

Dynamic Resource Allocation with Hidden Volatility^{*}

Felix Zhiyu Feng
University of Washington

Mark M. Westerfield
University of Washington

February 2020

Abstract

We study a firm's internal resource allocation using a dynamic principal-agent model with endogenous cash flow volatility. The principal supplies the agent with resources for productive use, but the agent has private control over both project volatility and resource intensity and may misallocate resources to obtain private benefits. The optimal contract can yield either overly risky or overly prudent project selection. It can be implemented with a constant pricing schedule (i.e., a static, decentralized, linear mechanism), giving the agent control over the resource quantities, project risk, and agent's equity share. The implementation rationalizes the use of hurdle rates above a firm's cost of capital and transfer prices above marginal cost, while showing that hurdle rates or transfer prices may not vary with the agent's risk choice.

JEL Classification: G31, D86, G34

Key Words: capital budgeting, transfer pricing, dynamic contracting, volatility control, cost of capital

^{*}We thank Bruno Biais, George Georgiadis, Will Gornall, Raymond C. W. Leung, Adriano Rampini, Jean-Charles Rochet, Bruno Strulovici, Curtis Taylor, Lucy White, Yao Zeng, Haoxiang Zhu, seminar and conference participants at Boston University, Cheung Kong Graduate School of Business, Chinese University of Hong Kong, Duke University, Tsinghua PBC School of Finance, University of British Columbia, University of Minnesota, University of Notre Dame, University of Texas-Dallas, the Finance Theory Group Spring 2018 meeting, the CICF, the Midwest Economic Theory conference and the Econometric Society North American Summer Meeting for helpful comments and suggestions. We also thank an anonymous referee and the editor (Ron Kaniel) for especially helpful reviews.

Feng can be reached at ffeng@uw.edu; Westerfield can be reached at mwesterf@uw.edu

1 Introduction

This paper analyzes dynamic resource allocation (capital or intermediate goods) within a firm. We focus on two interacting frictions. First, the divisional manager can hide his choice of production risk from headquarters. For example, he can choose a production process more subject to the vagaries of market demand, quality deviations, or technological risk. Second, the divisional manager consumes benefits from unproductive investments or unused resources. For example, he can obtain private benefits from financial slack, from executing fun but unproductive projects, or from diverting resources. Headquarters observes divisional output, but not the production technology or its inputs. Thus, the firm has a multi-tasking agency problem over inputs and risk.

The interaction of our two frictions, hidden risk choice and hidden resource allocation, allows for a moral hazard problem over volatility. Organizational settings in which project managers privately influence both the amount and riskiness of investment are ubiquitous. Nevertheless, the design of dynamic incentive plans that both allocate the appropriate amount of capital and induce managers to choose the desired amount of volatility has challenged formal analysis. In our model, the agent's ability to hide inputs complements his ability to hide risk, and we focus on the combination.

We solve for the optimal contract in a general, dynamic principal-agent setting, and show that it can be implemented with a simple scheme: a decentralized and static transfer pricing schedule for intermediate goods or capital. In a standard recursive optimal contract, the principal would allocate resources to the agent and command a certain level of total cash flow risk. Our implementation allows the agent to choose his own level of resources, the cash flow risk target, and even his own compensation structure (his pay-performance sensitivity or equity share). The principal simply supplies the requested resources to the agent at a cost based on the agent's announced choices.¹ The resource cost is subtracted from the project's output, and the remainder is split between the principal and agent based on the proposed equity share.

Crucially, the contract implementation is a static schedule of prices that depend only on the agent's current choices. The principal does not need to control any stochastic processes nor even keep track of the agent's performance history. In that sense, this is a theory of institutional design, rather than active management. By choosing his own level of resources,

¹For a concrete example, imagine the principal tells the agent: "Take however much capital you want. If you want a 10% equity share, your cost of capital is 8%. If you want a 20% equity share, your cost of capital is 7%." A higher equity share is equivalent to higher pay-for-performance incentives.

cash flow risk target, and compensation structure, the agent is effectively picking from the set of incentive compatible contracts. This happens because the capital or intermediate goods prices are calculated exactly so that the agent is indifferent across choices; all the available choices are incentive compatible and exactly meet the agent’s participation constraint. Being indifferent, the agent picks the principal’s desired set of resources, risk target, and compensation structure.

The implementation has several implications: first, interpreting resources as capital, we rationalize observed practices of capital budgeting. As documented by [Jagannathan, Matsa, Meier, and Tarhan \(2016\)](#), [Graham and Harvey \(2001\)](#), and others,² firms impose capital rationing on managers through a deliberately high cost of capital (hurdle rate), and they most commonly do not adjust the cost of capital for risk. These stylized facts contradict conventional wisdom and textbook corporate finance instruction. In our model, headquarters optimally raises the cost of capital above the replacement rate to limit the scale of the manager’s investment and the size of the agency problem, passing up otherwise positive NPV projects.

More surprisingly, our implementation rationalizes the practice of setting costs of capital without adjusting for cash flow risk. We obtain this result because the manager’s incentive compatibility condition equalizes the marginal product of capital in productive and unproductive projects. Thus, the cost of capital can be set equal to the marginal product of capital in unproductive projects. This can be done without further adjusting for the returns and risk in the productive projects, and the desired capital allocation follows.

As a matter of contract form, our cost of capital (hurdle rate) can be interpreted as a preferred return to investors. These preferred returns are standard in private equity contracts (see, e.g., [Metrick and Yasuda, 2010](#); [Robinson and Sensoy, 2013](#)). Alternatively, the hurdle rate can be interpreted as a type of Funding Value Adjustment or a Capital Value adjustment, as described in [Andersen, Duffie, and Song \(2019\)](#). In these arrangements, a previously decided cost of capital is deducted from the profit and loss statement of an operation before gains are distributed.

Similarly, interpreting resources as intermediate goods, our study rationalizes cost-based transfer pricing with a markup above marginal cost. While the seminal work of [Hirshleifer \(1956\)](#) argues that pricing at the marginal cost of production achieves optimality for the

²In Section 5.1, we summarize and condense the findings from [Jagannathan, Matsa, Meier, and Tarhan \(2016\)](#), [Graham and Harvey \(2001, 2011, 2012\)](#), [Jacobs and Sivdasani \(2012\)](#), and [Poterba and Summers \(1995\)](#).

headquarters, the majority of firms that practice cost-based pricing use a markup above the production cost (PriceWaterhouse, 1984; Tang, 1992). Existing models that rationalize the use of such a policy often restrict their analysis to a particular set of mechanisms or incomplete contracts; the resulting optimal transfer scheme is usually non-linear and must be administered and monitored by a central planner.³ In contrast, our transfer pricing scheme is static, linear, decentralized, and does not require dynamic monitoring because the principal does not need to actively manage resources or track the agent’s performance history. Our paper further demonstrates that resource rationing is an endogenous choice of the principal: she has unlimited access to the external markets, but engages in deliberate rationing of divisional resources anyway.

In more detail, we begin with a basic dynamic moral hazard framework in which the principal (she) has assets or projects that she hires the agent (he) to manage. The agent takes hidden actions that determine the cash flow from these projects; the principal observes output and rewards the agent with consumption after high output and imposes termination after low output. Separation is costly, which makes the principal effectively risk averse. Our novel assumption is that the agent receives capital or input goods from the principal but privately decides both the risk of the project and the amount of inputs that are actually invested. Any unused resources are allocated to generate private benefits for the agent. The resources at issue are liquid, and can be reallocated freely. This joint modeling of resource intensity and risk circumvents the issue that volatility is observable for Brownian motions, which has been the primary obstacle precluding the analysis of dynamic risk choice in a tractable framework. Thus, in our setting the principal must provide incentives to generate both the desired risk choice and the desired resource intensity to avoid misallocation. The agent’s choice is constrained by the observability of total risk/volatility; it is the *components* of risk/volatility that are unobservable and subject to an agent’s manipulation.

We show that the resulting optimal contract can lead to both overly risky or overly prudent risk choice, relative to the first-best. This contrasts with the standard result under full risk/volatility observability (see, e.g., DeMarzo and Sannikov, 2006; Biais, Mariotti, Plantin, and Rochet, 2007) that an effectively risk averse principal would strictly prefer lower volatility. In our paper, to reduce the likelihood of costly termination following poor performance, the optimal contract reduces the volatility of the agent’s continuation value. This volatility

³We discuss the relevant transfer pricing literature in more detail in Section 5.2, including the seminal mechanism design approach in Harris, Kriebel, and Raviv (1982) and subsequent studies such as Rajan and Reichelstein (2004), Hiemann and Reichelstein (2012), Hyde and Choe (2005), Sahay (2003) and Baldenius, Dutta, and Reichelstein (2007).

has two components: the volatility of the project’s cash flow and the agent’s exposure to it (his pay-performance sensitivity, or PPS). Exactly which component the principal reduces depends on the specific risk-return relationship, which determines how much additional risk the agent takes as incentives are made less intense. When the risk taken by the agent is very sensitive to incentives, the principal offers more incentives, increasing PPS and reducing risk below the first-best level. In contrast, when the risk taken by the agent is relatively insensitive to incentives, the principal relaxes incentives, reducing PPS and resulting in an overly risky project choice. In contrast to the risk adjustment, capital intensity is always (weakly) less than the first-best because more intensive use of capital implies higher cash flow volatility and requires higher pay-performance sensitivity.⁴

The basic framework of our model borrows from [DeMarzo and Sannikov \(2006\)](#) and [Biais, Mariotti, Plantin, and Rochet \(2007\)](#).⁵ Two other studies that examine volatility control in continuous time are [Cvitanic, Possamaï, and Touzi \(2017\)](#) and [Leung \(2017\)](#). [Cvitanic, Possamaï, and Touzi \(2017\)](#) assesses optimal control over a multi-dimensional Brownian motion when the contract includes only a terminal payment and is sufficiently integrable. They show that the principal can attain her optimal value (possibly in a limit) by maximizing over contracts that depend only on output and quadratic variation. The setup is similar to earlier work by [Cadenillas, Cvitanic, and Zapatero \(2004\)](#) and more broadly to the literature on delegated portfolio control such as [Carpenter \(2000\)](#), [Ou-Yang \(2003\)](#), and [Lioui and Poncet \(2013\)](#), which focus on exogenous compensation and/or information structures. [Leung \(2017\)](#) assumes that cash flow is made up of the agent’s private choice of project risk and an exogenous market factor that is unobservable to the principal. However, because the market factor in [Leung \(2017\)](#) is exogenous, after the contract is in place the agent has only limited ability to manipulate risk without being detected. Our paper is also broadly related to [Biais, Mariotti, Rochet, and Villeneuve \(2010\)](#), [DeMarzo, Livdan, and Tchisty \(2013\)](#), [Li and Williams \(2017\)](#), and others that study optimal incentives when the agent’s action generates discrete, verifiable jump risks. Finally, [Epstein and Ji \(2013\)](#) develop a volatility control model based on an ambiguity problem. In contrast, with standard preferences, we

⁴Our results help reconcile empirical evidence regarding the correlation between investment risk and pay-performance sensitivity, which has been particularly ambiguous among existing studies. [Prendergast \(2002\)](#) and [Edmans, Gabaix, and Jenter \(2017\)](#) summarize the related theoretical and empirical studies. See [Section 5.3](#) for a more detailed discussion of this line of research.

⁵Other papers that investigate agency problems and capital usage in the same model include [He \(2011\)](#), [DeMarzo, Fishman, He, and Wang \(2012\)](#), and [Malenko \(2019\)](#). We add to those papers by modeling an agency problem over resource intensity and the productivity of production inputs, as opposed to over mean cash flow or growth.

study the design of an optimal contract in an intuitive agency environment and show that our contract can be implemented with a simple structure resembling the observed practice of internal resource allocation.

2 The Model

In this section, we introduce a model of resource allocation and production within a firm. A division of the firm uses fixed assets and a flow of resources provided by headquarters to generate risky cash flows. The resources from headquarters can be liquid working capital, in which case we are modeling internal capital markets, or intermediate goods such as microchips or service hours, in which case we are modeling transfer pricing within a firm.

There are two important frictions. First, the division makes a hidden choice of production technology, in particular a hidden choice of risk or cash flow volatility. Second, the division can use the resources to generate private benefits.

Our goal is to obtain an optimal scheme that efficiently uses resources while controlling risk taking. To this end, we assume a *general, dynamic* contracting environment in which headquarters has a great deal of control and flexibility. We show that the optimal contract can actually be implemented with a set of *simple, static* rules.

2.1 The Basic Environment

There is a principal (she, headquarters) that has access to resources and an agent (he, the divisional manager) that has the skills to manage profitable projects. Both the principal and the agent are risk neutral. Time is continuous. The principal has unlimited liability, the ability to cover project costs, and a discount rate r . The agent has limited liability, no ability to access resources on his own, and a discount rate $\gamma > r$. The fixed assets of production have a value outside the division $L > 0$, and the agent has an outside option or value $R = 0$.⁶ The agent cannot borrow or save. A contract allows the principal to fund the agent's projects so that they share the surplus created.

The agent has a choice of production technologies indexed by volatility σ , such that $0 \leq \underline{\sigma} \leq \sigma \leq \bar{\sigma}$. Given a level of volatility σ_t and a flow of resources $K_t \geq 0$, the agent's

⁶Assuming $L > 0$ and $R = 0$ simplifies the exposition because the principal does not temporarily shut down production by granting the agent zero resources. In Section 6, we extend the model to include more general outside values.

production choice generates a cumulative cash flow Y_t that evolves as

$$dY_t = f(K_t) [\mu(\sigma_t)dt + \sigma_t dZ_t], \quad (1)$$

where Z_t is a standard Brownian motion. $\mu(\sigma)$ represents the efficient frontier: the highest expected cash flow that the agent can obtain given a particular level of volatility. Both σ_t and K_t can be instantaneously adjusted without cost; K_t represents liquid working capital, such as cash, machine-hours, etc., or intermediate goods like microchips or service hours.

Without the loss of generality, we assume the principal can raise an arbitrary amount of K_t with a flow cost rK_t ; i.e., we define a *unit* of K to be the amount that has a cost r . If K is consumed in production – K represents an intermediate good – then rK is the cost to headquarters of obtaining the flow of resources. If K is not consumed in production – K represents liquid capital that is returned after use – then rK is the opportunity cost to headquarters for the use of the resources.

Our specification of the cash flow process is designed to capture the notion that the agent is drawing from a limited selection of projects based on real assets to generate cash flows. Because the agent has a limited selection of underlying projects, each additional input generates decreasing cash flow output. We capture this with a decreasing-returns-to-scale technology:⁷ $f(K)$ is three-times differentiable with $f(0) = 0$; $f'(K) > 0$; $f''(K) < 0$; $\lim_{K \rightarrow 0} f'(K) = \infty$; $\lim_{K \rightarrow \infty} f'(K) = 0$. Because the agent's projects are based on real assets, as opposed to financial assets, they cannot profitably sustain infinite volatility. Instead, some risk taking is efficient, but beyond a certain point increasing volatility reduces average cash flows. We capture this with a hump-shaped returns technology: $\mu''(\sigma) < 0$, $\mu(\sigma)$ attains its maximum inside $(\underline{\sigma}, \bar{\sigma})$, and $\mu(\bar{\sigma}) < 0$.⁸

Our specification is flexible and allows a well-defined, interior first-best level of resource

⁷Allowing concavity in the production function can be more realistic than linearity (e.g., allowing for organizational frictions like a limited span of control), and the assumption generates a first-best with finite resource usage. Because our principal and agent are risk neutral, the first-best is achieved after some histories (see Section 3). In contrast, a linear production function (e.g., $f(K) = K$) implies infinite first-best resource usage. To compensate, we would need to make the agent risk averse, as in Sannikov (2008). Risk aversion can generate a principal's value function that is strictly concave, so that the first-best level of resources is never implemented.

⁸For example, if the agent is choosing among underlying projects each with a normally distributed returns and a linear cost of volatility, then standard results imply $\mu(\sigma) = \underline{\mu} + C\sqrt{\sigma^2 - \underline{\sigma}^2} - b\sigma$. If instead the agent faces a constant return to risk minus a convex volatility penalty, then we could have $\mu(\sigma) = b\sigma - \sigma^\alpha$.

usage and volatility $\{K^{FB}, \sigma^{FB}\}$, which are given by

$$\max_{K \geq 0; \sigma \geq \underline{\sigma}} [f(K)\mu(\sigma) - rK] \quad (2)$$

and characterized by the first-order conditions $0 = \mu'(\sigma^{FB})$ and $r = f'(K^{FB})\mu(\sigma^{FB})$.

2.2 The Agency Friction

The principal supplies resources K_t to the agent and a recommended level of project volatility σ_t . The agent chooses two hidden actions. First, the agent can allocate resources to productive projects or to an unproductive project that generates private benefits. The actual amount allocated to productive projects is $\hat{K}_t$ and is used to generate dY_t . The remaining resources, $K_t - \hat{K}_t$, are allocated to a zero-cash flow project that generates a flow of private benefits $\lambda(K_t - \hat{K}_t)dt$. Resource misallocation is (weakly) inefficient: resources cannot be used to generate private benefits in excess of their cost, $0 < \lambda \leq r$. The principal's incentive problem is to create a contract that induces the agent to allocate resources according to the principal's recommendation.

