

Optimal Dynamic Contracts with Moral Hazard and Costly Monitoring*

Tomasz Piskorski[†]
Columbia University

Mark M. Westerfield[‡]
University of Washington

July 2016

Abstract

We introduce a tractable dynamic monitoring technology into a continuous-time moral-hazard problem and study the optimal long-term contract between principal and agent. Monitoring adds value by allowing the principal to reduce the intensity of performance-based incentives, reducing the likelihood of costly termination. We present a novel characterization of optimal dynamic incentive provision when performance-based incentives may decline continuously to zero. Termination happens in equilibrium only if its costs are relatively low. In general, the intensity of both monitoring and performance-based compensation can be non-monotonic functions of the quality of past performance. Our results can also help explain puzzling empirical findings on the relationship between performance history and future pay-performance sensitivity and on the linkage between termination, performance, and monitoring. We also discuss implications of our model for optimal security design and endogenous financing constraints.

JEL Classification: D82, D86, M52.

Keywords: Monitoring, Dynamic Contracts, Managerial Compensation, Moral Hazard, Endogenous Financing Constraints

*We are grateful to the Editor, the Associate Editor and three anonymous referees for helpful comments and suggestions. We also thank seminar participants at the USC Marshall School, Columbia Business School, NYU Stern, NYU Economics Department, UC-Berkeley Haas School, ASU W.P. Carey School, Rochester Simon School, UT Austin McCombs School, and the North American Summer Meeting of the Econometric Society for their useful discussion and comments.

[†]Associate Professor of Finance and Economics, Columbia Business School, Uris Hall 810, Broadway 3022, NYC, NY 10027; Telephone: 212-854-4655; Fax: 212-662-8474; E-mail: tp2252@columbia.edu.

[‡]Assistant Professor of Finance and Business Economics, Foster School of Business, University of Washington, Box 353226, Seattle, WA 98195; Telephone: 206-543-4578; E-mail: mwesterf@uw.edu.

1 Introduction

Despite its fundamental role as an incentive device, the optimal use of costly monitoring has remained relatively unexamined in dynamic settings. In this paper we fill this gap by showing how compensation, monitoring, and contract termination should be jointly and optimally structured in a dynamic principal-agent framework with moral hazard. The important change relative to earlier work is the addition of a dynamic monitoring technology with variable intensity – a second signal in addition to observed performance – that allows the principal to detect, at a cost, undesirable activity on the part of the agent.

Our continuous time setting allows us to provide a tractable characterization of a fully history dependent optimal contract. As we discuss below, this characterization yields novel insights on the interplay of compensation, monitoring, and termination as a function of observed history of performance. In addition, our model helps explain two puzzling sets of empirical findings in managerial compensation literature – one having to do with the relationship between performance history and future pay-performance sensitivity and a second having to do with the empirical linkages between termination, performance, and monitoring. Finally, we also discuss how use of monitoring relaxes the endogenous financing constraint implied by the agency problem.

We begin with a dynamic agency problem similar to DeMarzo and Sannikov (2006) and Biais et al. (2007) in which the principal hires an agent to operate a project. A particular agent-asset match generates a project that increases the productive value of the asset, but the agent lacks the funds to purchase that asset directly. Instead, the agent is offered a compensation contract by the principal (a firm or outside investors) in exchange for undertaking the project. There is a moral hazard problem because the project’s cash flows are subject to shocks that are privately observed by the agent, and the agent receives a private benefit from taking actions that decrease the payoff from the project (e.g. shirking, diversion, asset misuse, etc.). To induce the agent not to divert resources for his own benefit, the principal can use standard performance-based incentives, where bad performance leads to termination of the contract and good performance leads to bonus payments to the agent. Termination is costly because the principal retains assets that are less valuable without the match to the agent.

We extend this setting by giving the principal access to a monitoring technology that allows her to detect, at a cost, undesirable activity by the agent. The probability of detection

depends on the resources the principal spends on monitoring and on the extent of the agent's bad acts. The monitoring technology can be interpreted as continuous or repeated audits of the project by the principal, or as direct involvement of the principal in the operations of the project. If the agent is caught shirking, the principal can impose penalties on the agent by reducing his future compensation, or at the extreme, by terminating the contract early. Thus, instead of having to observe a series of negative cash flows – either from bad behavior or bad luck – before terminating, the principal can probabilistically observe and punish bad behavior directly.

Active monitoring provides a second incentive device for the principal. Using the monitoring technology allows the principal to reduce the agent's share of the project's gains and losses (to lower performance-based incentives) and to reduce the likelihood of costly termination, while still maintaining incentive compatibility. Monitoring can be made intense enough, at a high cost, to completely substitute for termination-based incentives. Thus, termination may not occur in equilibrium.

The optimal contract efficiently trades off these two ways of providing incentives to the agent and maximizes the total surplus from the project. In our continuous time setting, the optimal contract takes a tractable form. The agent's continuation utility acts as a state variable summarizing the project's past performance that determines the nearness to termination and the efficient trade-off between monitoring and pay-for-performance incentives.

