ is increasing in x, that is,

$$1 - \Gamma'(\breve{\varphi}(M^*))\varphi'(x) \geq 0 \quad \forall x \in X. \tag{54}$$

This condition is implied by (53) because Γ' is increasing by assumption and because $\breve{\varphi}(M^*) \leq \varphi(M^*) \leq \varphi(\overline{x})$ by definition of the lower convex envelope and because φ is increasing by assumption. And this completes the proof. QED

Proof of Proposition 9 All arguments are in the main text except for the argument establishing that the solution to R is also a solution to the original problem P. The argument for this is analogous to the corresponding argument in Part 4. of the proof of Proposition 8. QED

References

Allen, Franklin, and Adelina Barbalau (2024). Security design: A review. Journal of Financial Intermediation 60, 101113.

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

Biais, Bruno, and Catherine Casamatta (1999). Optimal leverage and aggregate investment. Journal of Finance 54.4, 1291-1323.

Castro-Pires, Henrique, Deniz Kattwinkel, and Jan Knöpfle (2025). Who and how? Adverse selection with flexible moral hazard. Mimeo.

Cerreia-Vioglio, Simone, Fabio Maccheroni, and Massimo Marinacci (2017). Stochastic dominance analysis without the independence axiom. Management Science 63(4), 1097-1109.

Georgiadis, George, Doron Ravid, and Balazs Szentes (2024). Flexible moral hazard problems. Econometrica 92(2), 387-409.

Grossman, Sanford J. and Oliver D. Hart (1983). An analysis of the principal-agent problem. Econometrica 51(1), 7-49.

Hébert, Benjamin (2018). Moral hazard and the optimality of debt. Review of Economic Studies 85(4), 2214-2252.

Hellwig, Martin F. (2009). A reconsideration of the Jensen-Meckling model of outside finance. Journal of Financial Intermediation 18(4), 495-525.

Hellwig, Martin F. (2025). A principal-agent model with flexible moral hazard concerning risks as well as expected returns. Mimeo.

Hong, Chew Soo, and Naoko Nishimura (1992). Differentiability, comparative statics, and non-expected utility preferences. Journal of Economic Theory 56(2), 294-312.

Innes, Robert D. (1990). Limited liability and incentive contracting with ex-ante action choices. Journal of Economic Theory 52(1), 45-67.

Jensen, Michael C. and William H. Meckling (1976). Theory of the firm: managerial behavior, agency costs and ownership structure. Journal of Financial Economics 3, 305-360.

Kocherlakota, Narayana R. (2026). Optimal labor income taxation: a flexible moral hazard approach. Review of Economic Studies, rdag046.

Krähmer, Daniel (2026a). The hold-up problem with flexible unobservable investments. Journal of Economic Theory. 235, 106205.

Krähmer, Daniel (2026b). Optimal monopoly regulation with flexible investments. Mimeo.

Liu, Siwen (2025). Flexible moral hazard problems with adverse selection. Mimeo.

Machina, Mark J. (1982). "Expected utility" analysis without the independence axiom. Econometrica, 277-323.

Mattsson, Lars-Göran, and Jörgen W. Weibull (2023). An analytically solvable principal-agent model. Games and Economic Behavior 140, 33-49.

Winkler, Gerhard (1988). Extreme points of moment sets. Mathematics of Operations Research 13(4), 581-587.

Yang, Kai Hao, and Alexander K. Zentefis (2024). Monotone function intervals: theory and applications. American Economic Review 114(8), 2239-2270.