In developing the optimal contract, we initially restrict attention to contracts that implement zero misallocation. (i.e., $K = \hat{K}$ and private benefits are zero). This is without loss of generality as long as misallocation is inefficient ($\lambda \leq r$), which we show in Proposition 3. We discuss several generalizations of the model in Section 6.

Our agency problem can be interpreted in several different ways:

- Choosing $\hat{K}_t < K_t$ means shifting resources to enjoyable but unproductive projects. Two examples are a manager with a desire for the quiet life (e.g., [Bertrand and Mullainathan, 2003](#)) and a manager who prefers not to travel to make site inspections (e.g., [Giroud, 2013](#)).
- In Section 4, we demonstrate an implementation of the optimal contract that allows the agent to request any level of resources (K) from the principal and allocate some of it ($\hat{K}_t < K_t$) to unproductive projects. In that case, the agency friction can also be naturally interpreted as empire-building, since the manager enjoys the control of more resources than are needed for productive uses.
- A manager may need to exert effort to maintain the production efficient frontier ($\mu(\sigma)$) at a high level. For example, a manager might shirk by considering a smaller number

of potential projects. In doing so, he gains private benefits from shirking $\lambda\Delta_t$, and the efficient investment frontier is reduced to $f(K_t - \Delta_t)(\mu(\hat{\sigma}_t)dt + \hat{\sigma}_t dZ_t)$. Here, Δ_t plays the role of $K_t - \hat{K}_t$.

The cash flow process Y is observable to the principal; headquarters can fully observe profits and losses. As a practical matter, if headquarters can observe the history of the division's general accounting ledger, then the division is unable to lie about cash flow volatility by too much. In our model, the principal can exactly infer the overall volatility because of the continuous nature of the shocks (dZ_t).⁹ As a result, the principal can impose a particular level of overall volatility (e.g., by terminating the agent if the proper level is not observed). We make the more direct assumption that the principal simply controls the total cash flow volatility from (1), labeled Σ_t , with

$$\Sigma_t \equiv f(K_t)\sigma_t = f(\hat{K}_t)\hat{\sigma}_t. \quad (3)$$

The second equality is the constraint that the agent must achieve a total volatility with his hidden choices that is consistent with the principal's recommended choices.

The agency friction in our model comes from the fact that the principal does not observe the *source* of volatility: intensive resource use in productive projects or excessively risky projects. Put differently, the agent can generate the appearance of productive activity (cash flow volatility) while still using resources to generate private benefits. The agent can allocate $K_t - \hat{K}_t$ to the unproductive project while increasing volatility in the productive project ($\hat{\sigma} > \sigma_t$), keeping aggregate volatility (Σ_t) constant. In so doing, the agent enjoys total private benefits $\lambda(K_t - \hat{K}_t)$. Thus, the principal provides an incentive contract to induce the agent to choose the desired *components* of volatility; the agent must be induced not to take bad risks that hide bad resource allocation.

Remark 1 *Existing models of hidden risk taking generally focus on jump risk. One class of models links increased jump risk to private benefits, such as [Biais, Mariotti, Rochet, and](#)*

⁹For concreteness, consider a heuristic example: a principal observes $dY_t = a_t dt + b_t dZ_t$ with Y_0 known and $b_t > 0$. The principal does not observe either a_t or b_t directly. The ability to observe the path of Y implies the ability to observe the path of Y^2 . Since $d(Y_t^2) - 2Y_t dY_t = b_t^2 dt$, the principal is able to infer b_t .

An analogous constraint on observability holds in discrete time binomial models. In a binomial model, cash flows could be high ($+\Sigma = +f(K)\sigma$) or low ($-\Sigma = -f(K)\sigma$), with high cash flows realized with probability $p(\sigma)$. Then, the total cash flow at stake $\pm f(K)\sigma$ is observable to the principal after either a high or low outcome. However, the source of the cash flow at stake, σ or $f(K)$, would remain hidden because Σ is observable to the principal but $p(\sigma)$ is not. Thus, the agent could misallocate K to hide bad risk choice (choice of σ).

Villeneuve (2010) and Feng (2020). These models share the feature that stronger (down-side) incentives reduce jump risk and shirking. A second class of models links increased jump-risk to extra returns, such as DeMarzo, Livdan, and Tchistyi (2013), Moreno-Bromberg and Roger (2017), Li and Williams (2017), and Wong (2019). These models share the feature that stronger pay-for-performance incentives encourage jump risk, and, if there is an additional effort process, lower private benefits from shirking. The principal then imposes a second set of incentives to reduce jump risk.

Our model works differently from those just mentioned because we consider a moral hazard problem over the components of volatility. Excess volatility is linked to private benefits at the project level, despite the fact that private benefits have no direct effect on cash flow volatility. Instead, the effect is indirect: the agent allocates resources to the unproductive project and compensates with an excessively volatile productive project choice. Despite the direct/indirect distinction, our model has a positive association between excessive risk taking and obtaining private benefits, so it also has the property that stronger incentives reduce project volatility (see Section 3.1).¹⁰

2.3 The Objective Functions

Contracts in our model are characterized using the agent's continuation utility as the state variable. Denote the probability space as $(\Omega, \mathcal{F}, P)$, and the filtration as $\{\mathcal{F}_t\}_{t \geq 0}$ generated by the cash flow history $\{Y_t\}_{t \geq 0}$. Contingent on the filtration, a contract specifies a payment process $\{C_t\}_{t \geq 0}$ to the agent, a stopping time τ when the contract is terminated, a sequence of resources $\{K_t\}_{t \geq 0}$ under the agent's management, and a sequence of recommended volatility levels $\{\sigma_t\}_{t \geq 0}$. $\{C_t\}_{t \geq 0}$ is non-decreasing because the agent is protected by limited liability. All quantities are assumed to be integrable and measurable under the usual conditions.

Given a contract, the agent chooses a given set of policy rules $\{\hat{K}_t, \hat{\sigma}_t\}_{t \geq 0}$. The agent's objective function is the expected discounted value of consumption plus private benefits

$$W_t^{\hat{K}, \hat{\sigma}} = E_{\hat{K}, \hat{\sigma}} \left[\int_t^\tau e^{-\gamma(s-t)} \left(dC_s + \lambda(K_s - \hat{K}_s) ds \right) + e^{-\gamma\tau} R \middle| \mathcal{F}_t \right], \quad (4)$$

while the principal's objective function is the expected discounted value of the cash flow,

¹⁰In Section 6 and the Internet Appendix, we extend the model so that the private benefits project also generates volatility. In that version, using resources for private benefits forces the agent to reduce productive volatility, meaning private benefits are associated with insufficient *productive* risk taking. The principal thus uses stronger incentives to induce the agent to take more risk.

minus the cost of resources and payments to the agent

$$V_t^{\hat{K}, \hat{\sigma}} = E_{\hat{K}, \hat{\sigma}} \left[\int_t^\tau e^{-r(s-t)} (dY_s - rK_s ds - dC_s) + e^{-r\tau} L \middle| \mathcal{F}_t \right], \quad (5)$$

where both expectations are taken under the probability measure associated with the agent's choices. The optimal contract is defined as:

Definition 1 (Incentive compatible and optimal contracts) *A contract is incentive compatible if the agent maximizes his objective function by choosing $\{\hat{K}_t, \hat{\sigma}_t\}_{t \geq 0} = \{K_t, \sigma_t\}_{t \geq 0}$.*

A contract is optimal if it maximizes the principal's objective function over the set of contracts that 1) are incentive compatible, 2) grant the agent his initial level of utility W_0 ,¹¹ and 3) give $W_t^{K, \sigma} \geq R$.

This definition restricts our analysis to contracts that involve no resource misallocation because we have defined incentive compatible contracts to mean $\hat{K}_t = K_t$. As mentioned earlier, we discuss several generalizations in Section 6 and the Appendix.

3 The Optimal Contract

In this section, we derive the optimal contract. We begin by characterizing the properties of incentive compatible contracts and proceed to the principal's Hamilton-Jacobi-Bellman (HJB) equation. We end with a categorization of contract types and some comparative statics. Our discussion in the text is somewhat heuristic; proofs not immediately given in the text are in the Appendix.

3.1 Continuation Value and Incentive Compatibility

The following proposition summarizes the dynamics of the agent's continuation value W_t as well as the incentive compatibility condition:

Proposition 1 *Given any contract and any sequence of the agent's choices, there exists a predictable, finite process β_t ($0 \leq t \leq \tau$) such that W_t evolves according to*

$$dW_t = \gamma W_t dt - \lambda(K_t - \hat{K}_t)dt - dC_t + \beta_t \left(dY_t - f(\hat{K}_t)\mu(\hat{\sigma}_t)dt \right). \quad (6)$$

¹¹We leave the determination of W_0 outside the model. One commonly used value is the principal's maximum: if the principal's value function is $F(W)$, then $W_0 = \arg \max_W F(W)$.

The contract is incentive compatible if and only if

$$\{K_t, \sigma_t\} = \arg \max_{\hat{K}_t \in [0, K_t]; f(\hat{K}_t)\hat{\sigma}_t = f(K_t)\sigma_t} \left[\beta_t f(\hat{K}_t) \mu(\hat{\sigma}_t) - \lambda \hat{K}_t \right]. \quad (7)$$

If the contract is incentive compatible, then $\beta_t \geq 0$ and

$$\beta_t \geq \beta(\sigma_t, K_t) \equiv \frac{\lambda}{f'(K_t)} \times \frac{1}{\mu(\sigma_t) - \mu'(\sigma_t)\sigma_t}, \quad (8)$$

and (6) simplifies to

$$dW_t = \gamma W_t dt + \beta_t \Sigma_t dZ_t - dC_t. \quad (9)$$

The dynamics of W_t can be derived using standard martingale methods. The first three terms on the right hand side of (6) reflect the promise keeping constraint: because the agent has a positive discount rate, any utility not awarded today must be compensated with increased consumption in the future. The last term contains the agent's incentives, where β_t is the agent's pay-performance sensitivity (PPS): for every dollar of *excess* cash flow, the agent's continuation value changes by β_t dollars. Under an incentive compatible contract $\{\hat{K}_t, \hat{\sigma}_t\} = \{K_t, \sigma_t\}$, so (6) simplifies to (9).

The incentive compatibility condition (7) is a maximization over the discretionary part of the agent's instantaneous payoff. Given the evolution of the agent's continuation value (6), the agent chooses $\hat{K}_t$ and $\hat{\sigma}_t$ to maximize his flow utility:

$$\beta_t E[dY_t] + \lambda(K_t - \hat{K}_t)dt = \underbrace{\beta_t f(\hat{K}_t) \mu(\hat{\sigma}_t) dt}_{\text{Cash flow paid to the agent}} + \underbrace{\lambda(K_t - \hat{K}_t) dt}_{\text{Agent's private benefit}}, \quad (10)$$

subject to two constraints. First, because the agent cannot borrow on his own, he must choose $\hat{K}_t \in [0, K_t]$. Second, because the principal controls aggregate volatility (3), the agent must choose his controls such that $f(\hat{K}_t)\hat{\sigma}_t = f(K_t)\sigma_t = \Sigma_t$. The resulting maximization problem is given in (7).¹²

Condition (8) is the inversion of the maximization in (7). We denote $\beta(\sigma, K)$ the minimum level of PPS necessary for the agent to choose the principal's recommended values of K and

¹²The agent's problem in this model is effectively static: at any time t , the principal decides her optimal K_t and σ_t and uses β_t and Σ_t to implement such choices. This is a standard feature among dynamic contracting models in which the agent's actions have no persistent effect on cash flows. The dynamics of the optimal contract come from the principal's side, which we describe in the next subsection.

σ . The principal in fact chooses $\beta_t = \beta(\sigma_t, K_t)$ because the principal's value function is concave, which we show in the next section.

We now discuss the important features of incentive compatible contracts in our setting. First, the agency friction does not prevent the principal from implementing the first-best outcome. In fact, the principal can do so even without giving the agent a full share of the project's cash flow:

Property 1 *By choosing $\beta_t = \frac{\lambda}{r} \leq 1$ and $K_t = K^{FB}$, the principal implements $\{K^{FB}, \sigma^{FB}\}$. We define $\beta^{FB} \equiv \frac{\lambda}{r} \leq 1$ to be the level of incentives that, when combined with K^{FB} , implement the first-best policies in the second-best problem.*

To implement the first-best, the principal gives the agent the first-best level of resources, so that the marginal value of productive use is r (from Eq. 2). Then, the principal awards the agent a share of cash flows β^{FB} so that the agent's share of the productive use $r\beta^{FB}$ is equal to the marginal benefit from misallocation (λ). Substituting $\beta_t = \frac{\lambda}{r}$ into the agent's problem (7) produces the same outcome as the first-best optimization (2). No additional incentives are needed to implement the first-best level of volatility.

Second, the agency friction requires the principal to be somewhat moderate in her risk taking: she is precluded from implementing very-high or very-low volatility projects.

Property 2 *The principal cannot implement very-low volatility ($\sigma \leq \underline{\sigma} \equiv \arg \max \frac{\mu(\sigma)}{\sigma}$). It must be the case that $\frac{\partial}{\partial \sigma} \frac{\mu(\sigma)}{\sigma} < 0$.*

The principal never implements very-high volatility ($\sigma \geq \bar{\sigma} \equiv \max\{\sigma | \mu(\sigma) = 0\}$).

Our production technology exhibits decreasing returns to volatility ($\mu(\sigma)$ is hump-shaped). For very low levels of σ , production efficiency increases with additional volatility ($\frac{\partial}{\partial \sigma} \frac{\mu(\sigma)}{\sigma} \geq 0$), so the agent always has a profitable deviation. For any Σ fixed by the principal, increasing $\hat{\sigma}$ can both increase average cash flows and free up resources to be used for private benefits. Thus, the very-low volatility projects with $\frac{\partial}{\partial \sigma} \frac{\mu(\sigma)}{\sigma} > 0$ are never incentive compatible. Further, the principal never chooses to implement a very-high value of $\sigma > 0$ that generates negative expected cash flow. This follows from the fact that the principal can always generate zero cash flow with zero volatility by giving the agent zero resources.

Third, for moderate volatility projects, $\sigma_t \in (\underline{\sigma}, \bar{\sigma})$, stronger incentives are used to increase resource intensity and decrease volatility:

Property 3 For $\beta(\sigma, K)$ defined in (8), we have

$$\frac{\partial}{\partial \sigma} \beta(\sigma, K) = \frac{1}{\lambda} f'(K) \beta(\sigma, K)^2 \sigma \mu''(\sigma) < 0 \quad (11)$$

$$\frac{\partial}{\partial K} \beta(\sigma, K) = -\frac{f''(K)}{f'(K)} \beta(\sigma, K) > 0. \quad (12)$$

Since $\frac{\partial}{\partial \sigma} \beta(\sigma, K) < 0$ and $\frac{\partial}{\partial \sigma} \frac{\mu(\sigma)}{\sigma} < 0$ in the implementable region, stronger incentives are used to increase the efficiency of risk taking (the Sharpe ratio of cash flows).

Intuitively, we have assumed that there is a shortage of good projects but no shortage of bad projects, so our production technology exhibits decreasing returns to scale but our private benefits technology is linear. Thus, the marginal return to the productive use of resources ($f'(K)$) is decreasing, while the marginal value of resources in generating private benefits (λ) is constant. So, stronger incentives are needed to induce the agent to retain resources for productive purposes when projects are already large ($\beta_K > 0$). Agency acts to exacerbate decreasing returns to scale, as it is more difficult to prevent a given level of misallocation when there is a large amount of resources in use.

Similarly, the marginal return to productive projects is decreasing in volatility ($\mu''(\sigma) < 0$), which implies there is an increasing cash flow cost to the agent for excess risk taking. Because of the total volatility constraint, $f(\hat{K})\hat{\sigma} = \Sigma$, the agent takes excessive risks to conceal misallocation. Since the cash flow cost of excessive volatility is high when volatility is already high, weaker incentives can be used, and $\beta_\sigma < 0$.

We illustrate the solution to the agent's problem as well as the features we discussed above in Figs. 1 and 2. Fig. 1 depicts the minimum level of incentives ($\beta(\sigma, K)$) required for any given level of σ and K , as well the resulting expected cash flow to the principal, $CF(\sigma, K) = f(K)\mu(\sigma) - rK$.

With an abuse of notation, we invert (3) and (8) to find $K(\beta, \Sigma)$ and $\sigma(\beta, \Sigma)$. These are the level of resources K and project risk σ that are implemented by the principal's controls: incentive β and overall cash flow volatility Σ . Fig. 2 depicts $K(\beta, \Sigma)$ and $\sigma(\beta, \Sigma)$. Consistent with Properties 2 and 3, stronger incentives implement more resources with less risk, and projects with too little risk cannot be implemented without infinitely high β . We write cash flows from varying levels of β and Σ as $CF(\beta, \Sigma) = f(K(\beta, \Sigma))\mu(\sigma(\beta, \Sigma)) - rK(\beta, \Sigma)$. Cash flow is a concave function in both controls. In particular, stronger incentives do not necessarily lead to a larger cash flow because they may implement an inefficiently low volatility.