A first result is that the relationship between the intensity of monitoring and measures of past performance is not necessarily monotonic, contrary to the predictions of, for example, Costly State Verification models.¹ To understand this result it is useful to think about the optimal use of monitoring as being shaped by two competing forces: the *risk of costly termination* and the *inside stake*. When the project under-performs and the agent's continuation value falls, then 1) termination is more likely, but also 2) the agent has less to lose. Because termination is costly to the principal as well as the agent, the principal may wish to monitor more so as to lower the implied pay-for performance sensitivity and reduce the risk of costly termination. However, the effectiveness of monitoring also depends on how much the principal can punish the agent. When the agent has less to lose (a lower inside stake), monitoring is less effective, and so the principal may wish to monitor less. When the “cost of termination” effect dominates – which is particularly likely when the loss in termination is high – bad performance will lead to more monitoring and fewer pay-for-performance incentives. In

¹See, for example, Townsend (1979) and Gale and Hellwig (1985).

contrast, if the “inside stake” effect dominates, monitoring will actually decline and pay-for-performance compensation incentives will increase as the performance history deteriorates. When the cost of termination is not too high, these competing effects generate a U-shaped relationship between the agent’s continuation utility and pay-for performance intensity along with a corresponding hump-shaped relationship between monitoring and continuation utility.

A second result is that monitoring can significantly affect the relationship between performance and termination implied by the optimal contract. First, when the termination cost is high, monitoring can completely substitute for termination-based incentives. In this case, termination due to poor performance never occurs and the relationship between performance and termination completely disappears. Second, when termination happens in equilibrium, active use of monitoring can weaken the relationship between performance and termination. This is because when bad performance by the agent leads to an increase in monitoring, that performance will also lead to a correspondingly *smaller increase* in termination probability. Consequently, relative to an optimal contract without monitoring, active use of monitoring lowers the chances of termination and increases the range of performance where termination probability is relatively insensitive to past performance.

One application of our results is to managerial compensation, and our model can help explain two puzzling sets of empirical findings, as we discuss in detail in Section 5. The first empirical finding is that explicit pay-for-performance sensitivity for incumbent CEOs in financial distress has been found to be larger than for CEOs outside financial distress; However, in broad samples, pay-for-performance sensitivity appears to increase in the level of past performance. The second empirical finding is that variables such as the probability of managerial termination, firm performance, and monitoring, which intuitively should be tightly linked, are empirically not.

The U-shaped relationship between the level of past performance and future pay-for-performance sensitivity matches the first empirical result. Incumbent CEOs in financial distress usually have a history of bad performance and so have strong performance-based incentives. However, most firms are outside of distress and have managers without a history of failure – medium or high continuation values. For these managers, our model predicts that pay-for-performance sensitivity increases in the level of past performance.

In our model, the monitoring technology reduces the apparent explanatory power of performance for termination, matching the second empirical result. Monitoring lowers the probability of termination unconditionally and increases the range in which the probability

of future termination is insensitive to past performance. In addition, because monitoring can substitute for other incentives, the U-shaped relationship described above weakens any linear association between monitoring and past performance or the probability of termination.

Finally, our model can be used to study the implications of costly monitoring for optimal security design and endogenous financing constraints. As we discuss in Section 5 and Appendix C, active use of monitoring relaxes the endogenous financing constraint implied by the agency problem, allowing the creditor (the principal) to provide additional funds after negative cash flows to the borrower (the agent) without weakening incentives. This additional state-contingent financing acts as an insurance mechanism, which lowers the risk of liquidation (default). There is liquidation in equilibrium only if its costs are not too high.

This paper contributes to the literature on dynamic contracting. The studies that are most closely related to ours are DeMarzo and Fishman (2007a) and the continuous-time formulations of DeMarzo and Sannikov (2006) and Biais et al. (2007).² Our contributions relative to these models is the addition of costly monitoring technology and characterization of its impact on joint determination of compensation, monitoring, and termination as a function of observed history of performance implied by the optimal contract. In addition, we present a set of novel results on optimal dynamic incentives with performance-based incentives that may decline continuously to zero as reliance on monitoring increases.

On the technical side, our paper is related to the studies of dynamic agency problems in Brownian-Poisson settings. Sannikov (2008) provides a recent treatment of continuous-time techniques for the principal-agent problem in a pure Brownian setting, while Sannikov and Skrzypacz (2010) provide a general characterization of the set of equilibria of repeated games where players learn information via a continuous-time stationary process that has a Brownian component and a Poisson component. Other recent studies in such settings include, among others, formulations of Biais et al. (2010), He (2012), and Piskorski and Tchisty (2010).

Our paper is conceptually related to the contracting literature with costly state verification, as exemplified by static models of Townsend (1979) and Gale and Hellwig (1985). In particular, those papers, like ours, do not make *a priori* assumptions about the form

²Other recent related papers include, among others, Tchisty (2013), Clementi and Hopenhayn (2006), DeMarzo and Fishman (2007b), Sannikov (2007), Philippon and Sannikov (2007), Sannikov (2008), He (2009, 2012), Zhang (2009), Piskorski and Tchisty (2010), DeMarzo et al. (2012), Edmans et al. (2012), Edmans and Gabaix (2011), Malenko (2012), and Williams (2011).

of the contract. Unlike these papers, our model is set in dynamic, long-term environment³ with costly termination and a monitoring technology that allows for stochastic, rather than deterministic, state verification.⁴

Two important distinctions between our model and static CSV-style models are 1) our setting allows us to consider early termination of the relationship and 2) neither optimal monitoring intensity nor pay-for-performance sensitivity need to be monotonic functions of the quality of past performance. Because direct monitoring is distinct from termination or liquidation in our setting, incentive substitution means that monitoring can *reduce* the probability of performance-based termination.⁵ These two distinctions help us explain the puzzling empirical findings.