Figure 1: The top-left plot displays $\beta(\sigma, K)$ as a function of σ for $K = K^{FB} = 8.56$ (solid line), $K = \frac{2}{3}K^{FB}$ (dashed line), and $K = \frac{1}{3}K^{FB}$ (dotted line). The top-right plot displays $\beta(\sigma, K)$ as a function of K for $\sigma = \frac{3}{2}\sigma^{FB}$ (solid line), $\sigma = \sigma^{FB} = 0.12$ (dashed line), and $\sigma = \frac{1}{2}\sigma^{FB}$ (dotted line). The bottom two plots display average cash flows $CF(K, \sigma) = f(K)\mu(\sigma) - rK$ as a function of σ for each of the three values of K used in the top row, and as a function of K for each of the three values of σ used in the top row. These plots are generated using $\lambda = 0.02$, $r = 0.03$, $f(K) = 3K^{\frac{1}{2}}$, and $\mu(\sigma) = 0.07 + 0.5(\sigma^2 - 0.05^2)^{\frac{1}{2}} - 0.55\sigma$.

3.2 The Principal's Value Function

Given the results of Proposition 1, the principal's problem is to maximize her objective function (5), subject to the incentive compatibility constraint (7), the law of motion for W_t (9), and the agent's participation constraint ($W_t \geq R = 0$). The agent's continuation utility W is a sufficient state variable to characterize the principal's maximal payoff under the optimal contract. Let $F(W)$ be the principal's expected payoff (5) for the optimal contract given the agent's continuation value. $F(W)$ can be fully characterized with an ODE with boundary conditions summarized in the following proposition:¹³

¹³To reduce the level of mathematical formalism needed and maintain focus on the economic insight, we make a few largely technical assumptions to simplify the proofs. See the Appendix for details.




Figure 2: The top-left plot displays $K(\beta, \Sigma)$ as a function of β for $\Sigma = \Sigma^{FB} = 1.05$ (solid line), $\Sigma = \frac{2}{3}\Sigma^{FB}$ (dashed line), and $\Sigma = \frac{1}{3}\Sigma^{FB}$ (dotted line). The top-right plot displays $\sigma(\beta, \Sigma)$ as a function of β for $\Sigma = \Sigma^{FB} = 1.05$ (solid line), $\Sigma = \frac{2}{3}\Sigma^{FB}$ (dashed line), and $\Sigma = \frac{1}{3}\Sigma^{FB}$ (dotted line). The bottom two plots display average cash flows $CF(\beta, \Sigma)$ as a function of β for each of the three values of Σ used in the top row, and as a function of Σ for $\beta = \beta^{FB} = \frac{\lambda}{r} = 0.67$ (solid line), $\beta = \frac{2}{3}\beta^{FB}$ (dashed line), and $\beta = \frac{1}{3}\beta^{FB}$ (dotted line). These plots are generated using $\lambda = .02$, $r = .03$, $f(K) = 3K^{\frac{1}{2}}$, and $\mu(\sigma) = 0.07 + .5(\sigma^2 - 0.05^2)^{\frac{1}{2}} - 0.55\sigma$.

Proposition 2 *A solution to the principal's problem $F(W)$ exists, is unique, and is strictly concave on $W \in [R = 0, W^C]$, where W^C is chosen so that $F'(W^C) = -1$ and $F''(W^C) = 0$. $F(W)$ solves*

$$rF(W) = \max_{K; \sigma \in (\underline{\sigma}, \bar{\sigma})} \left[f(K)\mu(\sigma) - rK + \gamma W F'(W) + \frac{1}{2}\beta^2(K, \sigma)f^2(K)\sigma^2 F''(W) \right]. \quad (13)$$

$F(W)$ is C^3 for all $W \in [0, W^C]$, and has $F(0) = L$. The maximizers $\{K, \sigma\}$ are the optimal policies, and $dC_t = \max(W_t - W^C, 0)$. The agent's continuation utility evolves as in (9), which has a unique weak solution.

Proposition 2 is a standard dynamic programming solution. First, the principal pays the agent only when the agent's continuation utility exceeds a given threshold, so that

$dC_t = \max(W_t - W^C, 0)$, with $F'(W^C) = -1$ and $F''(W^C) = 0$. This payment boundary exists because the agent is risk neutral with respect to consumption and more impatient than the principal. W^C represents the point beyond which the cost of saving for the agent (due to his impatience) exceeds the benefit of avoiding contract termination after series of negative shocks.¹⁴

Second, the principal chooses $K > 0$ and has $F(0) = L$. This means that the project is always running and that if the agent's continuation utility drops to his outside value, the contract is terminated.¹⁵

Finally, for $W \in [R, W^C)$, applying Ito's Lemma to $dF(W)$, we obtain

$$dF(W_t) = \gamma W_t F'(W_t) dt + \frac{1}{2} \beta_t^2 f(K_t)^2 \sigma_t^2 F''(W_t) dt + \beta_t f(K_t) \sigma_t F'(W_t) dZ_t.$$

Combined with cash flows, this yields the principal's Hamilton-Jacobi-Bellman (HJB) equation, (13). Because the principal's value function is concave, the $F''(W)$ term is always negative and represents the principal's cost of providing incentives to the agent. We discuss properties of the HJB equation and the dynamics of the endogenous variables $K(W)$, $\beta(W)$, $\sigma(W)$, and $\Sigma(W)$ in the next subsection.¹⁶

3.3 The Contract Description

The optimal contract admits two different types of solutions. In one type, the agency problem induces the principal to implement levels of project volatility above the first-best, and in the other type, project volatility is below the first-best. The key distinguishing factor is the

¹⁴These boundary conditions are the same as those in [DeMarzo and Sannikov \(2006\)](#), in which a detailed argument can be found. The principal can always make a lump-sum payment of dC to the agent, moving the agent from W to $W - dC$. This transfer benefits the principal only if $F(W - dC) - dC \geq F(W)$, and so we have no transfers if $F'(W) \geq -1$. Thus, we define $W^C = \min\{W | F'(W) \leq -1\}$, and the smooth-pasting condition is $F'(W^C) = -1$. Since W^C is optimally chosen and the principal has linear utility, we have the super-contact condition $F''(W^C) = 0$. See [Dumas \(1991\)](#) for a general discussion of the smooth-pasting and super-contact conditions.

¹⁵In a more general model with $R > 0$, the principal can temporarily shut down the project by setting $K = 0$ and allowing the agent's continuation value (W) to reflect upwards. We discuss this generalization in Section 6. Our proof of Proposition 2 shows that this temporary shutdown only ever occurs at $W_t = R$, and only for $R > 0$.

¹⁶ $\beta(W)$ is understood to mean $\beta(\sigma(W), K(W))$, with $\beta(\sigma, K)$ defined in (8). Similarly for $\Sigma(W)$ from (3).

elasticity of incentives with respect to volatility, evaluated at the first-best:

$$\varepsilon_{\beta,\sigma}^{FB} \equiv \left. \frac{\partial \ln \beta(\sigma, K)}{\partial \ln \sigma} \right|_{\sigma=\sigma^{FB}}, \quad (14)$$

$\varepsilon_{\beta,\sigma}^{FB} < 0$ is a constant and does not depend on K . Then, we have

Proposition 3 *The optimal contract can take two different forms. The two forms share the features that the risk of default/separation induces the principal to reduce incentives, resources, and aggregate cash flow volatility:*

1. *At $W = W^C$, the optimal contract implements the first-best, with $\beta(W^C) = \beta^{FB} = \frac{\lambda}{r}$, $K(W^C) = K^{FB}$, and $\Sigma(W^C) = \Sigma^{FB}$.*
2. *For $W < W^C$, incentives, resources, and aggregate cash flow volatility are all lower than the first-best levels, with $\beta(W) < \beta^{FB}$, $K(W) < K^{FB}$, and $\Sigma(W) < \Sigma^{FB}$.*
3. *In response to a small change in W , $\{K(W), \Sigma(W), \beta(W)\}$ all change in the same direction. If L is high enough, then $\{K(W), \Sigma(W), \beta(W)\}$ are all monotonically increasing.*

However, the two contract forms treat project volatility differently:

- a. *Under- σ (“overly prudent”): If $\varepsilon_{\beta,\sigma}^{FB} < -1$, then*
 - $\sigma(W) < \sigma^{FB}$ for all $W < W^C$, with equality at $W = W^C$.
 - *In response to a small change in W , $\sigma(W)$ changes in the same direction as $\{K(W), \Sigma(W), \beta(W)\}$. If L is high enough, then $\sigma(W)$ is monotonically increasing.*
- b. *Over- σ (“overly risky”): If $\varepsilon_{\beta,\sigma}^{FB} > -1$, then*
 - $\sigma(W) > \sigma^{FB}$ for all $W < W^C$, with equality at $W = W^C$.
 - *In response to a small change in W , $\sigma(W)$ changes in the opposite direction as $\{K(W), \Sigma(W), \beta(W)\}$. If L is high enough, then $\sigma(W)$ is monotonically decreasing.*

Finally, if we generalize Definition 1 to allow for resource misallocation (private benefits) in optimal contracts, then all optimal contracts implement zero misallocation, except possibly at W^C .

Proposition 3.1 is derived from the basic properties of dynamic principal-agent models: the payment boundary (W^C) exists where the cost of paying the agent (reducing future continuation utility by promise keeping) exactly equals the benefit (from the reduced possibility that a future sequence of negative shocks moves W_t down to R). Because the agent is risk neutral with respect to consumption, the principal is effectively risk neutral with respect to risk in the agent's continuation utility at W^C , so $F''(W^C) = 0$. Then, there is no cost to providing incentives (plugging $F'' = 0$ into Eq. 13) and we achieve the first-best at W^C .

Between the left and right boundaries, the cost of incentives prevents the principal from implementing the first-best project choice. The principal is effectively risk averse because of the cost of liquidation, so $F''(W) < 0$ and the flow cost of incentives (from Eq. 13) is

$$\frac{1}{2}\beta^2(K, \sigma)\Sigma(K, \sigma)^2 F''(W) = \frac{1}{2}\lambda \left(\frac{f(K)}{f'(K)} \right)^2 \left(\frac{\sigma}{\mu(\sigma) - \sigma\mu'(\sigma)} \right)^2 F''(W) \quad (15)$$

There are two useful ways of understanding the principal's choices. The first is to examine the principal's volatility controls, $\{\Sigma, \beta\}$. These give us volatility and incentive choices at the relationship level. The second is to examine the cash flow inputs, $\{K, \sigma\}$. These give us inputs and risk choices at the production level. The mapping between $\{K, \sigma\}$ and $\{\Sigma, \beta\}$ is given by the formulas for Σ and β (Eqs. 3 and 8).

It is clear that the cost of incentives is an increasing function of β and Σ : increasing either increases the volatility of the agent's continuation value (left hand side of Eq. 15). As the principal becomes more effectively risk averse ($F''(W)$ becomes more negative), she reduces the volatility of the agent's continuation utility, meaning she reduces β and Σ . The result is Proposition 3.2.

A novel result of Proposition 3 is that the optimal contract can implement levels of project-based risk (σ) that are higher or lower than the first-best. The volatility of the agent's continuation value is *not* a purely increasing function of project-level risk. The reason is that σ affects the volatility of the agent's continuation value through two opposing mechanisms: on the one hand, Property 3 shows that $\beta_\sigma < 0$. So, implementing a larger σ (which leads to a less risk-efficient cash flow) requires weaker incentives. On the other hand, cash flow volatility is increasing in project-level volatility since $\Sigma = f(K)\sigma$. The volatility of the agent's continuation value is a product of these two effects, $\Sigma_\sigma > 0$ and $\beta_\sigma < 0$, and whether the optimal contract implements a higher or lower σ relative to the first-best depends on which effect dominates.

The elasticity of incentives with respect to project volatility ($\varepsilon_{\beta, \sigma}^{FB}$) captures the effect of

σ on continuation value volatility. When $\varepsilon_{\beta,\sigma}^{FB} < -1$, incentives can be made much weaker for high levels of project volatility; enough weaker that total volatility of the agent's continuation volatility ($\beta\Sigma$) is *lower* for high project-level volatility. When $\varepsilon_{\beta,\sigma}^{FB} > -1$, incentives can be made only slightly weaker for high levels of project volatility; the dominant effect is that higher project-level volatility causes *higher* total continuation volatility. The critical distinction here is between cash flow volatility and continuation value volatility. The agency problem dictates that it is the risk of default and termination that generates losses, and therefore the agent's continuation value volatility that generates risk for the principal.

While resource levels (K) also affect continuation value volatility through both the incentives and project volatility channels, both channels operate in the same direction (Σ_K and β_K are both positive). More resources imply more total volatility, and stronger incentives are necessary to prevent shirking with high resource intensity because the marginal return to productive inputs is lower. Thus the optimal contract always features *under-investment* ($K_t \leq K^{FB}$) relative to the first-best.

To complete the discussion of Proposition 3.3, we examine the impact of small changes in W . If a given change in W causes $F''(W)$ to become more negative – the principal becomes effectively more risk averse – then she acts to reduce Σ and β , and makes the corresponding change in K and σ . Thus, the instantaneous change in controls is perfectly correlated (1 or -1). Finally, we are able to show the slightly limited result that if L is high enough, then the contract's policy variables are all monotonic. For these economies, decreasing W brings the principal closer to default, making her more risk averse and increasing the cost of incentives (making F'' more negative).¹⁷

Last but not least, we note that because misallocation is assumed to be weakly inefficient ($\lambda \leq r$), the optimal contract is robust to considering positive resource misallocation in equilibrium, as shown in the final statement of Proposition 3.

To summarize the results in this subsection, we illustrate an optimal contract in Fig. 3. We label solutions for $\varepsilon_{\beta,\sigma}^{FB} > -1$ and $\sigma_t \leq \sigma^{FB}$ as “under- σ ”; and solutions for $\varepsilon_{\beta,\sigma}^{FB} < -1$ and $\sigma_t \geq \sigma^{FB}$ as “over- σ .”

¹⁷We have been unable to prove this result for general values of L or to find a numerical counter-example. Our proof uses the sufficient condition that if L is high enough, then $F'''(W) > 0$, and the policy variables are monotonic. However, $F'''(W) > 0$ is not necessary for the result nor generically true in models built on DeMarzo and Sannikov (2006) and Biais, Mariotti, Plantin, and Rochet (2007). These models have constant mean and volatility for cash flows and the property that for some parameter values, particularly intermediate values of L , the principal can become less risk averse when close to default. Piskorski and Westerfield (2016) demonstrates how one can use additional assumptions to obtain sharper results in this class of models.




Figure 3: All plots are generated using $f(K) = 2K^{\frac{1}{2}}$, $L = 1$, $R = 0$, $\lambda = 0.02$, $r = 0.03$, and $\gamma = 0.05$. The “Under- σ ” column uses $\mu(\sigma) = \frac{1}{3}\sigma^{\frac{1}{2}} - 0.37\sigma$ and generates $W^C = 3.19$; the “Over- σ ” column uses $\mu(\sigma) = -3\sigma^2 + \sigma$ and generates $W^C = 3.27$. In the top row, the solid line is $F(W)$, and the dashed line is the right-boundary condition $rF(W^C) = \max[CF(K, \sigma)] - \gamma W^C$. In the second row, the solid line is $\beta(W)$ and the dashed line is $\Sigma(W)$.

4 The Implementation

4.1 Description

In this section, we show that the optimal contract can be implemented with a simple two-part tariff on production inputs (resources). The most critical features of the implementation are that the two-part tariff is a *static* function of the agent's visible choices, and the principal *does not keep track* of the agent's continuation utility. In other words, after setting up the two-part tariff, the principal's role is simply to apply the static functions. In that sense, this is a theory of institutional design, rather than active management.

The principal offers resources (liquid capital or intermediate goods) to the agent as a two-part tariff:

- The fixed assets of production (the assets that have liquidation value L) are assigned a rental price ϕ .
- The variable resources used in production, K , are assigned a unit price θ . For capital, θ is a hurdle rate or capital charge; for intermediate goods, θ is a transfer price.

The agent can freely request any level of resources, cash flow volatility, and even his pay-for-performance sensitivity (i.e., any combination of $\{\tilde{K}_t, \tilde{\Sigma}_t, \tilde{\beta}_t\}$). The tilde notation is used to indicate that those quantities are choices of the agent. The tariff is adjusted based on the agent's choices: $\theta = \theta(\tilde{\Sigma}_t, \tilde{\beta}_t)$ and $\phi = \phi(\tilde{\Sigma}_t, \tilde{\beta}_t)$, and the cost of resources is deducted from the project's cash flows:

$$dY_t^{NEW} \equiv dY_t - \tilde{K}_t \theta(\tilde{\Sigma}_t, \tilde{\beta}_t) dt - \phi(\tilde{\Sigma}_t, \tilde{\beta}_t) dt. \quad (16)$$

The agent retains a fraction of *net* project cash flows, $\tilde{\beta}_t dY_t^{NEW}$, which is placed into a cash account with balance M_t that the agent controls. This account grows at interest rate γ but the agent can freely withdraw consumption from the account as long as the balance is positive.¹⁸ Thus, the account balance evolves according to:

$$dM_t = \gamma M_t dt + \tilde{\beta}_t dY_t^{NEW} - dC_t. \quad (17)$$

The original agency frictions remain: the agent can both allocate resources so as to obtain private benefits and change project-level risk. A portion of the obtained resources ($\hat{K}_t \leq \tilde{K}_t$)

¹⁸As in DeMarzo and Sannikov (2006) or Biais, Mariotti, Plantin, and Rochet (2007).

is allocated to the productive project with $\hat{\sigma}_t$, and the agent obtains the corresponding private benefit $\lambda(\tilde{K}_t - \hat{K}_t)$, subject to the chosen volatility constraint: $f(\hat{K}_t)\hat{\sigma}_t = \tilde{\Sigma}_t$. The agent is prohibited from borrowing on the account, so we require $M_t \geq 0$.

The most critical property of the implementation is that it only involves static functions of the agent's choices $\{\tilde{\Sigma}_t, \tilde{\beta}_t\}$. The equilibrium price levels $\{\theta(\tilde{\Sigma}_t, \tilde{\beta}_t), \phi(\tilde{\Sigma}_t, \tilde{\beta}_t)\}$ are time varying because the agent's choices are time varying, but the price functions do not change. The principal does not need make any dynamic adjustment to those functions based on the agent's specific policy choices or his performance history. In fact, she does not even need to track the agent's continuation utility; the account M_t performs this role. The evolution of M_t is set so that $W_t = M_t$ in equilibrium, despite the fact that the agent can withdraw consumption from the account at any time.

We formally summarize the implementation as follows:

Definition 2 (Implementation and two-part tariff) *An implementation is a two-part tariff $\{\phi(\tilde{\Sigma}, \tilde{\beta}), \theta(\tilde{\Sigma}, \tilde{\beta})\}$ and an account M such that for all $t \geq 0$, the agent chooses a set of actions $\{\tilde{\Sigma}_t, \tilde{\beta}_t, \tilde{K}_t, \hat{K}_t, dC_t\}$ to maximize his expected utility subject to (16) and (17).*

An optimal implementation is a two-part tariff $\{\phi(\tilde{\Sigma}, \tilde{\beta}), \theta(\tilde{\Sigma}, \tilde{\beta})\}$ and account M that induces the agent to choose $\{\tilde{\Sigma}_t, \tilde{\beta}_t, \tilde{K}_t, \hat{K}_t, dC_t\}$ as in the optimal contract, with $dC_t = \max(M_t - W^C, 0)$, and the agent quitting when $M_t = 0$.

Interpreting K as capital, our two-part tariff matches the common practice of placing a capital charge in the profit-and-loss accounting for a desk or division. Sometimes this is called KVA (Capital Value Adjustment) or FVA (Funding Value Adjustment).¹⁹ The goal is to measure the performance of a desk or division on an on-going basis, often for the purpose of determining compensation. In private equity, a similar mechanism is a preferred return to outside investors, in which outside investors capture a given return on capital and the residual is then split given a sharing rule with the general partner (see, e.g., [Metrick and Yasuda, 2010](#); [Robinson and Sensoy, 2013](#)). More generally, our implementation rationalizes the use of price controls, rather than quantity controls, in capital budgeting problems.

Interpreting K as an intermediate good, such as microchips or service hours, our two-part tariff matches the common practice of using cost-based transfer pricing, but imposing a markup over marginal cost.²⁰

¹⁹See, e.g., [Andersen, Duffie, and Song \(2019\)](#) for a full description of Funding Value Adjustments, including examples and some history.

²⁰69% of firms according to [PriceWaterhouse \(1984\)](#). [PriceWaterhouse \(1984\)](#) and [Tang \(1992\)](#) contain extended survey evidence.

4.2 Optimality

An optimal implementation uses the agent's choice of $\{\tilde{\Sigma}, \tilde{\beta}\}$ to determine the two part tariff $\{\phi, \theta\}$. We have:

Proposition 4 *There exists an optimal implementation with the following properties:*

$$\theta(\tilde{\Sigma}, \tilde{\beta}) = \frac{\lambda}{\tilde{\beta}} \quad (18)$$

$$\phi(\tilde{\Sigma}, \tilde{\beta}) = f(K(\tilde{\Sigma}, \tilde{\beta}))\mu(\sigma(\tilde{\Sigma}, \tilde{\beta})) - \theta(\tilde{\Sigma}, \tilde{\beta})K(\tilde{\Sigma}, \tilde{\beta}) \quad (19)$$

where $K(\tilde{\Sigma}, \tilde{\beta}), \sigma(\tilde{\Sigma}, \tilde{\beta})$ in (18) and (19) jointly solve the incentive compatibility (8) and feasibility (3) conditions:

$$\tilde{\beta} = \frac{\lambda}{f'(K)} \frac{1}{\mu(\sigma) - \sigma\mu'(\sigma)} \quad (20)$$

$$\tilde{\Sigma} = f(K)\sigma \quad (21)$$

and $\phi = \theta = \infty$ if the solution to Eqs. (20) and (21) does not exist.²¹ In addition, $M_0 = W_0$. Moreover, in equilibrium, $E[dY_t^{NEW}] = 0$ for all t .

The optimal implementation is the result of several intuitions about how the agency problem is constructed. First, every new element of the adjusted cash flow dY^{NEW} is observable to both the principal and the agent. Thus, the underlying hidden action problem is the same in the implementation as in the original contract setting. Put differently, while it is usually assumed that the principal takes the agent's output from the underlying productive technology as the performance criterion, in fact the principal can look at any augmented performance criterion that she likes. In our setting, the augmentation is a resource charge.

Second, the value of the resource charge is simple ($\tilde{\beta}\theta = \lambda$) exactly because the agent is already optimizing over resource usage. The agent's hidden-action optimality implies that the marginal value of resources used in productive output ($E[\beta_t dY_t]$) is the same as the marginal value of resources used for private benefits (λ). So, to induce the agent to allocate resources correctly, we do not need to exactly measure the marginal value in productive use; we can instead set the resource charge equal to the marginal value of private benefits. Thus, we have $\tilde{\beta}\theta = \lambda$, which gives the agent no incentive to ask for additional resources for private

²¹The solution exists if the given $\tilde{\beta}$ and $\tilde{\Sigma}$ is part of any incentive compatible contract (see Definition 1). However, some pairs of $\tilde{\beta}$ and $\tilde{\Sigma}$ are not part of any valid contract (e.g., $\beta = 0$ and $\Sigma > 0$).

benefits. If he were to misallocate resources, he would reduce $E[\beta_t dY_t]$ by exactly the value of private benefits.

The resource charge has two additional features to note:

- The resource charge (hurdle rate, transfer price, unit cost) is always higher than the resource cost to the principal. The implementation has $\theta_t = \lambda/\beta_t$, and the optimal contract has $\beta_t \leq \lambda/r$; together these imply $\theta_t \geq r$, with equality only at $W = W^C$. Thus, the principal always sets the resource charge above the cost at which she obtains the resource from external markets or other divisions.
- $\partial\theta/\partial\tilde{\Sigma} = 0$. The resource cost is adjusted for the agent's choice of cash flow share, but it is not additionally adjusted for risk. Again, this is exactly because the agent is already optimizing: the marginal product of resources in productive use must equal the marginal value of resources used for private benefits, and the private value is constant. There is no additional modification for risk. In particular, the principal does not change the cost of capital to adjust for risk.

Third, the value of ϕ is set so as to make the agent indifferent across any pair of incentive compatible projects. To generate this indifference, the principal sets the fixed fee $\phi(\tilde{\Sigma}, \tilde{\beta})$ to extract the entire surplus that remains after the variable cost (θ) is deducted. More specially, given the agent's choices of $\{\tilde{\Sigma}, \tilde{\beta}\}$, and implied value of θ , the agent chooses $\{\tilde{K}, \hat{K}, \tilde{\sigma}\}$ so as to maximize $E[dY^{NEW}]$ (with private benefits set to zero). However, ϕ is set so that this maximum $E[dY^{NEW}]$ is achieved at zero for any $\{\tilde{\Sigma}, \tilde{\beta}\}$ the agent might choose. Because the agent's expected flow utility is zero for any choice of $\{\tilde{\Sigma}, \tilde{\beta}\}$, the agent is indifferent and makes the principal's desired choice.

In other words, the agent receives his full value from the contract when it is signed, and the entire flow of expected producer surplus goes to the principal. Since the agent cannot do better than the principal's desired choices, the agent does as the principal recommends.²²

²²To see this indifference tradeoff in more detail, imagine the agent has continuation utility W'_t , which corresponds in the optimal contract to $\Sigma(W'_t)$ and $\beta(W'_t)$. Given incentive compatibility, choosing $\Sigma(W'_t)$ and $\beta(W'_t)$ results in $dW'_t = \gamma W'_t dt + \Sigma(W'_t)\beta(W'_t)dZ_t$ (from Eq. 9). Now, let us allow the agent to choose $\Sigma(W''_t)$ and $\beta(W''_t)$, associated with another value of continuation utility W''_t . These new choices also lead to an incentive compatible resource allocation, and they result in $dW'_t = \gamma W'_t dt + \Sigma(W''_t)\beta(W''_t)dZ_t$. However, the agent is indifferent between these two outcomes because $E[dZ_t] = 0$, and so the expected change in continuation utility is the same under both choices. Further, this is not a product of risk neutrality over consumption: the Hamilton-Jacobi-Bellman formulation implies that all agents are risk neutral with respect to continuation utility. This is exactly analogous to setting the agent's flow continuation utility (see, e.g., Eq. 10) equal to zero at the optimum, a property that the martingale representation theorem implies any incentive compatible contract must have.

We can understand the agent's ability to choose his compensation structure (his own pay-performance sensitivity $\tilde{\beta}$) through a trade-off with resource allocation. The ability to choose $\tilde{\beta}_t$ is surprising because the private benefit from resource misallocation, λ , is fixed. For example, in [DeMarzo and Sannikov \(2006\)](#) the desired pay-performance sensitivity is implemented with inside equity that is chosen so that the marginal benefit from reporting additional cash flow is equal to the marginal benefit from resource misallocation. Both are constant. However, in our implementation, θ is designed such that the agent has to trade off two different controls, $\tilde{\beta}_t$ and $\hat{K}_t$. If $\tilde{\beta}_t$ is chosen to be very small, the cost of resources is very high, so the agent has to use the resources productively to avoid a loss of continuation value. If $\tilde{\beta}_t$ is chosen to be high, then resources are cheap, but the gains to resource misallocation are lower than the agent's chosen cash flow residual (pay-performance sensitivity).

We illustrate the shape of the two-part tariff in Figs. 4 and 5. Fig. 4 depicts θ and ϕ as functions of β and Σ , i.e., the level of the two-part tariff necessary to implement a particular pair of (β, Σ) . Fig. 5 depicts θ and ϕ as functions of W , i.e., the equilibrium level of tariff under the optimal contract for each different state. The patterns plotted are consistent with the proceeding descriptions. In particular, the resource cost θ decreases in W because when the agent's continuation utility grows, the agency friction is less severe, so the principal imposes lower costs.

The last part of the implementation, the cash balance, is standard (e.g., [DeMarzo and Sannikov, 2006](#); [Biais, Mariotti, Plantin, and Rochet, 2007](#)). It is a mechanism to turn the stock of promised future consumption (W) into an explicit account (M) from which the agent can consume. Since the account has rate of return γ , which is equal to the agent's time discount rate, the agent is indifferent between the timing of consumption (dC_t versus $e^{\gamma s}dC_{t+s}$ for $s \geq 0$). Thus, the agent is willing to consume as per the optimal contract. The agent quits at $M_t = 0$, receiving outside value $R = 0$ and matching the optimal contract; this is because the agent is not allowed $M_t < 0$ and so cannot tolerate any positive volatility at $M_t = 0$.

Finally, we have written the implementation so that the cost of variable resources (θ_t) is deducted from the project cash flow instead of the agent's continuation value or cash account (M_t). This is not required; we could deduct the cost of variable resources from the agent's account directly. The difference is whether we interpret the cost as being paid by the agent or by the project. We have consistently interpreted the resource cost as a divisional charge (a preferred return on invested capital, a funding value adjustment, or a transfer price), meaning it is paid by the project, not the agent personally. We can also take either β_t or Σ_t




Figure 4: These plots are calculated using the values from Fig. 2 and depict the functions in Proposition 4. The two left plots display $\theta(\beta, \Sigma)$ and $\phi(\beta, \Sigma)$ as functions of β for $\Sigma = \Sigma^{FB} = 1.05$ (solid line), $\Sigma = \frac{2}{3}\Sigma^{FB}$ (dashed line), and $\Sigma = \frac{1}{3}\Sigma^{FB}$ (dotted line). Note that $\theta(\beta, \Sigma)$ does not vary with Σ so the lines overlap. The two right plots display $\theta(\beta, \Sigma)$ and $\phi(\beta, \Sigma)$ as functions of Σ for $\beta = \beta^{FB} = \frac{\lambda}{r} = 0.67$ (solid line), $\beta = \frac{2}{3}\beta^{FB}$ (dashed line), and $\beta = \frac{1}{3}\beta^{FB}$ (dotted line). These plots are generated using $\lambda = .02$, $r = .03$, $f(K) = 3K^{\frac{1}{2}}$, and $\mu(\sigma) = 0.07 + .5(\sigma^2 - 0.05^2)^{\frac{1}{2}} - 0.55\sigma$.

out of the agent's choice set; i.e., we have presented the most decentralized implementation by giving the agent the choice over both β_t and Σ_t as well as K_t , but that is not required.

5 Empirical Discussion

5.1 Capital Budgeting and the Cost of Capital

In addition to the qualitative features of our implementation, the equivalence of a Funding Value Adjustment or the private equity preferred return discussed earlier, our implementation captures some broadly observed stylized facts surrounding the use of hurdle rates for capital




Figure 5: These plots are calculated using the values from Fig. 2 and depict the functions in Proposition 4. The plots display $\theta(W) = \theta(\beta(W), \Sigma(W))$ and $\phi(W) = (\beta(W), \Sigma(W))$. These plots are generated using $\lambda = .02$, $r = .03$, $f(K) = 3K^{\frac{1}{2}}$, and $\mu(\sigma) = 0.07 + .5(\sigma^2 - 0.05^2)^{\frac{1}{2}} - 0.55\sigma$.

budgeting:²³

- Most or almost all firms use DCF methods with a hurdle rate. That hurdle rate is substantially above *both* the econometrician-estimated and firm-estimated cost of capital. For example, Jagannathan, Matsa, Meier, and Tarhan (2016) reports an average hurdle rate of 15% compared to an average cost of capital of 8%. They find that this is not likely to be caused by behavioral biases or driven by managerial exaggeration.
- Firms engage in deliberate capital rationing. This rationing is often a response to non-financial constraints; more than half of firms report that they pass up apparently positive NPV projects because of constraints on managerial time and expertise (55.3%, according to Jagannathan, Matsa, Meier, and Tarhan, 2016).
- Most firms (58.8%, according to Graham and Harvey, 2001) use a company-wide discount rate that is not adjusted for project risk, and very few firms (15%, according to Graham and Harvey, 2001) use divisional hurdle rates.

Our model is consistent with these results on hurdle rates and capital rationing: first and foremost, we demonstrate in Section 4 that the equilibrium hurdle rate is indeed higher than the firm's true (external) cost of capital (i.e., $\theta_t = \lambda/\beta_t \geq r$). This is optimal because the agency problem imposes a constraint on the use of managerial time and expertise. The

²³This list is a summary of results in Jagannathan, Matsa, Meier, and Tarhan (2016), Graham and Harvey (2001, 2011, 2012), Jacobs and Sivdasani (2012), and Poterba and Summers (1995).

principal must offer the agent a portion of residual cash flow in order to induce the desired project choice and capital usage. This portion, combined with limited liability on the part of the manager, creates the possibility of termination, which entails the loss of a high NPV project. To avoid the larger loss, the principal accepts the smaller loss of reducing the scale of the agent’s activity in order to lower the volatility of the agent’s residual claim. Moreover, our model is consistent with firms using hurdle rates that do not vary with project risk. Again, this is exactly because the agent is already optimizing: the marginal product of capital in productive use must equal the marginal value of capital in private benefits, and the private value is constant. There is no additional modification for risk.

5.2 Transfer Pricing

Our implementation also provide a rationale for observed practices of transfer pricing. In particular, the majority of firms (69%, according to [PriceWaterhouse, 1984](#)) that use cost-based transfer pricing imposes a markup over the marginal cost. Our paper demonstrates the optimality of such a markup policy: the principal (i.e., the headquarters) implements the optimal allocation and risk policies via a transfer price (θ_t) that is higher than the marginal cost of the input (r). Importantly, the markup implements an optimal contract, which is different from existing models that rationalize the use of the such policy among a set of restricted mechanisms or incomplete contracts.²⁴

The transfer pricing literature points to two primary determinants of transfer prices: taxes (which are important for multinational firms but relatively less so for domestic-only firms) and incentives. Consequently, a two-book accounting system is a common practice. In this paper, we implicitly assume that firms can completely separate the two prices in their accounting while still including the tax effects of transfer pricing when projecting cash flows. [Hiemann and Reichelstein \(2012\)](#) and [Hyde and Choe \(2005\)](#) are part of a long literature that analyzes the challenges associated with these accounting schemes.

²⁴See, e.g., [Göx and Schiller \(2006\)](#) and [Baldenius \(2008\)](#) for a summary of the incomplete contract models on transfer pricing. Restricted mechanisms are also common. For example, [Sahay \(2003\)](#) demonstrates the advantage of the markup strategy by comparing it to an alternative, “multiplicative” strategy in a standard two-division framework. [Baldenius, Dutta, and Reichelstein \(2007\)](#) argues that the “competitive hurdle rate” mechanism (i.e., hurdle rate above the headquarters’ cost of capital) optimally solves the division’s moral hazard problem while assuming an exogenous capital constraint. Other examples include [Vaysman \(1996\)](#), [Vaysman \(1998\)](#), [Baldenius, Reichelstein, and Sahay \(1999\)](#), , and [Pfeiffer, Schiller, and Wagner \(2011\)](#). In contrast, we study a dynamic agency problem and the internal transfer scheme that implements the optimal contract. We expand from existing studies by showing the optimality of the cost-markup among any arbitrary pricing mechanism in our model.