The rest of this paper is organized as follows. Section 2 introduces the model, and Section 3 derives the optimal contract. Section 4 discusses the properties of the contract, including the timing of monitoring, performance-based pay, termination, and model robustness. Section 5 discusses the relationship between the model and empirical work on managerial compensation and implications of our model for security design and endogenous financing constraints. Section 6 concludes.

All proofs are in Appendix A. Appendix B contains a discussion of the level of commitment necessary in the contract and a renegotiation-proof alternative. Finally, Appendix C contains a securities implementation in which monitoring allows for the addition of short-term financing to the standard debt contract.

³Wang (2005) extends the static analysis of Townsend (1979) to a discrete time dynamic risk sharing model with costly state verification and provides a partial characterization of the optimal contract. In contrast, our model focuses on a principal-agent problem with costly termination, and our continuous time setting allows for a tractable characterization of the optimal contract.

⁴In costly state verification models, the agency problem is that the agent may choose not to reveal the true state to the principal, and the principal can pay a fixed cost to observe the realization of the state. In our setting, cash flows are observable but the actions that generated those cash flows are not. Monitoring is simultaneous with the agent's actions and stochastic: the probability of observing the agent's shirking depends on the resources spent on monitoring as well as on the extent of shirking.

⁵The static CSV-style models do not allow to explicitly consider the early termination of the contract. However, this literature often interprets the act of monitoring as bankruptcy potentially leading to managerial replacement. In our setting, monitoring-based incentives and termination play distinct roles and can serve as substitutes.

2 Setup

In this section, we set up a continuous time contracting problem with costly monitoring between the principal (outside investors or a firm) and an agent (a manager). The agent has the ability to increase the cash flows from assets owned by the principal, but the agent does not have the funds to purchase the assets directly. As a result, the agent operates the assets in exchange for a compensation contract.

2.1 The Principal, the Agent, and the Project

Time is continuous and infinite. We assume that the principal is risk-neutral with time-discount rate r and has unlimited access to capital. The agent is also risk-neutral with time-discount rate γ . The agent has limited liability and his consumption (dC_t) must be non-negative for all t . We assume that $\gamma > r$ to capture the fact that a borrowing-constrained, undiversified agent has a higher inter-temporal marginal rate of substitution than an unconstrained principal or diversified investors.

The agent has the ability to operate a project that uses the principal's assets in a manner that adds value. The agent has no initial wealth, and the principal must incur the cost of $I > 0$ to set up the project (e.g. to purchase the assets). If the agent operates the project, a standard Brownian motion (Z) drives the cumulative cash flow process (Y) for the project. We assume that Y_t evolves according to

$$dY_t = (\mu - s_t)dt + \sigma dZ_t, \tag{1}$$

where $\mu > 0$ is the project's expected cash flow under the agent's control, $\sigma > 0$ is the project's volatility and $s_t \geq 0$ is the agent's shirking or diversion process.

While the realizations of cash flow process (Y) are publicly observable, the processes Z and s are only observable by the agent. Because of the agent's informational advantage, he can gain a private benefit (s) from taking actions that reduce the cash flow of the project.⁶

⁶There is no direct loss from diversion because the agent's private gain is the same as the cash flow loss to the project. This does not imply that diversion, shirking, or asset misuse is efficient, as we will show. Moreover, if the agent is expected to consume the entire value of the project in private benefits, he will be unable to obtain the assets from the principal necessary to start the project. Finally, it is straightforward to introduce an additional cost of diversion in our model by assuming that the agent's private benefits accumulate at the rate $\chi s_t dt$ where $\chi \in (0, 1)$, which we discuss in Section 4.3.

This private benefit can be interpreted in many ways, including 1) diversion from the project's cash flows to the agent's consumption, 2) relaxing costly effort, or 3) asset misuse in which case the agent uses a portion of the assets in a way that reduces cash flows but provides a private benefit.

The match between the principal and the agent improves their joint opportunities. Alone, the agent has an outside option with total value $V_R > 0$. The principal receives L if the contract with the agent is terminated. This reflects the value of assets in the hands of the principal and may include project liquidation costs or a cost of replacing the agent. The sum of the participants' outside options is less than the continuation value of the project, making termination costly:

$$\gamma V_R + rL < \mu. \quad (2)$$

2.2 Monitoring

In a standard dynamic principal-agent model, the compensation contract will award the agent consumption based on the performance Y of the project. We add a second incentive device, monitoring, to this setting. We assume that the principal sets a benchmark level of private benefits (diversion, shirking), s_t^b for each time t , and can then pay a cost θm_t to partially determine whether the agent engages in diversion beyond the principal's benchmark level. $m_t \geq 0$ is the principal's chosen level of monitoring, and $\theta > 0$ parameterizes the cost-effectiveness of monitoring. By choosing to monitor the agent, the principal gains access to a Poisson process N that hits with intensity $\lambda(m_t, s_t - s_t^b)$:

$$\lambda(m_t, s_t - s_t^b) = m_t \max(0, s_t - s_t^b). \quad (3)$$

We use s_t^b to represent the principal's choice of how strict is the monitoring technology; this can be distinct from the level of private benefits that the principal recommends to the agent, which we denote s_t^r . Later, we will show that the principal will optimally set $s^b = s^r = 0$, but we do not assume it.

This technological specification implies: the principal pays a positive amount to catch any diversion beyond the benchmark level ($s_t > s_t^b$); the principal never detects extra diversion above the benchmark if there isn't any; and, the probability that the principal detects deviations above the benchmark level is proportional to the intensity of monitoring and to

the amount of diversion. This monitoring technology can represent continuous or repeated auditing of the agent or the direct involvement of the principal (firm headquarters, the board of directors, or investors) in the operations of the project.

We note that our monitoring technology is different from that of the standard costly state verification models. In a typical costly state verification model, principal can pay a fixed cost to observe the realization of the state. In our setting, monitoring is simultaneous with the agent's actions and stochastic: the probability of observing the agent's diversion/shirking depends on the resources spent contemporaneously and on the extent of diversion/shirking. In addition, in our setting, the principal never learns the true realization of the state: the principal does not observe the extent of diversion, instead learning only that diversion took place at a given time.

We assume that the agent can be made to suffer more if termination occurs following discovery of bad acts (diversion, shirking, asset misuse) than following simple bad performance. In this case the principal is able to (but is not required to) reduce the agent's outside value to V_F , where by assumption $V_F < V_R$. The fact that monitoring-based termination is damaging to the agent can reflect loss of reputation, or some penal costs, as in Diamond (1984).

2.3 Formulation of the Optimal Contracting Problem

At time 0 the principal offers the agent a compensation contract in exchange for operating the project. The principal can commit to the contract terms including the choice of monitoring.⁷ The contract is constructed to deliver expected value of $V_0 \geq V_R$ to the agent. This initial value can represent the relative bargaining powers of the agent and the principal.

Under the terms of the contract, the principal offers the agent a non-decreasing consumption process C that is a function of the principal's information set: the history of cash flow realizations (Y) and detected shirking/diversion (N). Similarly, the recommended diversion process s^r , the benchmark level used in the monitoring technology s^b , and the principal's monitoring process m are also functions of the history of cash flow realizations and monitoring signals. Because the agent observes the output shocks Z separately from the output process, the manager's diversion process s can also be a function of the history output shocks. The contract continues until either the agent chooses to quit (time τ^q), the contract termi-

⁷We discuss commitment and re-negotiation proof contracts in Appendix B.

nates because of under-performance (time τ^d), or the contract terminates because the agent is caught shirking (diverting more than the recommended level) and the principal decides to terminate the relationship (time τ^f).⁸ The principal is not required to terminate the contract if the agent is caught shirking.

The principal's objective function is the expected value of cash flows from the project, including any costs of monitoring, minus consumption paid to the agent:

$$E \left[\int_0^\tau e^{-rt} (dY_t - \theta m_t dt - dC_t) + e^{-r\tau} L \right]$$

where $\tau = \min(\tau^d, \tau^f, \tau^q)$. The agent's objective function is the expected value of consumption plus private benefits:

$$E \left[\int_0^\tau e^{-\gamma t} (dC_t + s_t dt) + e^{-\gamma\tau} V_\tau \right],$$

where we write $V_\tau = V_R$ if the agent quits or the contract terminates because of under-performance ($\tau = \min\{\tau^q, \tau^d\}$) and $V_\tau \in [V_F, V_R]$, chosen by the principal, if the agent is caught shirking and the principal chooses to terminate the contract ($\tau = \tau^f$).

The contract is incentive compatible if the agent maximizes his objective function by following the principal's recommended action without quitting early. The contract is optimal if it maximizes the expected profit of the principal in the class of all contracts that are incentive compatible and provide the agent with his initial expected utility V_0 .

Consequently, we have the following definition of the optimal contracting problem.

Definition 1 *The contract $\xi = \{\tau^d, \tau^f, V_\tau, C, s^r, s^b, m\}$ specifies a (random) termination time of the relationship $\tau^d \leq \infty$, a firing time for the agent $\tau^f \leq \infty$, the agent's outside option in the event of firing (V_τ), the consumption of the manager C , the recommended diversion process s^r , the benchmark level s^b , and the monitoring choice m , all as a function of the history of cash flow realizations (Y) and detected shirking (N).*

The contract is incentive compatible if the agent maximizes his objective function by

⁸More formally, the processes C , m , s^b , and s^r are progressively measurable functions with respect to (Y, N) , while the processes s is a progressively measurable function with respect to (Z, N) . (τ^d, τ^f, τ^q) are stopping times such that $\tau^f \geq \min\{t \geq 0 : N_t = 1\}$, τ^q is (Z, N) -measurable, and τ^d and τ^f are (Y, N) -measurable. All processes fulfill the standard integrability conditions so that the stochastic integrals in Section 3 are well defined.

following the principal's recommendation ($s = s^r$) and does not quit early ($\tau^q \geq \tau^d$).

The contract is optimal if it maximizes the principal's profit subject to the incentive compatibility constraint and the constraint that the agent receives his expected utility of V_0 .

The above definition of the optimal contract implies that it maximizes the total surplus. Further, our definition of incentive compatible implies $s = s^r$. This is without loss of generality from the revelation principle: if the agent's optimal action differs from the recommendation, we can simply change the recommendation without changing the monitoring technology (in particular, without changing s^b).