The seminal work in the incentive-based transfer pricing literature is [Harris, Kriebel, and Raviv \(1982\)](#), while [Rajan and Reichelstein \(2004\)](#) summarizes the follow-up literature. [Harris, Kriebel, and Raviv \(1982\)](#) analyze a mechanism design problem in which divisions have hidden information about productivity and preferences that diverge from those of headquarters. Importantly, the transfer pricing scheme is a menu of combinations of price and quantity that is not linear. As argued by [Rajan and Reichelstein \(2004\)](#), the scheme is centralized because the firm’s central office commits to allocation and compensation rules for all possible combinations of (multi-)divisional reports. Their allocation is also of a bang-bang nature in that a division receives either zero resource or highest possible amount of resources. Our implementation differs in major ways. First, headquarters only commits to a price (determined from the agent’s choice of compensation structure), while the agent selects the resource quantity. Thus, our scheme is more decentralized. Second, the agent’s choice is continuous, and because the price is almost always above the marginal cost, the agent almost always asks for and receives fewer resources than in the first-best. Third, our model is dynamic but the price schedule is static; thus, the agent’s choice of project and productivity (volatility and mean cash flow) can vary over time, but the optimal transfer pricing scheme does not need to be continuously updated. All three of these features improve on matching observed practice: stated prices, with markups above marginal cost, that do not change continuously and allow for discretion by the manager.

5.3 Risk Taking and Pay-Performance Sensitivity

Dynamic adjustment in risk taking (σ) can lead to overly risky or overly prudent investments. The source of this result is the difference between the volatility of the agent’s continuation value, which is what drives the principal’s termination/default risk, and the volatility of the cash flow. The volatility of the agent’s continuation value is the product of project cash flow volatility and the intensity of incentives. Thus, depending on model specification, risk reduction can mean reducing project volatility or allowing project volatility to increase in order to reduce incentive intensity. One result is that actions that look like risk shifting can actually be risk reducing.

The dynamic risk adjustment can generate strong empirical differences between the static and dynamic responses to incentives; the time series and the cross section might look very different because of different sources of variation. Consider, for example, the static $\beta(K, \sigma)$ function in (8): we have $\beta_\sigma < 0$, meaning that for any given state of the world, stronger




Figure 6: This plot depicts β as a function of σ (Eq. 8, solid lines) and the locus of points $\{\beta(W), \sigma(W)\}$ (the solution to Eq. 13, dashed line). The plot is generated using the under- σ specification from Fig. 3: $f(K) = 2K^{\frac{1}{2}}$, $L = 1$, $R = 0$, $\lambda = 0.02$, $r = 0.03$, $\gamma = 0.05$, and $\mu(\sigma) = \frac{1}{3}\sigma^{\frac{1}{2}} - 0.37\sigma$.

incentives *cause* lower project risk taking. However, when we consider the path of outcomes over time, the correlation between incentives and *total* risk (β and Σ) is *positive* (see, e.g., Fig. 3 and Proposition 3.4). At the same time, the correlation between incentives and *project* risk (β and σ) can be positive (in the under- σ specification) or negative (in the over- σ specification). Thus, the causal results and dynamic correlations can have opposite signs. We illustrate this in Fig. 6, where the dashed line depicts how the economy evolves as a function of W , and the solid lines are the function $\beta(K, \sigma)$ for two different values of K .

Empirically, this result means that our model predicts different results for time series and cross-sectional tests in the under- σ specification, because the source of variation is different.²⁵

²⁵The empirical literature has indeed found contradictory results across different settings. While some studies find that firms increase risk taking following poor performance (e.g., Eisdorfer, 2008), others find no such effect (e.g., Andrade and Kaplan, 1998) or a reduction of risk (e.g., Gilje, 2016). The asset management literature paints a similarly mixed picture (e.g., Rauh, 2008, Huang, Sialm, and Zhang, 2011, Aragon and Nanda, 2011, and their discussions of the literature). Note that while these empirical studies address different institutional backgrounds and may follow different assumptions, our model suggests that more sophisticated controls are needed to assess the relationship between risk choice and performance. The relationship should be consistent in the time series for each firm, but differ in the cross section based on unobservable characteristics, like investment opportunities.

If the dominant source of variation is the history of success or failure (which changes over time), then we should see a positive relationship between the agent’s share and volatility. This is movement along the dynamic curve, and may be more likely in time series variation. If the dominant source of variation is firm characteristics that are stable and not the history of success or failure, then we might see a negative relationship between the agent’s share and volatility. If our model is correct, then an experiment or identification that uncovers the causal mechanism between the agent’s share and volatility would show the static curve, which we predict to be a negative relationship.

The over- σ specification does deliver another empirical test. In settings with easily scalable investment and in which the relevant risk is agent-separation, rather than default, increasing project risk after failure should be more common. Investment funds, especially mutual funds and hedge funds, would seem to be good examples. To the extent that fund manager skill is real, the primary danger to fund value is that the high-skill manager leaves. In fact, many empirical results (e.g., [Chevalier and Ellison, 1997](#), [Aragon and Nanda, 2011](#), [Huang, Sialm, and Zhang, 2011](#)) find increasing project risk after failure. However, those studies often attribute the increase in risk to convex incentive schemes. Our mechanism is different: increasing project risk actually decreases termination risk. A useful empirical test would be to distinguish changes in inside and outside values, and to see to what extent that difference impacts investment risk.

6 Further Discussion and Conclusion

6.1 Model Robustness

In the Appendix and Internet Appendix, we present and discuss several extensions to our model.

In the Appendix, we generalize the outside options of the principal and the agent from $L > 0$ and $R = 0$ to the more general $L \geq 0$ and $R \geq 0$. The primary outcome is that the principal can choose $K = 0$ for some histories, meaning the firm temporarily shuts down. This temporary shutdown only ever occurs at $W_t = R$, at which point the agent’s continuation value reflects upward and the firm re-starts operations.

In the Internet Appendix:

1. We allow private benefits to generate cash flow volatility, so that total volatility is $\Sigma_t = f(\hat{K}_t)\mu(\hat{\sigma}_t) + \lambda(K_t - \hat{K}_t)\epsilon$. We show that if ϵ is high enough, we change the

sign of $\frac{\partial}{\partial \sigma} \beta(K, \sigma)$ so that stronger incentives are used by the principal to increase risk taking.

2. We allow the private benefits function ($\lambda(K_t - \hat{K}_t)$ in the main model) to have an arbitrary form ($\Lambda(K_t, \hat{K}_t)$). This allows contracts that have overinvestment ($K_t > K^{FB}$) and efficient diversion ($K_t > \hat{K}_t$ in an optimal contract).²⁶
3. We allow for an additional effort choice by the agent. This changes the dimensionality of the agent's problem so that the principal uses two controls (β and Σ) to induce three choices by the agent ($\hat{K}$, $\hat{\sigma}$, and $\hat{e}$). The principal can only implement a two-dimensional curve in a three-dimensional space. The primary result is that the economics of the $(\hat{K}, \hat{\sigma})$ choice differ from that of hidden effort.

The purpose of these extensions is to demonstrate model robustness – that the agency problem over the components of volatility does not depend for its existence on the specific reduced-form assumptions we have made in our presentation. The contract implementation is substantively unchanged as a result of the first two modifications, except for the particular values of $\{\phi(\Sigma, \beta), \theta(\Sigma, \beta)\}$.²⁷ With an arbitrary private benefits function, the conceptual form of the implementation, a static set of formulae for the cost of resources, remains unchanged, although that cost is no longer linear in K . If we add an additional effort choice by the agent with arbitrary cost function, then the agent's problem implies a non-linear private benefits function, so the implementation also involves a non-linear pricing scheme.

6.2 Conclusion

Continuous-time principal-agent models have developed rapidly with applications to expanding areas of economic research. Despite the progress, almost all of the existing models involve the agent controlling the drift of the output/cash flow process. We add to the literature by considering the optimal contract when the agent controls the volatility. In the model, overall cash flow is made up of two components: the individual risk of the project and resource intensity, both of which are observable only to the agent. The principal must incentivize

²⁶This result complements other studies featuring overinvestment under the optimal contract in which other market frictions and elements are introduced on top a dynamic moral hazard problem; for example, [Ai and Li \(2015\)](#) and [Bolton, Wang, and Yang \(2019\)](#) with limited commitment, [Szydlowski \(2019\)](#) and [Gryglewicz, Mayer, and Morellec \(2020\)](#) with multitasking, [Gryglewicz and Hartman-Glaser \(2020\)](#) with real options.

²⁷To allow for shutdown, one only needs to have $\{\phi = 0, \theta = \infty\}$ as one possibility for the two-part tariff.

the agent to choose the desirable level of project risk and resource intensity. Such a setting represents a unique and different environment from a drift-control model such as modified versions of [DeMarzo and Sannikov \(2006\)](#) or [Biais, Mariotti, Plantin, and Rochet \(2007\)](#).

Interestingly, we find that the optimal incentives can be implemented with a simple resource charge (hurdle rate or transfer price) against which the agent's performance, the realized cash flow, is measured. The agent is allowed to propose the level of resources and risk he wants, and even his own incentive power. Moreover, the principal is freed from the need to keep track of the agent's performance. A natural direction of future research is to explore the implications for asset pricing and portfolio intermediation (e.g., [Buffa, Vayanos, and Woolley, 2015](#) with an optimal contract). Our resource charge is analogous to the required return of limited partners or outside investors, and the joint determination of contracts and equilibrium asset prices is a natural question to explore.

Appendix

A Proofs and Derivations

In this appendix, we provide proofs of propositions included in the main text. We assume only that $L \geq 0$ and $R \geq 0$. The restriction to $L > 0$ and $R = 0$, as discussed in the text of Section 2, is a corollary.

To reduce the level of mathematical formalism, we maintain the following assumptions:

Assumption A.1 $f(K)$ and $\mu(\sigma)$ satisfy the following assumptions in addition to those mentioned in Section 2.1:

1. $\frac{d^2}{dK^2} \left(\frac{f(K)}{f'(K)} \right)^2 \geq 0$ for all $K > 0$.
2. $\frac{d^2}{d\sigma^2} \left(\frac{\sigma}{\mu(\sigma) - \sigma\mu'(\sigma)} \right)^2 \geq 0$ for all $\sigma \geq \underline{\sigma}$.
3. There is a minimum positive level of resources that the principal can grant the agent: $K \in \{0 \cup [k_0, \infty)\}$ for some $k_0 > 0$ very small.

The first line ensures that decreasing returns occurs smoothly enough for the principal's problem to be strictly concave, so that there are no jumps in K_t . The second ensures that the variance of the agent's continuation value is convex in the standard deviation of the project's cash flow. Altogether, these restrictions are economically innocuous and consistent with most commonly used production functions and risk-return relationships. Examples of $f(K)$ that meet the conditions include $f(K) = \ln(1+K)$ and $f(K) = K^\alpha$, $\alpha \in (0, 1)$. Examples of $\mu(\sigma)$ include $\mu(\sigma) = \sigma^a - b\sigma$, $a \in (0, \frac{1}{2})$; $\mu(\sigma) = b\sigma - \sigma^a$, $a > 1$; and $\mu(\sigma) = \underline{\mu} + C\sqrt{\sigma^2 - \underline{\sigma}^2} - b\sigma$, which is the mean-variance efficient frontier if the agent has access to several projects with normally distributed cash flows. In that case, we assume C and $\underline{\mu}$ are not too large.

The final condition simplifies the proof of the existence and uniqueness of the principal's Hamilton-Jacobi-Bellman ODE. One should think of k_0 as being very small: e.g., the principal cannot allocate less than one penny of capital without allocating zero capital. Our assumption is a restriction on the principal rather than on the agent: incentive compatibility conditions are required at $K = k_0$. This assumption can be relaxed without changing the economics of the problem, but at the cost of a great deal of added formalism. See [Piskorski and Westerfield \(2016\)](#) for such a proof, when the principal's control can go continuously to zero.

A.1 $R > 0$ and Temporary Shutdown

Once we generalize our assumptions to $L \geq 0$ and $R \geq 0$, the principal can choose $K = 0$ for some histories. This represents a temporary shutdown of the firm by the principal. Our principal's proposition is:

Proposition 5 *A solution to the principal's problem $F(W)$ exists, is unique, and is concave on $W \in [R, W^C]$, where W^C is chosen so that $F'(W^C) = -1$ and $F''(W^C) = 0$. $F(W)$ solves*

$$rF(W) = \max_{K \in \{0 \cup (k_0, \infty)\}; \sigma \in (\underline{\sigma}, \bar{\sigma})} \left[f(K)\mu(\sigma) - rK + \gamma W F'(W) + \frac{1}{2}\beta^2(K, \sigma)f^2(K)\sigma^2 F''(W) \right]. \quad (22)$$

$F(W)$ is $\mathcal{C}$ for all W and $\mathcal{C}^3$ for all $W \in (R, W^C)$, and has $F(R) = L > L^$ if $K(R) > 0$ and $F(R) = L^*$ if $K(R) = 0$. Further, $K(W) > 0$ for $W \in (R, W^C]$. The agent's continuation utility evolves as in (9), which has a unique weak solution.*

This Proposition is what we prove in the Appendix. Proposition 2 is a corollary.

The key result is that termination is optional for the principal: if the principal chooses $K(\tilde{W}) = 0$, then the law of motion for W_t (9, with $\Sigma = 0$) implies that W_t reflects upwards at $\tilde{W}$. Thus, shutdown is temporary. The principal only optimally chooses $K = 0$ at $W = R$, and only when $L \leq L^*$, for some constant L^* .

The fact that the principal only chooses $K = 0$ at $W = R$ arises because setting $K = 0$ is costly, so the principal wishes to delay paying that cost as long as possible. There are two costs to setting $K = 0$. One is an opportunity cost because any time with $K = 0$ is time the principal might otherwise have positive expected cash flow. The second cost is incurred by causing W_t to reflect early: the reflection is closer to the agent's consumption boundary, and the agent will be awarded consumption sooner.

The fact that the principal only chooses $K = 0$ if L is low is economically straightforward: the principal only avoids default and termination if her value in default is low. If L is high enough, the principal simply accepts default rather than pay the opportunity cost associated with $K = 0$. Note that because of the costs associated with termination, the principal's value function is concave regardless of whether termination occurs in equilibrium. If $L > L^*$, there is a direct cost associated with termination as long as L is less than the discounted, first-best cash flow. This cost makes volatility undesirable, and the principal's value function is concave. If $L < L^*$, there is a direct cost to termination that the principal avoids in equilibrium, choosing instead to pay the opportunity cost of shutting down the project. The principal forgoes the project's cash flow, allowing the agent's continuation value to reflect upwards.

Our model's temporary shutdown is different from that of [Zhu \(2013\)](#), who shows that the principal can relax incentives to either pay the agent with private benefits instead of cash or relax incentives to prevent termination/default after bad cash flow realizations. [Zhu \(2013\)](#) has fixed project scale, so when incentives are relaxed the principal cannot also restrict private benefits to the agent. Further, the principal might have to shut down incentives before termination would otherwise happen (incentives can be relaxed at $W > R$) to ensure that the agent's continuation value is high enough to support both private benefits and the continued project.

While our model does feature an incentives shutdown, the mechanism is very different because the project's size is endogenous. In our model, the principal has to pay an ongoing cost of resources (rK_t) in order to fund the project; the level of resources determines the project's scale and the potential private benefits available to the agent. Since paying the agent with private benefits delivers benefits that are less than the principal's resource cost ($\lambda \leq r$), the principal always couples zero incentives with zero project size. This prevents the agent from receiving additional negative cash flow shocks, so the agent's continuation value drifts upward, and the project can continue. Because this shutdown is costly, the principal delays as much as possible, and our firm does not shut down until the last moment, $W_t = R$.