In what follows we will restrict our attention to contracts in which the benchmark level used in the monitoring process is the same as the recommended level of diversion ($s^b = s^r$). In Section 4.3 we show that this restriction is without loss of generality.

3 Derivation of the Optimal Contract

3.1 Incentive Compatibility

In order to characterize the optimal contract we will consider as a state variable the value that the agent places on the remainder of the contract at any time t until termination at time $\tau = \min(\tau^d, \tau^f, \tau^q)$. For any diversion strategy and quitting time (s, τ^q) , that value is

$$V_t^s = E^s \left[\int_t^\tau e^{-\gamma(u-t)} (dC_u + s_u du) + e^{-\gamma(\tau-t)} V_\tau \middle| Z, N \right], \quad (4)$$

where the expectation is taken under the measure induced by the agent's strategy. This value includes all of the agent's expected consumption up until the contract is terminated, plus his value at termination – when contract is terminated because of under-performance, after the agent quits, or with fraud and involuntary separation.

Next, we define the change in the agent's continuation utility when he is caught diverting above the recommended level s_t^r (which is also the benchmark level s^b) as ψ_t , where $\psi_t > 0$ is a punishment/loss. Because the principal can at most punish the agent by reducing his continuation utility to V_F , we have:

$$\psi_t \leq V_t - V_F. \quad (5)$$

Then, to understand incentive compatible contracts, we must know how the continuation utility evolves if the agent follows the principal's recommendation as well as what conditions guarantee this. The following proposition provides a useful characterization of incentive compatible contracts.

Proposition 1 *For any contract, if the agent follows the principal's recommendation ($s = s^r$), then there exist predictable processes (β_t, ψ_t) so that V_t evolves before termination ($t \leq \tau$) as:*

$$dV_t^{s^r} = \gamma V_t^{s^r} dt - dC_t - s_t^r dt + \beta_t (dY_t^{s^r} - (\mu - s_t^r) dt) - \psi_t dN_t^{s^r}. \quad (6)$$

The contract is incentive compatible if and only if

$$\begin{aligned} \beta_t &\geq 1 - \psi_t m_t & \text{if } s_t^r = 0, \\ \beta_t &\in [1 - \psi_t m_t, 1] & \text{if } s_t^r > 0, \end{aligned} \quad (7)$$

and $V_t \geq V_R$ for $t \leq \tau^d$.

The equation (6) describes how the agent's continuation utility must evolve for any contract for all $t \leq \tau$ and $s = s^r$. It can be written as

$$dV_t^{s^r} = \overbrace{(\gamma V_t^{s^r} dt - dC_t - s_t^r dt)}^{\text{Promise Keeping}} + \overbrace{\beta_t (dY_t^{s^r} - (\mu - s_t^r) dt)}^{\text{Pay for Performance}} - \overbrace{\psi_t dN_t^{s^r}}^{\text{Punishment for Shirking}}.$$

The first term reflects the fact that the principal must be consistent in his promises: if the principal does not award the agent consumption today, then the agent's future consumption must grow at a rate equal to the agent's discount rate of γ . The agent's reward can come either in the form of consumption granted by the principal directly or through recommended diversion from the project's cash flows.

The second term reflects the *pay-for-performance* compensation of the agent. The agent receives a β_t fraction of project cash flows in the form of future promises. If the agent increases diversion by an additional δdt , then the project's cash flows decline and he forgoes a pay-for-performance benefit of $\delta \beta_t dt$.

The third term reflects the agent's stochastic punishment for shirking and is due to monitoring. The agent does not face any exogenous cost of diversion: if the agent diverts s_t from

the project, the agent is able to consume s_t . However, the monitoring imposes an *endogenous cost of shirking* of the agent based on the likelihood of being caught and the extent of punishment. From (3), if the agent diverts an additional δdt above the recommended level, then his probability of being caught is $\lambda(m_t, \delta) = \delta m_t dt$. If caught, the agent is punished through the loss of continuation value equal to ψ_t , making the total expected value of punishment equal to $E \left[\psi_t dN_t^{s_t^r + \delta} \right] = \psi_t \lambda(m_t, \delta) dt = \delta \psi_t m_t dt$. Thus the agent's effective efficiency of diversion, his gain from diversion minus the expected cost of punishment, is proportional to $(1 - m_t \psi_t)$.

Together, this means that if the agent diverts an additional δdt above the recommended level, then he forgoes a benefit of $\delta \beta_t dt$ and obtains an expected gain of $\delta (1 - m_t \psi_t) dt$. Hence, the agent will not divert resources above the recommended level if $\beta_t \geq 1 - \psi_t m_t$. On the other hand if the agent decides to divert δdt less than the recommended amount, he loses δdt consumption and receives an $\beta_t \delta dt$ in the form of future utility. Thus, the agent does not divert less than the recommended amount if $\beta_t \leq 1$. Combining the positive and negative deviations, we see that the agent will follow the principal's recommendation if (7) is met: $\beta_t \geq 1 - \psi_t m_t$ if $s_t^r = 0$ and $\beta_t \in [1 - \psi_t m_t, 1]$ if $s_t^r > 0$.

Finally if $V_t^{s^r} \geq V_R$ for all $t \leq \tau^d$, it is not optimal for the agent to quit early because doing so results in (weakly) lower utility.

If the agent does follow the principal's recommendation and does not quit early, (6) simplifies for all $0 \leq t \leq \tau^d$ to

$$dV_t = \gamma V_t dt - dC_t - s_t^r dt + \beta_t \sigma dZ_t. \quad (8)$$

The Poisson process that detects shirking does not report false positives, so any adjustment of the agent's utility when caught shirking is an out-of-equilibrium event.

3.2 Optimality

We note that in search for an optimal contract we can restrict our attention to contracts that implement zero diversion:

Proposition 2 *Any incentive compatible contract can be weakly improved by setting $s^r = 0$.*

This result holds because any given level of diversion can be replaced by consumption without changing incentives, making both parties weakly better off. In contrast, relaxing incentives creates unbounded diversion; thus we can focus on contracts with $s^r = 0$. However, this result requires that potential diversion is unbounded; otherwise, the principal could both award utility through finite diversion and also relax incentives. We will discuss alternative model setups, including placing exogenous bounds on the quantity and/or efficiency of diversion, in Section 4.3.

Now, we are ready to provide an intuitive derivation of the profit-maximizing way to deliver to the agent his continuation utility of V_t . Let $F(V)$ be the highest expected payoff the principal can obtain from an incentive compatible contract that provides the agent with utility equal to V . To simplify our exposition, we assume that the function F is concave and $\mathcal{C}^2$ for $V > V_R$ and at least $\mathcal{C}^1$ for $V = V_R$. We will then verify that assumption and formalize our results in Theorems 3 and 4.

Property 1 *The principal pays the agent only when $V_t \geq V_C$. The V_C is chosen so that it satisfies the smooth pasting condition $F'(V_C) = -1$ and the super-contact condition $F''(V_C) = 0$. If $V_0 > V_C$ an immediate transfer $V_0 - V_C$ is made.*

To understand the above property (as in DeMarzo and Sannikov (2006)), we observe that the principal can always make a lump sum payment to the agent of dC . This moves the agent from a continuation utility of V to $V - dC$. Let us assume that the principal awards the agent consumption when V moves above some level V_C . Then, as described in Dumas (1991), we have the smooth-pasting condition $F(V_C + dC) = F(V_C) - dC$, or $F'(V > V_C) = -1$ for any such V_C . In addition, since V_C is optimally chosen, $F''(V_C)$ must be continuous at V_C (the super-contact condition); in this case, $F''(V_C) = 0$.

Property 1 implies that for $V > V_C$, we have $F(V) = F(V_C) - (V - V_C)$. The continuity of the first and second derivatives of F at V_C ensure that if F is $\mathcal{C}^2$ on $V \in (V_R, V_C)$, then it is $\mathcal{C}^2$ for $V > V_R$. We will now analyze how the principal's value function behaves on $V \in (V_R, V_C)$.

Property 2 *The principal chooses $\beta_t = 1 - \psi_t m_t$ to minimize the volatility of the agent's continuation value. The principal chooses $\psi_t = V_t - V_F$ and fires the agent upon the first detection of excess diversion in order to minimize the cost of monitoring.*

We can understand the above property as the result of minimizing the cost of a given level of performance-based incentives. The agent's continuation utility evolves according to (8), and so Ito's lemma implies that for any $V_t \in (V_R, V_C)$:

$$dF(V_t) = \gamma V_t F'(V_t) dt + \frac{1}{2} \beta_t^2 \sigma^2 F''(V_t) dt + \beta_t \sigma F'(V_t) dZ_t.$$

Since the principal also receives cash flows of $\mu - \theta m_t$ we have a standard Hamilton-Jacobi-Bellman equation for $V \in (V_R, V_C)$:

$$rF(V) = \max_{\substack{\beta \geq 1 - \psi m \\ m \geq 0 \\ \psi \leq V - V_F}} \left[\mu - \theta m + \gamma V F'(V) + \frac{1}{2} \beta^2 \sigma^2 F''(V) \right]. \quad (9)$$

Because F is concave, β takes the minimum incentive compatible value from (7), and so $\beta_t = 1 - \psi_t m_t$. Because the monitoring process does not report false positives, the optimal contract implements maximal punishment in order to minimize the cost of monitoring. The principal fires the agent upon the first detection of and sets $\psi_t = V_t - V_F$.

Intuitively, the principal is effectively averse to risk (the value function is concave) because there is a cost of termination. To minimize the probability of termination, the principal minimizes the volatility in the agent's continuation value, subject to the incentive compatibility constraint. Then, since the agent is always induced not to shirk and the monitoring process does not report false positives, punishment is always an out of equilibrium event. Thus, punishment is always made as harsh as possible to increase the effectiveness of resources spent on monitoring.