A.2 Proof of Proposition 1

This proof is a slightly modified version of a proof in [Piskorski and Westerfield \(2016\)](#), which in turn is closely based on a proof in [Sannikov \(2008\)](#). First, define the agent's total expected utility received under a contract conditional on his information at time t as:

$$U_t = E^{\hat{K}, \hat{\sigma}} \left[\int_0^\tau e^{-\gamma u} dC_u + \int_0^\tau e^{-\gamma u} \lambda \hat{\Delta}_u du + e^{-\gamma \tau} R | \mathcal{F}_t \right],$$

where $\hat{\Delta}_u = K_u - \hat{K}_u$. We note that the process $U = \{U_t, \mathcal{F}_t; 0 \leq t < \tau\}$ is an $\mathcal{F}_t$ -martingale. The expectation is taken with respect to the probability measure induced by $\{\hat{K}, \hat{\sigma}\}$, such that $\int_0^T \Sigma_t dZ_t^{\hat{K}, \hat{\sigma}} dt = \int_0^T dY_t - \int_0^T f(\hat{K}_t) \mu(\hat{\sigma}_t) dt$ is a martingale for all $T > 0$. Recall here that σ_t , $\hat{\sigma}_t$, K_t , and $\hat{K}_t$ are all bounded from below by zero and from above by positive constants, so Σ_t is also bounded below by zero and above by a positive constant. Then, by the martingale representation theorem for Lévy processes, there exists a $\mathcal{F}_t$ -predictable, integrable process β such that

$$U_t = U_0 + \int_0^t e^{-\gamma u} \beta_u \Sigma_u dZ_u^{\hat{K}, \hat{\sigma}}. \quad (23)$$

Recall the agent's continuation value $W_t^{\hat{K}, \hat{\sigma}}$ defined in (4). We have

$$U_t = \int_0^t e^{-\gamma u} dC_u + \int_0^t e^{-\gamma u} \lambda \hat{\Delta}_u du + e^{-\gamma t} W_t^{\hat{K}, \hat{\sigma}} \quad (24)$$

for $t \leq \tau$. Differentiating (23) and (24), we obtain

$$dU_t = e^{-\gamma t} \beta_t \Sigma_t dZ_t^{\hat{K}, \hat{\sigma}} = e^{-\gamma t} dC_t + e^{-\gamma t} \lambda \hat{\Delta}_t dt - \gamma e^{-\gamma t} W_t^{\hat{K}, \hat{\sigma}} dt + e^{-\gamma t} dW_t^{\hat{K}, \hat{\sigma}},$$

therefore

$$dW_t^{\hat{K}, \hat{\sigma}} = \gamma W_t^{\hat{K}, \hat{\sigma}} dt - dC_t - \lambda \hat{\Delta}_t dt + \beta_t \Sigma_t dZ_t^{\hat{K}, \hat{\sigma}}.$$

This equation also implies the evolution of promised value given in (6), and the evolution given in (9) for $\{\hat{K}, \hat{\sigma}\} = \{K, \sigma\}$.

Next, define $\tilde{U}_t$ to be the payoff to a strategy $\{\tilde{K}, \tilde{\sigma}\}$ that consists of following an arbitrary strategy until time $t < \tau$ and then $\{K, \sigma\}$ thereafter, then

$$\tilde{U}_t = \int_0^t e^{-\gamma u} dC_u + \int_0^t e^{-\gamma u} \lambda \tilde{\Delta}_u du + e^{-\gamma t} W_t^{K, \sigma}.$$

Differentiating $\tilde{U}_t$ and combining terms yields

$$\begin{aligned} e^{\gamma t} d\tilde{U}_t &= \lambda (K_t - \tilde{K}_t) dt + \beta_t \Sigma_t dZ_t^{K, \sigma} \\ &= \lambda (K_t - \tilde{K}_t) dt + \beta_t (f(\tilde{K}_t) \mu(\tilde{\sigma}_t) - f(K_t) \mu(\sigma_t)) dt + \beta_t \Sigma_t dZ_t^{\tilde{K}, \tilde{\sigma}} \end{aligned}$$

where the second equality reflects a change in the probability measure from the one induced by $\{K, \sigma\}$ to the one induced by $\{\tilde{K}, \tilde{\sigma}\}$. Recall here that both $\{K, \sigma\}$ and $\{\tilde{K}, \tilde{\sigma}\}$ must generate the same Σ , per the agent's constraints. If (7) does not hold on a set of positive measure, then the agent could chose $\{\tilde{K}, \tilde{\sigma}\}$ such that

$$\beta_t f(\tilde{K}_t) \mu(\tilde{\sigma}_t) - \lambda \tilde{K}_t > \beta_t f(K_t) \mu(\sigma_t) - \lambda K_t.$$

So, the drift of $\tilde{U}$ is always nonnegative and strictly positive on a set of positive measure, which implies

$$E^{\tilde{K}, \tilde{\sigma}} [\tilde{U}_t] > \tilde{U}_0 = W_0^{K, \sigma},$$

and that the strategy $\{K, \sigma\}$ would not be optimal for the agent. If (7) does hold for the strategy $\{K, \sigma\}$ then $\tilde{U}_t$ is a super-martingale (under the measure induced by $\{\tilde{K}, \tilde{\sigma}\}$) for any strategy $\{\tilde{K}, \tilde{\sigma}\}$, that is,

$$E^{\tilde{K}, \tilde{\sigma}} [\tilde{U}_t] \leq \tilde{U}_0 = W_0^{\tilde{K}, \tilde{\sigma}},$$

which proves that choosing $\{K, \sigma\}$ is optimal for the agent if and only if (7) holds for the strategy $\{K, \sigma\}$.

Finally, incentive compatibility requires that $\beta \geq 0$ because otherwise it is impossible for (7) to hold at $\hat{K} = K > 0$. If $K_t = 0$, then $\hat{K}_t = 0$ and no deviation is possible, so the value of β_t is irrelevant and we say $\beta_t = 0$.

Substituting the constraint $\hat{\sigma} = \frac{\Sigma}{f(\hat{K})}$ into (7) yields a new maximization problem

$$K = \arg \max_{0 \leq \hat{K} \leq K} \beta f(\hat{K}) \mu \left(\frac{\Sigma}{f(\hat{K})} \right) - \lambda \hat{K}. \quad (25)$$

The first order condition is

$$\beta f'(\hat{K})\mu(\hat{\sigma}) - \beta f'(\hat{K})\mu'(\hat{\sigma})\hat{\sigma} - \lambda = 0. \quad (26)$$

Meanwhile, the second order condition is

$$\beta f''(\hat{K})(\mu(\hat{\sigma}) - \mu'(\hat{\sigma})\hat{\sigma}) + \beta f(\hat{K})\mu''(\hat{\sigma})\hat{\sigma}^2 < 0, \quad (27)$$

where the inequality follows from $f''(K) < 0$, $\mu''(\sigma) < 0$, and $\beta \geq 0$ (which implies $\mu(\hat{\sigma}) - \mu'(\hat{\sigma})\hat{\sigma} > 0$). Therefore, there is a unique maximum described by the first-order condition. Rearranging (26), and remembering $0 \leq \hat{K} \leq K$, yields (8). Substitution of $\{\hat{\sigma}, \hat{K}\} = \{\sigma, K\}$ into (6) yields (9).

A.3 Proof of Propositions 2 and 5

We prove Proposition 5, of which Proposition 2 is a corollary. We start with four preliminary results:

Lemma 6 *Any solution to the principal's value function problem must have consumption awarded to the agent at W^C such that $W \leq W^C$, with $F'(W^C) = -1$ and $F''(W^C) = 0$.*

These are the value matching and super-contact conditions.

Lemma 7 *Let $\hat{F}$ be a C^2 solution to (22). Assume that the boundary conditions in Lemma 6 are met at some $\hat{W}^C$, but ignore the boundary condition at $W = R$. Then $\hat{F}''(W) < 0$ for all $W < \hat{W}^C$ on which $\hat{F}$ is defined.*

Proof. Taking the derivative (from the left or the right) of the HJB equation (22) with respect to W and using the envelope theorem yields

$$0 = (\gamma - r)F'(W) + \gamma W F''(W) + \frac{1}{2}\Sigma^2(W)\beta^2(W)F'''(W) \quad (28)$$

If $F'' = 0$ and $F''' < 0$, then (28) implies $F' > 0$. Moreover, from (22), $F'' = 0$ and $F' > 0$ together imply $rF > \max[f(K)\mu(\sigma) - rK]$, which is impossible since $\max[f(K)\mu(\sigma) - rK]$ characterizes the solution under the first-best scenario. Therefore $F'' = 0$ and $F''' < 0$ are jointly impossible: if $F'' = 0$, then it must be that $F''' > 0$ (i.e., F'' can only cross zero from below). Therefore, if $F''(\hat{W}^C) = 0$ for some $\hat{W}^C$, then $F'' < 0$ for $W < \hat{W}^C$. ■

Lemma 8 *Let $\hat{F}$ and $\tilde{F}$ be C^2 solutions to (22). Assume that the boundary conditions in Lemma 6 are met at $\hat{W}^C$ and $\tilde{W}^C$, respectively, but ignore the boundary conditions at R . Then the following four statements are equivalent:*

- $\tilde{W}^C < \hat{W}^C$

- $\tilde{F}(W) > \hat{F}(W)$ for all $W \leq \tilde{W}^C$ such that $\hat{F}$ and $\tilde{F}$ both exist.
- $\tilde{F}'(W) < \hat{F}'(W)$ for all $W \leq \tilde{W}^C$ such that $\hat{F}$ and $\tilde{F}$ both exist.
- $\tilde{F}''(W) > \hat{F}''(W)$ for all $W \leq \tilde{W}^C$ such that $\hat{F}$ and $\tilde{F}$ both exist.

The arguments in [Piskorski and Westerfield \(2016\)](#), Lemma 8, are sufficient.

Lemma 9 *Let $\hat{F}$ be a $\mathcal{C}^2$ solution to (22). Assume that the boundary conditions in Lemma 6 are met at some $\hat{W}^C$, but ignore the boundary condition at $W = R$. Then, if the principal chooses $K(\tilde{W}) = 0$, the law of motion for W_t (9, with $\Sigma = 0$) implies that W_t reflects upward at $\tilde{W}$. The principal optimally chooses $K = 0$ only at $W = R$.*

Proof. First, any solution to the ODE characterized by (22) with $K(\tilde{W}) = 0$ for any $\tilde{W} > R$ can be improved by moving the $K = 0$ transition point to the left. The principle of optimality²⁸ implies that at $\tilde{W}$ for which $K(\tilde{W}) = 0$, we must have continuous $F(W)$ (the value-matching condition) and continuous $F'(W)$ (the smooth pasting condition). In addition, if the highest $\tilde{W}$ for which $K(\tilde{W}) = 0$ lies in the interior of W (i.e., $\tilde{W} > R$), then $F''(W)$ must be continuous as well (the super-contact condition). Finally, there can be at most one value of $\tilde{W}$ for which a transition exists.

Our problem makes the continuity of $F''(W)$ at $\tilde{W}$ impossible. With $\tilde{W} > R$ such that $K(\tilde{W}) = 0$, (22) becomes $0 = -rF(W) + \gamma WF'(W)$ for $R \leq W \leq \tilde{W}$. Taking the derivative, solving for $F'(W)$, and substituting that back into $0 = -rF(W) + \gamma WF'(W)$ yields $rF(W) = -\frac{\gamma^2}{\gamma-r}W^2F''(W)$ for $R \leq W \leq \tilde{W}$. Since $F(\tilde{W})$ is bounded above by the first-best, it must be the case that $F''(\tilde{W})$ is bounded from below by a constant and this bound does not depend on k_0 . However, inspection of the first-order condition of (22) for K , remembering that σ and β are bounded, implies that $F''(\tilde{W})$ is arbitrarily negative if k_0 is arbitrarily small. This is a contradiction, so the second derivative cannot be continuous, and the transition from $K = k_0$ to $K = 0$ cannot be optimal on the interior of W .

In addition, there cannot be a transition from $K > k_0$ to $K = 0$ on the interior of W , using the continuity of $F''(W)$ and the maximization conditions for K . ■

We now proceed to analyze the HJB equation (22).²⁹ A necessary condition for optimality (Lemma 9) is that $K \geq k_0$ on $W \in (R, W^C]$, so we first consider that region first. Re-writing (22), we have

$$F''(W) = \min_{K \geq k_0; \sigma \in (\underline{\sigma}, \bar{\sigma})} \frac{rF(W) - f(K)\mu(\sigma) - rK - \gamma WF'(W)}{\frac{1}{2}\beta(K, \sigma)^2 K^2 \sigma^2} \quad (29)$$

²⁸See [Dumas \(1991\)](#) for a detailed theoretical discussion or [Piskorski and Westerfield \(2016\)](#) or [Zhu \(2013\)](#) for applications in a similar setting.

²⁹This part of the proof is based on a similar proof in [Sannikov \(2008\)](#), altered to this setting and extended to include the possibility that the agent's volatility might be zero.

The right hand side can be written as the function $H_{K,\sigma}(W, F(W), F'(W))$, which is differentiable in all of its arguments. Since K and β are bounded away from zero and infinity, $H_{K,\sigma}$ has uniformly bounded derivatives in $F'(W)$ and $F(W)$, and $H_{K,\sigma}$ is Lipschitz continuous in $F(W)$ and $F'(W)$. It follows that solutions to (29) exist, are $\mathcal{C}^2$ and are unique and continuous in initial conditions.³⁰ Inspection of (29) and (28) shows that the solution is at least $\mathcal{C}^3$, rather than just $\mathcal{C}^2$ on $W \in (R, W^C)$.

To complete our solution, we need to show that our remaining condition can be met: that for any L , either there exists a W^C such that $F(R; W^C) = L$ and $K(R) > 0$ or there exists a W^C such that $F(R; W^C) \geq L$ and $K(R) = 0$. We need only consider $K = 0$ at R from Lemma 9.

Lemma 8 shows that proposed solutions to (22) that obey the right boundary condition can be ranked by the proposed value $\hat{W}^C$. For $\hat{W}^C = R$, we have the highest solution with $rF(R; \hat{W}^C) = \max[f(K)\mu(\sigma) - rK] - \gamma R$. As we increase $\hat{W}^C$, $F(W; \hat{W}^C)$ declines for all W ; thus, $\lim_{W \rightarrow R} F(W; \hat{W}^C)$ declines as $\hat{W}^C$ increases.

Within the set of solutions that obey the right boundary condition, we consider the subset with $K(R; \hat{W}^C) > 0$. Label the infimum value of $F(R; \hat{W}^C)$ as L^* , and the corresponding value of $\hat{W}^C$ as W^{C*} . Any solution with $K(R; \hat{W}^C) > 0$ must be better than the principal's value with $K = 0$ and immediate payment (flow $\gamma R dt$), which is finite. Thus, L^* and W^{C*} are both finite.

Because solutions are continuous in initial conditions, L^* is also the maximum value of $F(R; \hat{W}^C)$ among all proposed solutions with $K(R) = 0$. Then, consider the proposed solution with W^{C*} such that $K(R, W^{C*}) = 0$ and $F(W; W^{C*}) = L^*$. This solution can be implemented for any value of L : because $K(R; W^{C*}) = 0$, there is no termination and L is never realized. This solution is preferred to all other no-termination solutions by construction.

If, instead, $L > L^*$, then the principal prefers to allow termination, with $K(R) > 0$. Since L is greater than the infimum value of $F(R; \hat{W}^C)$ among all proposed solutions with $K(R) > 0$, and solutions to the ODE are continuous in initial conditions and ordered (Lemma 8), there exists exactly one solution with $\hat{W}^C < W^{C*}$ such that $F(R; \hat{W}^C) = L$.

This shows existence and uniqueness of solutions to the HJB equation with a given liquidation value L : If $L > L^*$, then $K(W) > 0$ and solutions exist and are unique by the arguments given above. If $L < L^*$, then we use the solution that generates $K(R) = 0$ and $F(R) = L^*$, which exists and is unique.

Standard existence and uniqueness results (see, e.g., Karatzas and Shreve, 1998), expanded to include Sticky Brownian Motions (see arguments from e.g., Harrison and Lemoine, 1981 or Engelbert and Peskir, 2014) are sufficient to show that (9) has a unique solution and

³⁰To see these conditions directly, use one of the first order conditions (K or σ) to solve out F'' , and observe that, for regions in which $\{K, \sigma\}$ are interior, (29) is a first-order ODE that can be solved through direct integration, taking $K(W)$ and $\sigma(W)$ as unknown, bounded functions. Similarly, in regions in which $K = k_0$, we have constant or bounded coefficients and standard results imply existence and uniqueness there (see e.g., Piskorski and Westerfield, 2016 or Zhu, 2013). We use the more powerful Lipschitz continuity to show that the solution is $\mathcal{C}^2$.




Figure 7: We generate right-to-left shooting solutions to the HJB ODE equation by varying the right boundary point W^C along the line described by $rF(W^C) = \max(CF) - \gamma W^C$ (dot-dash grey line). Our parameter choices implies $L^* = 0$. The solid line is the solution with $F(R=0) = 0 = L^*$. The dotted lines are also solutions to the HJB, but they have $K(W > R) = 0$ and achieve lower values for the principal than the solid black line. The small circles indicate values of W with $K(W) = 0$, at which point W reflects upward. The higher, dashed lines are solutions with a high value of L and default in equilibrium, $F(R) = L > L^*$ and $K(R=0) > 0$. The plot uses $\mu(\sigma) = 0.07 + .5(\sigma^2 - 0.05^2)^{\frac{1}{2}} - 0.55\sigma$ and $f(K) = 3K^{\frac{1}{2}}$.

that W_t is a Sticky Brownian Motion near W_R if $K(W_R) = 0$.

The proof is completed with a standard dynamic-programming verification argument. We illustrate $F(W)$ under various parameter choices in Fig. 7, with the consumption boundary $rF(W^C) = \max(CF) - \gamma W^C$ obtained by substituting in the consumption boundary conditions ($F'(W^C) = -1$ and $F''(W^C) = 0$) into the HJB equation.