Property 3 *The principal chooses to monitor whenever*

$$F''(V_t) < -\frac{\theta}{\sigma^2(V_t - V_F)}. \quad (10)$$

When (10) is satisfied, we have

$$m_t = \frac{\theta}{(V_t - V_F)^2 \sigma^2 F''(V_t)} + \frac{1}{(V_t - V_F)} \quad (11)$$

$$\beta_t = 1 - \psi_t m_t = -\frac{\theta}{(V_t - V_F) \sigma^2 F''(V_t)} \quad (12)$$

Plugging $m_t = \frac{1}{\psi_t}(1 - \beta_t)$ and $\psi_t = V_t - V_F$ from Property 2 back into the Hamilton-Jacobi-Bellman equation (9) yields a refined optimization for $V \in (V_R, V_C)$:

$$rF(V) = \max_{\beta \in [0,1]} \left[\mu - \frac{\theta}{V - V_F}(1 - \beta) + \gamma VF'(V) + \frac{1}{2}\beta^2\sigma^2 F''(V) \right]. \quad (13)$$

The concavity of F implies that we can use the first-order conditions and possible corner solutions to find the optimal value of β . Taking the first order condition yields (11) and (12). Solving for $\beta_t = 1$ or $m_t = 0$, we find that the principal chooses to monitor when (10) is satisfied.

Monitoring allows the principal to relax pay-for-performance incentives. The cost of monitoring is direct (θm_t), while the cost of pay-for-performance incentives is that the principal may induce (inefficient) termination. When the cost implied by inefficient termination is high ($F''(V_t)$ is sufficiently negative), the principal makes monitoring more intense.

The principal can completely replace pay-for-performance incentives with monitoring: set $\beta_t = 0$ and $m_t = \frac{1}{V_t - V_F}$. In principle, one might think that a contract with full monitoring everywhere ($\beta_t = 0$ for all histories) might be optimal for some parameter values but Properties 1 and 3 combine to show otherwise. At any point at which the agent consumes, the principal's value function must have $F''(V_t) = 0$, which implies zero monitoring at that point.

The ability to completely replace pay-for-performance incentives with monitoring implies that the agent's continuation utility (8) may have zero volatility and that the principal's HJB equation (13) may have singularities. In particular, the optimal value of β (from 12) has the property that $\beta \downarrow 0$ is coincident with $F'' \downarrow -\infty$. To analyze this possibility, we will first consider a restricted problem that does not have a singularity before returning to the general problem:

Definition 2 (Restricted Problem) *Consider the solution to our baseline model under the restriction that the principal must choose $\beta_t \in \{0 \cup [\epsilon, 1]\}$ for some small $\epsilon > 0$. We denote the principal's value function $F_\epsilon(V)$ and use the ϵ subscript to indicate quantities associated with this optimization (β_ϵ , $V_{\epsilon,C}$, etc.).*

$$rF_\epsilon(V) = \max \left\{ \begin{array}{l} \mu - \frac{\theta}{V - V_F} + \gamma VF'_\epsilon(V), \\ \max_{\beta_\epsilon \in [\epsilon, 1]} \left[\mu - \frac{\theta}{V - V_F}(1 - \beta_\epsilon) + \gamma VF'_\epsilon(V) + \frac{1}{2}\beta_\epsilon^2\sigma^2 F''_\epsilon(V) \right] \end{array} \right\} \quad (14)$$

The purpose of the restricted problem is to remove possible singularities by separating regions with $\beta = 0$ from those with $\beta > 0$.⁹ A potential singularity is created when we allow β to *approach* zero; if β is equal to zero, the volatility term does not exist, and there is a well-behaved first-order ODE. This ODE is the top component of (14). Similarly, if β is positive but kept away from zero, the unrestricted HJB equation has positive and bounded coefficients, so it too is well-behaved. This is the bottom component of (14). In sum, our restricted problem creates a gap between the lowest positive value of β and zero, separating the contracting solution into two ODEs that are each well-behaved.

It is important that, for on any policy choice that is admissible in the restricted problem, the HJB equations for the unrestricted and restricted problems (13 and 14) are the same. Thus Properties 1, 2, and 3 (for interior values of β_ϵ) are all inherited by the restricted problem, with the additional restriction that

$$\beta_{\epsilon,t} = \epsilon \quad \text{iff} \quad F''(V_t) \leq -\frac{\theta}{(V_t - V_F)\sigma^2\epsilon} \quad (15)$$

Now, we can locate areas where β_ϵ can potentially be zero:

Property 4 *For all small $\epsilon > 0$, the principal chooses $\beta_\epsilon(V) > 0$ for all $V \in (V_R, V_{\epsilon,C}]$. $\beta_\epsilon(V)$ may be zero only at $V = V_R$.*

From Property 3, $\beta_\epsilon(V_{\epsilon,C}) = 1$; thus to have $\beta_\epsilon = 0$, it must be that there is a level of the agent's continuation value such that it is optimal for β_ϵ to jump from $\beta_\epsilon \in [\epsilon, 1]$ to $\beta_\epsilon = 0$. There are several necessary conditions for such a jump inside $V \in (V_R, V_{\epsilon,C})$ (i.e. on the interior): the value-matching, smooth-pasting, and super-contact conditions are that F_ϵ , F'_ϵ , and F''_ϵ are all continuous across the jump. Importantly, value-matching and smooth pasting are implied by the existence of a jump, but super-contact is a first-order condition for the optimality of the jump location and is only necessary on the interior of V .¹⁰ This continuity implies that both the left and right side values of β_ϵ must satisfy the same maximization (14) at the location of the jump. We use algebra and bounding arguments to show in the appendix that this super-contact condition implies such jumps are impossible except at $V = V_R$.

⁹Economically, the restricted problem for small ϵ is essentially the same as the unrestricted problem: if $\epsilon = .001$, then we are assuming that the agent cannot receive less than a tenth of a penny per dollar in performance pay unless he receives zero performance pay. We allow arbitrarily small ϵ .