Corollary 10 *If $R = 0$, then $L^* = 0$. $L > 0$ implies the principal chooses $K > 0$ for all W , and the boundary value is $F(R=0) = L$, and $F(W)$ is $\mathcal{C}^3$.*

Proof. If $R = 0$ and $K = 0$, then $dW = 0dt$, and no further consumption is realized for either party. Thus, the principal chooses to obtain L if $L \geq 0$, i.e., $L^* = 0$. ■

A.4 Proof of Proposition 3

The following preliminary lemma is useful throughout the proof:

Lemma 11 *If L is large enough, then $F'''(W) > 0$.*

Proof. Taking the derivative of the HJB equation (22) with respect to W and using the envelope theorem yields

$$0 = (\gamma - r) F'(W) + \gamma W F''(W) + \frac{1}{2} g^2(K) h^2(\sigma) F'''(W). \quad (30)$$

Using (30), $F''(W) \leq 0$ and $F'(R) < 0$ are sufficient to show that $F'''(W) > 0$. $F''(W) \leq 0$ follows from Lemma 7, so we need only show that if L is large enough, then $F'(R; L) < 0$. Lemma 8 shows that $F'(R; L)$ is monotonically decreasing in L . If $L = CF(K^{FB}, \sigma^{FB})/r$, then $W^C = R$ and $F'(R) = -1$. Thus, there exists an $L^* < CF(K^{FB}, \sigma^{FB})/r$ such that if $L > L^*$, we have $-1 < F'(R; L) < 0$. ■

We now prove Proposition 3. First, direct substitution of the boundary conditions at W^C into the principal's HJB equation (22) yields the result for $W = W^C$.

Next we begin with the results for K and σ . We examine Σ and β below. For K and σ , we have from (22):

$$\begin{aligned} rF(W) &= \max_{K, \sigma} \left[CF(K, \sigma) + \gamma W F'(W) + \frac{1}{2} g(K)^2 h(\sigma)^2 F''(W) \right] \\ FOC(K) : & CF_K(K, \sigma) + g'(K) g(K) h^2(\sigma) F''(W) \\ FOC(\sigma) : & CF_\sigma(K, \sigma) + g^2(K) h'(\sigma) h(\sigma) F''(W). \end{aligned} \quad (31)$$

where $CF(K, \sigma) = f(K)\mu(\sigma) - rK$, $g(K) = f(K)/f'(K)$, $h(\sigma) = \lambda\sigma/(\mu(\sigma) - \sigma\mu'(\sigma))$. Assumption A.1 is sufficient to ensure concavity of the HJB equation with respect to K and σ . We remember that $CF_\sigma(K, \sigma^{FB}) = 0$ for any K and $CF_K(K, \sigma) > 0$ implies $K < K^{FB}$ for any σ .

Because $g'(K) \geq 0$ and $F''(W) < 0$ for $W < W^C$, $FOC(K)$ implies $K < K^{FB}$ for $W < W^C$. In addition, $K(W)$ increases in W if and only if $F''(W)$ increases in W , by the implicit function theorem applied across both FOC conditions. $F'''(W)$ is positive if L is large enough, from Lemma 11.

For σ , a preliminary: direct substitution of (8) into (14) yields $\varepsilon_{\beta, \sigma}^{FB} = \frac{\sigma^2 \mu''(\sigma)}{\mu(\sigma) - \sigma \mu'(\sigma)} \Big|_{\sigma=\sigma^{FB}} < 0$. Further, some algebra shows that $\varepsilon_{\beta, \sigma}^{FB} < -1$ if and only if $h'(\sigma^{FB}) < 0$.

Then, we observe that if $h'(\sigma(\tilde{W})) = 0$ for any $\tilde{W}$, then $FOC(\sigma)$ implies $CF_\sigma = 0$. In turn, $CF_\sigma = 0$ implies $\sigma(\tilde{W}) = \sigma^{FB}$, regardless of K . In fact, σ^{FB} solves $FOC(\sigma)$ everywhere, so $\sigma(W) = \sigma^{FB}$ for all W , and also $h'(\sigma^{FB}) = 0$. If, instead, $h'(\sigma^{FB}) \neq 0$, then, since $h'(\sigma(W))$ must be continuous in W , $h'(\sigma(\tilde{W})) = 0$ is impossible, and the sign of $h'(\sigma(W))$ must be the same for all W .³¹

Then, if $h'(\sigma^{FB}) > 0$, we have $h'(\sigma(W)) > 0$ for all $W < W^C$. So, $FOC(\sigma)$ im-

³¹More intuitively, $\mu(\sigma)$ is both a benefit function for the principal through cash flows and a cost function through incentives and volatility. If, by coincidence, the benefits are maximized unconditionally at the same value of σ for which the costs are minimized unconditionally, then that value of σ is always the solution to the principal's problem. Otherwise, there are two non-interacting regions of the parameter space: $h' > 0$ in which the cost of incentives pushes σ down, or $h' < 0$ in which the cost of incentives pushes σ up.

plies $\sigma(W) < \sigma^{FB}$ for any $W < W^C$. In addition, $\sigma(W)$ increases in W if and only if $F''(W)$ increases in W by the implicit function theorem. Similarly, if $h'(\sigma^{FB}) < 0$, we have $h'(\sigma(W)) < 0$ for all $W < W^C$. So, $FOC(\sigma)$ implies $\sigma(W) > \sigma^{FB}$ for any $W < W^C$. In addition, $\sigma(W)$ decreases in W if and only if $F''(W)$ increases in W by the implicit function theorem. $F'''(W)$ is positive if L is large enough, from Lemma 11.

Second, (22) can be written as

$$\begin{aligned} rF(W) &= \max_{\Sigma, \beta} \left[CF(\Sigma, \beta) + \gamma W F'(W) + \frac{1}{2} \Sigma^2 \beta^2 F''(W) \right] \\ FOC(\Sigma) : & CF_{\Sigma}(\Sigma, \beta) + \Sigma \beta^2 F''(W) \\ FOC(\beta) : & CF_{\beta}(\Sigma, \beta) + \Sigma^2 \beta F''(W) \end{aligned} \quad (32)$$

where $CF(\beta, \Sigma) = CF(K(\beta, \Sigma), \sigma(\beta, \Sigma))$. Because the first-best has $CF_K(K^{FB}, \sigma^{FB}) = 0$ and $CF_{\sigma}(K^{FB}, \sigma^{FB}) = 0$, there are corresponding values of $\{\beta^{FB}, \Sigma^{FB}\}$ that have $CF_K(K(\beta^{FB}, \Sigma^{FB}), \sigma(\beta^{FB}, \Sigma^{FB})) = 0$ and $CF_{\sigma}(K(\beta^{FB}, \Sigma^{FB}), \sigma(\beta^{FB}, \Sigma^{FB})) = 0$.

The second order conditions for the HJB equation (32), evaluated at $\{\beta^{FB}, \Sigma^{FB}\}$ are

$$\begin{aligned} SOC(\beta)|_{(\beta^{FB}, \Sigma^{FB})} : & CF_{KK} \cdot (K_{\beta})^2 + CF_{\sigma\sigma} \cdot (\sigma_{\beta})^2 < 0 \\ SOC(\Sigma)|_{(\beta^{FB}, \Sigma^{FB})} : & CF_{KK} \cdot (K_{\Sigma})^2 + CF_{\sigma\sigma} \cdot (\sigma_{\Sigma})^2 < 0, \end{aligned} \quad (33)$$

where the inequalities follow from $CF_{KK} < 0$ and $CF_{\sigma\sigma} < 0$.

Because the first-best is unique, with $CF_{\beta} = 0$ and $CF_{\Sigma} = 0$, we also have that $CF_{\beta} > 0$ implies $\beta < \beta^{FB}$ and $CF_{\Sigma} > 0$ implies $\Sigma < \Sigma^{FB}$. Then, $F''(W) < 0$ for $W < W^C$, $FOC(\Sigma)$, and $FOC(\beta)$ together imply $\Sigma(W) < \Sigma^{FB}$ and $\beta(W) < \beta^{FB}$ for $W < W^C$. In addition, $\Sigma(W)$ and $\beta(W)$ increase in W if and only if $F''(W)$ increases in W , by the implicit function theorem. $F'''(W)$ is positive if L is large enough, from Lemma 11.

Finally, we generalize the definition of optimal contract to allow for misallocation in equilibrium. Assuming that the principal offers the agent a recommended level of resource misallocation, δ_t , to go with the assigned level of resources, $K_t + \delta_t$. Then the contract is incentive compatible if it implements $\hat{K}_t = K_t$ and $\hat{\delta}_t = \delta_t$. The incentive compatibility condition (7) becomes

$$\{K_t, \delta_t, \sigma_t\} = \arg \max_{\hat{K}_t + \hat{\delta}_t = K_t + \delta_t; f(\hat{K}_t) \hat{\sigma}_t = f(K_t) \sigma_t} \left[\beta_t f(\hat{K}_t) \mu(\hat{\sigma}_t) + \lambda (K_t + \delta_t - \hat{K}_t) \right]. \quad (35)$$

Note that the agent-controlled portion of the right hand side is the same as in the original IC condition (7). The first order conditions imply that β_t becomes

$$\beta_t = \beta(\sigma_t, K_t) = \frac{\lambda}{f'(K_t)} \times \frac{1}{\mu(\sigma_t) - \mu'(\sigma_t) \sigma_t}. \quad (36)$$

The evolution of the agent's continuation value becomes

$$dW_t = \gamma W_t dt + \beta_t \Sigma_t dZ_t - \lambda \delta_t dt - dC_t. \quad (37)$$

The principal's HJB equation (22) becomes

$$rF(W) = \max_{K + \delta \in \{0 \cup [k_0, \infty)\}; \sigma \in (\underline{\sigma}, \bar{\sigma})} \left[f(K) \mu(\sigma) - rK - r\delta + (\gamma W - \lambda \delta) F'(W) + \frac{1}{2} \beta^2(K, \sigma) f^2(K) \sigma^2 F''(W) \right] \quad (38)$$

For any given choice of K , the objective function of (38) is linear in δ with coefficient $-r - \lambda F'(W)$. Since $F'(W) \geq -1$, we have $-r - \lambda F'(W) < 0$. Therefore, $\delta = 0$ is (weakly) optimal.

A.5 Proof of Proposition 4

Consider the implementation in two stages. First, we solve the agent's problem conditional on an arbitrary choice (sequence) of $\{\tilde{\Sigma}, \tilde{\beta}\}$, as in Proposition 1 with a modified cash flow technology. The principal offers the agent the adjusted technology with

$$d\tilde{Y}_t^{NEW} = f(\hat{K}_t) \mu \left(\frac{\tilde{\Sigma}_t}{f(\hat{K}_t)} \right) dt - \tilde{K}_t \theta(\tilde{\Sigma}_t, \tilde{\beta}_t) dt - \phi(\tilde{\Sigma}_t, \tilde{\beta}_t) dt + \tilde{\Sigma}_t dZ_t, \quad (39)$$

where we have set $\hat{\sigma} = \frac{\tilde{\Sigma}}{f(\hat{K})}$. Then, the probability measure induced by $\{\hat{K}, \hat{\sigma}, \tilde{K}\}$ is such that

$$\tilde{\Sigma}_t dZ_t^{\hat{K}, \hat{\sigma}, \tilde{K}} = \left(d\tilde{Y}_t^{NEW} - f(\hat{K}_t) \mu \left(\frac{\tilde{\Sigma}_t}{f(\hat{K}_t)} \right) dt + \tilde{K}_t \theta(\tilde{\Sigma}_t, \tilde{\beta}_t) dt + \phi(\tilde{\Sigma}_t, \tilde{\beta}_t) dt \right)$$

is an $\mathcal{F}_t$ -martingale (following the construction in Proposition 1). The incentive compatibility arguments in Proposition 1 also follow, with the modification that the agent's incentive compatibility condition is

$$\{\tilde{K}, \hat{K}\} = \arg \max \left[\tilde{\beta}_t f(\hat{K}_t) \mu \left(\frac{\tilde{\Sigma}_t}{f(\hat{K}_t)} \right) - \tilde{\beta}_t \tilde{K}_t \theta(\tilde{\Sigma}_t, \tilde{\beta}_t) + \lambda(\tilde{K}_t - \hat{K}_t) \right]. \quad (40)$$

Substituting $\theta = \frac{\lambda}{\tilde{\beta}_t}$ (from Eq. 18) yields

$$\{\tilde{K}, \hat{K}\} = \arg \max \left[\tilde{\beta}_t f(\hat{K}_t) \mu \left(\frac{\tilde{\Sigma}_t}{f(\hat{K}_t)} \right) - \lambda \hat{K}_t \right], \quad (41)$$

which is identical to (7), with $\tilde{\beta}$ replacing β and $f(\hat{K})\hat{\sigma} = \tilde{\Sigma}$ as the constraint. As in Proposition 1, the first-order condition is necessary and sufficient and yields the solution analogous to (8). That is, $\hat{K}_t = K(\tilde{\Sigma}_t, \tilde{\beta}_t)$, which solves

$$K(\tilde{\Sigma}_t, \tilde{\beta}_t) : f'(K_t) \left[\mu \left(\frac{\tilde{\Sigma}_t}{f(K_t)} \right) - \mu' \left(\frac{\tilde{\Sigma}_t}{f(K_t)} \right) \left(\frac{\tilde{\Sigma}_t}{f(K_t)} \right) \right] = \frac{\lambda}{\tilde{\beta}_t}. \quad (42)$$

Further, the maximization condition does not depend on $\tilde{K}$, so the agent is indifferent across choices of $\tilde{K}_t \geq \hat{K}_t$, while preferring any $\tilde{K}_t \geq \hat{K}_t$ to any $\tilde{K}_t < \hat{K}_t$. Thus, we can set $\tilde{K}_t = \hat{K}_t = K(\tilde{\Sigma}_t, \tilde{\beta}_t)$. Following Proposition 1, the agent's continuation utility evolves as in (9).

Second, we solve for the agent's optimal choice of $\{\tilde{\Sigma}_t, \tilde{\beta}\}$ given (42). Substituting $\tilde{K} = \hat{K} = K(\tilde{\Sigma}, \tilde{\beta})$ back into (39) we obtain

$$d\tilde{Y}_t^{NEW} = f(K(\tilde{\Sigma}_t, \tilde{\beta}_t)) \mu \left(\frac{\tilde{\Sigma}_t}{f(K(\tilde{\Sigma}_t, \tilde{\beta}_t))} \right) dt - K(\tilde{\Sigma}_t, \tilde{\beta}_t) \theta(\tilde{\Sigma}_t, \tilde{\beta}_t) dt - \phi(\tilde{\Sigma}_t, \tilde{\beta}_t) dt + \tilde{\Sigma}_t dZ_t. \quad (43)$$

Substituting in (18) and (19), we see (43) reduces to $dY_t^{NEW} = \tilde{\Sigma}_t dZ_t$, meaning the drift is equal to zero for any choice of $\{\tilde{\Sigma}, \tilde{\beta}\}$. Plugging $dY_t^{NEW} = \tilde{\Sigma}_t dZ_t$ into (17) yields

$$dM_t = \gamma M_t dt + \tilde{\beta}_t \tilde{\Sigma}_t dZ_t - dC_t. \quad (44)$$

Standard arguments comparing (44) to (9, with $\{\tilde{\Sigma}_t, \tilde{\beta}_t\}$) imply that these results are sufficient for $W_t = M_t$, $dC_t = \max(M_t - W^C, 0)$, and for the agent to quit at $M_t = 0$ (see, e.g., DeMarzo and Sannikov, 2006; Biais, Mariotti, Plantin, and Rochet, 2007). By the law of iterated expectations, the agent is indifferent to the volatility of his continuation value. Since the agent is indifferent across $\{\tilde{\Sigma}_t, \tilde{\beta}_t\}$, he chooses the principal's desired values.

References

- Ai, H. and R. Li (2015). Investment and CEO compensation under limited commitment. *Journal of Financial Economics* 116(3), 452–472.
- Andersen, L., D. Duffie, and Y. Song (2019). Funding value adjustments. *Journal of Finance* 74(1), 145–192.
- Andrade, G. and S. N. Kaplan (1998). How costly is financial (not economic) distress? evidence from highly leveraged transactions that became distressed. *Journal of Finance* 53(5), 1443–1493.
- Aragon, G. O. and V. Nanda (2011). Tournament behavior in hedge funds: High-water marks, fund liquidation, and managerial stake. *Review of Financial Studies* 25(3), 937–974.
- Baldenius, T. (2008). Internal pricing. *Foundations and Trends in Accounting* 3(4), 223–313.
- Baldenius, T., S. Dutta, and S. Reichelstein (2007). Cost allocation for capital budgeting decisions. *The Accounting Review* 82(4), 837–867.
- Baldenius, T., S. Reichelstein, and S. A. Sahay (1999). Negotiated versus cost-based transfer pricing. *Review of Accounting Studies* 4(2), 67–91.
- Bertrand, M. and S. Mullainathan (2003). Enjoying the quiet life? Corporate governance and managerial preferences. *Journal of Political Economy* 111(5), 1043–1075.
- Biais, B., T. Mariotti, G. Plantin, and J.-C. Rochet (2007). Dynamic security design: Convergence to continuous time and asset pricing implications. *Review of Economic Studies* 74, 345–390.
- Biais, B., T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010). Large risks, limited liability, and dynamic moral hazard. *Econometrica* 78, 73–118.
- Bolton, P., N. Wang, and J. Yang (2019). Optimal contracting, corporate finance, and valuation with inalienable human capital. *Journal of Finance* 74(3), 1363–1429.
- Buffa, A. M., D. Vayanos, and P. Woolley (2015). Asset management contracts and equilibrium prices. Unpublished working paper, NBER.
- Cadenillas, A., J. Cvitanić, and F. Zapatero (2004). Leverage decision and manager compensation with choice of effort and volatility. *Journal of Financial Economics* 73(1), 71–92.
- Carpenter, J. N. (2000). Does option compensation increase managerial risk appetite? *Journal of Finance* 55(5), 2311–2331.

- Chevalier, J. and G. Ellison (1997). Risk taking by mutual funds as a response to incentives. *The Journal of Political Economy* 105(6), 1167–1200.
- Cvitanić, J., D. Possamaï, and N. Touzi (2017). Moral hazard in dynamic risk management. *Management Science* 63(10), 3328–3346.
- DeMarzo, P. and Y. Sannikov (2006). Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance* 61, 2681–2724.
- DeMarzo, P. M., M. J. Fishman, Z. He, and N. Wang (2012). Dynamic agency and the q theory of investment. *Journal of Finance* 67(6), 2295–2340.
- DeMarzo, P. M., D. Livdan, and A. Tchistyi (2013). Risking other people’s money: Gambling, limited liability, and optimal incentives. Unpublished working paper. Stanford University, University of California-Berkeley, and UIUC.
- Dumas, B. (1991). Super contact and related optimality conditions. *Journal of Economic Dynamics and Control* 75(4), 675–685.
- Edmans, A., X. Gabaix, and D. Jenter (2017). Executive compensation: A survey of theory and evidence. In: *The Handbook of the Economics of Corporate Governance* 1, 383–539.
- Eisdorfer, A. (2008). Empirical evidence of risk shifting in financially distressed firms. *Journal of Finance* 63(2), 609–637.
- Engelbert, H.-J. and G. Peskir (2014). Stochastic differential equations for sticky brownian motion. *Stochastics* 86(6), 993–1021.
- Epstein, L. G. and S. Ji (2013). Ambiguous volatility and asset pricing in continuous time. *Review of Financial Studies* 26(7), 1740–1786.
- Feng, F. Z. (2020). Dynamic compensation under uncertainty shocks and limited commitment. *Journal of Financial and Quantitative Analysis*. Forthcoming.
- Gilje, E. P. (2016). Do firms engage in risk-shifting? empirical evidence. *Review of Financial Studies* 29(11), 2925–2954.
- Giroud, X. (2013). Proximity and investment: Evidence from plant-level data. *Quarterly Journal of Economics* 128(2), 861–915.
- Göx, R. F. and U. Schiller (2006). An economic perspective on transfer pricing. *Handbooks of Management Accounting Research* 2, 673–695.
- Graham, J. and C. Harvey (2011). Duke/CFO magazine global business outlook. U.S. topline tables, March 2011, www.cfosurvey.org/pastresults.htm.
- Graham, J. and C. Harvey (2012). Duke/CFO magazine global business outlook. U.S. topline tables, June 2012, www.cfosurvey.org/pastresults.htm.

- Graham, J. R. and C. R. Harvey (2001). The theory and practice of corporate finance: Evidence from the field. *Journal of Financial Economics* 60(2), 187–243.
- Gryglewicz, S. and B. Hartman-Glaser (2020). Investment timing and incentive costs. *Review of Financial Studies* 33(1), 309–357.
- Gryglewicz, S., S. Mayer, and E. Morellec (2020). Agency conflicts and short- versus long-termism in corporate policies. *Journal of Financial Economics* 136(3), 718–742.
- Harris, M., C. H. Kriebel, and A. Raviv (1982). Asymmetric information, incentives and intrafirm resource allocation. *Management science* 28(6), 604–620.
- Harrison, J. M. and A. J. Lemoine (1981). Sticky brownian motion as the limit of storage processes. *Journal of Applied Probability* 18(1), 216–226.
- He, Z. (2011). A model of dynamic compensation and capital structure. *Journal of Financial Economics* 100(2), 351–366.
- Hiemann, M. and S. Reichelstein (2012). Transfer pricing in multinational corporations: An integrated management and tax perspective. In W. Schön and K. Konrad (Eds.), *Fundamentals of International Transfer Pricing in Law and Economics*, pp. 3–18. Heidelberg: Springer-Verlag.
- Hirshleifer, J. (1956). On the economics of transfer pricing. *The Journal of Business* 29(3), 172–184.
- Huang, J., C. Sialm, and H. Zhang (2011). Risk shifting and mutual fund performance. *Review of Financial Studies* 24(8), 2575–2616.
- Hyde, C. E. and C. Choe (2005). Keeping two sets of books: The relationship between tax and incentive transfer prices. *Journal of Economics & Management Strategy* 14(1), 165–186.
- Jacobs, M. T. and A. Sivdasani (2012). Do you know your cost of capital? *Harvard Business Review* 90, 118–124.
- Jagannathan, R., D. A. Matsa, I. Meier, and V. Tarhan (2016). Why do firms use high discount rates? *Journal of Financial Economics* 120, 445–463.
- Karatzas, I. and S. Shreve (1998). *Brownian Motion and Stochastic Calculus* (2nd ed.). Springer-Verlag New York.
- Leung, R. (2017). Continuous-time principal-agent problem with drift and stochastic volatility control. Unpublished working paper. CKGSB.
- Li, R. and N. Williams (2017). Optimal contracts with hidden risk. Unpublished working paper. University of Massachusetts-Boston and University of Wisconsin-Madison.

- Lioui, A. and P. Poncet (2013). Optimal benchmarking for active portfolio managers. *European Journal of Operational Research* 226(2), 268–276.
- Malenko, A. (2019). Optimal dynamic capital budgeting. *Review of Economic Studies* 86(4), 1747–1778.
- Metrick, A. and A. Yasuda (2010). The economics of private equity funds. *Review of Financial Studies* 23(6), 2303–2341.
- Moreno-Bromberg, S. and G. Roger (2017). Scale effects in dynamic contracting. Unpublished working Paper, University of Zurich and University of Sydney.
- Ou-Yang, H. (2003). Optimal contracts in a continuous-time delegated portfolio management problem. *Review of Financial Studies* 16(1), 173–208.
- Pfeiffer, T., U. Schiller, and J. Wagner (2011). Cost-based transfer pricing. *Review of Accounting Studies* 16(2), 219.
- Piskorski, T. and M. M. Westerfield (2016). Optimal dynamic contracts with moral hazard and costly monitoring. *Journal of Economic Theory* 166, 242–281.
- Poterba, J. M. and L. H. Summers (1995). A CEO survey of U.S. companies’ time horizons and hurdle rates. *MIT Sloan Management Review* (Fall), 43–53.
- Prendergast, C. (2002). The tenuous trade-off between risk and incentives. *Journal of Political Economy* 110(5), 1071–1102.
- PriceWaterhouse (1984). Transfer pricing practices of american industry.
- Rajan, M. V. and S. Reichelstein (2004). A perspective on “Asymmetric information, incentives and intrafirm resource allocation”. *Management Science* 50(12).
- Rauh, J. D. (2008). Risk shifting versus risk management: Investment policy in corporate pension plans. *Review of Financial Studies* 22(7), 2687–2733.
- Robinson, D. T. and B. A. Sensoy (2013). Do private equity fund managers earn their fees? Compensation, ownership, and cash flow performance. *Review of Financial Studies* 26(11), 2760–2797.
- Sahay, S. A. (2003). Transfer pricing based on actual cost. *Journal of Management Accounting Research* 15(1), 177–192.
- Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *Review of Economic Studies* 75, 957–984.
- Szydlowski, M. (2019). Incentives, project choice and dynamic multitasking. *Theoretical Economics* 14(3), 813–847.

- Tang, R. Y. (1992). Transfer pricing in the 1990s. *Strategic Finance* 73(8), 22.
- Vaysman, I. (1996). A model of cost-based transfer pricing. *Review of Accounting Studies* 1(1), 73–108.
- Vaysman, I. (1998). A model of negotiated transfer pricing. *Journal of Accounting and Economics* 25(3), 349–384.
- Wong, T.-Y. (2019). Dynamic agency and endogenous risk-taking. *Management Science* 65(9), 4032–4048.
- Zhu, J. Y. (2013). Optimal contracts with shirking. *Review of Economic Studies* 80, 812–839.

Internet Appendix

Dynamic Resource Allocation with Hidden Volatility

by Felix Zhiyu Feng and Mark M. Westerfield

Journal of Financial Economics, Forthcoming

February, 2020

B Model Robustness

B.1 Private Benefits with Volatility

We can extend the model to consider situations in which private benefits generate cash flow volatility. In the model introduced in Section 2, resources can be used to generate a flow of private benefits to the agent equal to $\lambda(K_t - \hat{K}_t)dt$. We now additionally assume that these private benefits create cash flow volatility, so that

$$\Sigma_t = f(\hat{K}_t)\hat{\sigma}_t + \lambda(K_t - \hat{K}_t)\epsilon, \quad (\text{B.1})$$

where ϵ is the cash flow volatility per unit of private benefits. The arguments of Section 3 follow, with

$$\beta(\sigma, K) = \frac{\lambda}{f'(K)(\mu(\sigma) - \sigma\mu'(\sigma)) + \lambda\epsilon\mu'(\sigma)}. \quad (\text{B.2})$$

If $\lambda\epsilon - \sigma f'(K) < 0$, then $\frac{\partial}{\partial\sigma}\beta(\sigma, K) < 0$, as in the baseline model.

In contrast, if $\lambda\epsilon - \sigma f'(K) > 0$, then $\frac{\partial}{\partial\sigma}\beta(\sigma, K) > 0$. This is interpretable: if re-allocating resources from productive to unproductive use greatly increases volatility, then a matching reduction must be made in productive cash flow volatility. Incentives are used to push the other way – to encourage more volatility in the productive assets. However, it is also the case that stronger incentives are used to produce less *efficient* risk taking. Recall from Section 3 that $\mu(\sigma)$ is hump shaped, and $\frac{\mu(\sigma)}{\sigma}$ is decreasing in the relevant region. If $\beta_\sigma > 0$, then stronger incentives and high volatility are associated with a lower return-to-risk ratio.

Similarly, if $\beta_\sigma > 0$ everywhere, then the arguments in Section 3.3 are sufficient to show that $\sigma \leq \sigma^{FB}$. In this case, the principal always reduces project level risk to reduce the volatility of the agent's continuation utility.

The structure of the implementation is unchanged, with different values of $\{\phi, \theta\}$, as long as $\{\beta, \Sigma\}$ is invertible as in Proposition 4.

B.2 A General Private Benefits Function

We can also extend the model in Section 2 to consider general, non-linear specifications for the agent's private benefits from resource misallocation. In the main model, the linearity of private benefits implies that the agent derives a constant marginal benefit from resource misallocation. Combined with the decreasing returns to scale of $f(K)$, this constant benefit can be interpreted as "there is a shortage of good projects but no shortage of bad projects". Generalizing to a non-linear private benefits function can be interpreted to mean that bad projects are variable in their utility.

We can repeat the arguments from Section 3 to derive the incentive compatibility condition with a generic private benefit. Let $\Lambda(\hat{K}, K)$ denote the generic private benefit from resource misallocation and be $\mathcal{C}^2$ with the appropriate shape/concavity. Let $\Lambda_{\hat{K}}(K, K)$ denote $\Lambda_{\hat{K}}(\hat{K}, K)|_{\hat{K}=K}$, which is the principal's desired marginal value of private benefits. Then,

$$\beta(\sigma_t, K_t) = \frac{\Lambda_{\hat{K}}(K_t, K_t)}{f'(K_t)} \times \frac{1}{\mu(\sigma_t) - \mu'(\sigma_t)\sigma_t}. \quad (\text{B.3})$$

The optimal contract will differ from the one in Section 3 in two ways. The first is that required incentives may increase or decrease with project size:

$$\frac{\partial}{\partial K} \beta(\sigma, K) = \left(\frac{\Lambda_{\hat{K}K} + \Lambda_{\hat{K}\hat{K}}}{\Lambda_{\hat{K}}} - \frac{f''(K)}{f'(K)} \right) \beta(\sigma, K), \quad (\text{B.4})$$

which may be positive or negative, unlike (12). If the marginal private benefit of misallocating resources declines with project size, and the decline is faster than the decline in the marginal productivity of productive use, then the agent's gain from misallocation is smaller when managing a larger project. Consequently, the incentives required to prevent misallocation can be lower as project size grows. An example uses $\Lambda(\hat{K}, K) = \lambda(K - \hat{K})/K$ over a limited range of K , with $f(K) = K^\alpha$, $0 < \alpha < 1$. This implies $\beta_K = -\frac{\alpha}{K}\beta < 0$. One can interpret the assumed Λ to mean that the agent's private benefit depends on the *fraction*, instead of level, of misallocated resources.

Second, the fact that β may decrease in K implies the possibility of over-investment in the optimal contract, or $K(W) > K^{FB}$. The intuition closely follows the intuition for overly risky and overly prudent project choice: as W moves away from W^C , the principal seeks to lower the volatility of W . If β is decreasing in K , the principal may achieve a lower level of continuation volatility with a high K and low β . Consequently, when the principal wants to reduce the volatility of the agent's continuation utility, she may find it optimal to either decrease total cash flow volatility (lower σ and/or lower K) or decrease incentives (higher σ and/or higher K), resulting in different combinations of overly prudent/risky project choices and under/over investment in the equilibrium.

Finally, we have been focusing on optimal contracts following Definition 1, i.e. a contract that eliminates resource misallocation in equilibrium because misallocation is inefficient.

We can consider the case in which misallocation is efficient: for example when $\Lambda(\hat{K}, K) = \lambda(\hat{K} - K)$ with $\lambda > r$. However, this specification changes the right boundary of $F(W)$ in an uninteresting way. If resources used for the agent's private benefit are $\delta \leq \bar{\delta}$, the principal's HJB equation becomes

$$rF(W) = \max_{K, \sigma, \delta} \left[f(K)\mu(\sigma) - rK - r\delta + (\gamma W - \lambda\delta)F'(W) + \frac{1}{2}\beta^2(K, \sigma)f^2(K)\sigma^2F''(W) \right].$$

The principal pays private benefits whenever the agent's continuation utility exceeds W^S , defined by $F'(W^S) = -\frac{r}{\lambda}$, instead of the current boundary condition, $F'(W^C) = -1$.

B.3 Dimensionality

Our model uses two-dimensional incentives (β and Σ) to implement a two-dimensional agent's choice (K and σ). What would happen if the agent had more choices, particularly an effort choice? We have assumed that the agent implements a portfolio of projects with a particular risk-return tradeoff, but what if the agent could work to improve the tradeoff? What if the relationship between K and σ were not separable? We might assume that, instead of (1), we have

$$dY_t = \alpha(K_t, \sigma_t, e_t)dt + \varphi(K_t, \sigma_t)dZ_t, \quad (\text{B.5})$$

where e_t is effort. The agent's private benefits function might then be $B(K_t - \hat{K}_t, e_t)$. In this setup, the principal then has two controls (β and $\Sigma = \varphi(K_t, \sigma_t)$) to implement a three-dimensional choice by the agent. The incentive compatibility condition (7) becomes

$$\{K, \sigma, e\} = \arg \max_{\hat{K}, \hat{\sigma}, \hat{e}} \left[\beta\alpha(\hat{K}, \hat{\sigma}, \hat{e}) + B(K - \hat{K}, \hat{e}) \right] \quad (\text{B.6})$$

under the constraint that $\Sigma = \varphi(\hat{K}, \hat{\sigma})$.

Because the principal only has two controls, but the agent has three choices, the set of choices the principal can actually implement is a two-dimensional curve in a three-dimensional space. Depending on the functional forms of α and φ , this implementable set might be very far from the first-best. Informally, the principal can condition on output to provide incentives that resources be used efficiently and that total risk is as desired, but her controls are not more granular than that; adding an effort choice need not increase the size of the set of implementable actions.

A key insight of our model of volatility is that it represents a unique setting that cannot be captured by an economically reasonable hidden effort model, such as a simple variation of [DeMarzo and Sannikov \(2006\)](#). One might think that our model could be made simpler by giving the principal control over volatility and resources and also giving the agent a hidden

effort choice over drift. For example, using $dY_t = f(K_t)e_t dt + \Sigma_t dZ_t$ instead of (1). However, this change of variables would put all the economics into the private benefits function in a completely uninterpretable way.³² Instead, we give the agent a second choice over per-unit-of-resources volatility, give the principal a second control over total volatility, and impose economically reasonable assumptions on the production function $f(K)\mu(\sigma)$. These choices are structured so that equilibrium output could be obtained with positive probability under both the agent's true action and under the principal's desired action, thus allowing us to use the martingale methods of Sannikov (2008). In sum, our model demonstrates that agency problems over the composition of volatility are possible and interesting – total cash-flow volatility (Σ), project volatility (σ), and the agent's continuation value volatility ($\beta\sigma$) are all meaningfully different.

³²In the example, we would have $e_t = \frac{f(\hat{K}_t)}{f(K_t)}\mu\left(\frac{\Sigma_t}{f(\hat{K}_t)}\right)$ as the hidden action instead of $\hat{K}_t$. The agent's private benefits function becomes $B(e, K, \Sigma) = \lambda\left(K - \hat{K}(e, K, \Sigma)\right)$, where $\hat{K}(e, K, \Sigma)$ is an inversion of $e(\hat{K}, K, \Sigma)$. Then, first, $B(e, K, \Sigma)$ has unsigned derivatives (i.e. B_K and B_Σ must change signs at arbitrary point over the relevant range of the problem). Second, even the level of private benefits is difficult to assess because the constraint $B(e(K, K, \Sigma), K, \Sigma) = 0$ is a technological constraint that must be imposed exogenously and has no clear meaning.