¹⁰See e.g. Dumas (1991) for a general discussion of these conditions or Zhu (2013, Theorem 2) for another application in a similar setting.

Intuitively, the principal can always improve a solution with full monitoring at some $\hat{V}$ ($\beta_\epsilon(\hat{V}) = 0$) by shifting the point of full monitoring to some smaller $\hat{V}$ – every solution satisfying (14) with a jump to $\beta_\epsilon = 0$ is dominated by another solution with a jump at lower $\hat{V}$.¹¹ The only remaining possibility is that the principal may employ full monitoring at $V = V_R$.

Property 5 *There exists an L_ϵ^* such that the principal will choose $\beta_\epsilon(V_R) = 0$ and $m_\epsilon(V_R) = \frac{1}{V_R - V_F}$ if and only if $L \leq L_\epsilon^*$. In addition, $L_\epsilon^* \geq \frac{1}{r} \left(\mu - \frac{\theta}{V_R - V_F} - \gamma V_R \right)$.*

This property can be understood by contradiction. The principal can always offer the agent a package of monitoring ($m_t = \frac{1}{\psi_t}$) and compensation ($dC_t = \gamma V_t$) so that the agent's continuation value is constant and termination is impossible. This gives the principal a value of $\frac{1}{r} \left(\mu - \frac{\theta}{V_t - V_F} - \gamma V_t \right)$. If the principal's value of termination is lower than this benchmark, then the principal never allows termination. The true value of L^* is higher because the principal can (and will) implement full monitoring only temporarily, delaying the agent's consumption to push the continuation value up before monitoring is relaxed.

Having solved for β_ϵ , we can fully described the behavior of $F_\epsilon(V)$ on the set $V \in [V_R, V_{\epsilon,C})$ using the HJB equation (14). The smooth pasting and super-contact conditions (Property 1) can be substituted back into (14) to yield $rF_\epsilon(V_{\epsilon,C}) = \mu - \gamma V_{\epsilon,C}$. The left boundary is $F_\epsilon(V_R) = \max\{L, L_\epsilon^*\}$.

We formalize these findings in the following theorem:

Theorem 3 *Define L_ϵ^* to be the maximum value of $F_\epsilon(V_R)$ among all solutions to (14) that have some $V_{\epsilon,C}$ which satisfies Property 1 and $\beta_\epsilon(V_R) = 0$. Then, there is a unique, concave function that satisfies (14) on $[V_R, V_{\epsilon,C}]$, is $\mathcal{C}^2$ on $(V_R, V_{\epsilon,C}]$ and at least $\mathcal{C}^1$ at V_R , with $F'(V_{\epsilon,C}) = -1$, $F''(V_{\epsilon,C}) = 0$, and $F_\epsilon(V_R) = \max\{L, L_\epsilon^*\}$. That function characterizes the optimal contract in the ϵ -restricted problem for small ϵ .*

¹¹Solutions to (14) with partial monitoring have a more negative second derivative than solutions with full monitoring. If partial and full monitoring solutions have a proposed crossing, they must cross with equal level and slope (e.g. Dumas (1991)). Then, there is always a slightly higher partial-monitoring solution that crosses the full-monitoring solution twice, once from below to the left of the proposed crossing, and once from above to the right of the proposed crossing. The crossing from below has a higher slope, and therefore higher values to the right of the crossing. This means that we can always improve a full-monitoring solution by shifting to the point where the principal engages full monitoring to lower V . Zhu (2013) uses a similar argument.

The agent's continuation utility evolves as (8) with $\beta_\epsilon(V_t)$ given by (12), and (8) has a unique (weak) solution. The agent is awarded consumption at $V > V_{\epsilon,C}$.

If $L > L_\epsilon^*$, then the contract terminates at $V_t = V_R$. If $L \leq L_\epsilon^*$, then the contract never terminates and V_t moves upward at V_R .

Solutions to the restricted problem follow a standard form: the final ODE is a composition of two ODEs with bounded or constant coefficients, the top and bottom lines of (14). The transition between the two ODEs is at a fixed location, V_R .

Next, we want to use the solutions to the restricted problem ($\beta_\epsilon \in \{0 \cup [\epsilon, 1]\}$) in order to solve the unrestricted problem ($\beta \in [0, 1]$). The monotone convergence theorem thus tells us that $\lim_{\epsilon \rightarrow 0} F_\epsilon$ exists in a point-wise sense: The principal is strictly better off as $\epsilon \downarrow 0$, which means that F_ϵ is non-decreasing in ϵ and bounded from above by the first-best. However, point-wise convergence does not necessarily preserve differentiability or policy choice. Thus, we need in addition to show that we can use the ϵ -restricted solutions to exhibit a unique solution to the unrestricted problem (13).

We will first show the result and then illustrate the construction we use to control the unrestricted ODE:

Theorem 4 Define $L^* = \lim_{\epsilon \rightarrow 0} L_\epsilon^*$, which is also the infimum among all solutions to (13) with $\beta > 0$ and valid right boundary (Property 1). Then there is a unique, concave function on $[V_R, V_C]$ that satisfies (13), is $\mathcal{C}^2$ on $V \in (V_R, V_C]$ and at least $\mathcal{C}^1$ at V_R , and has $F'(V_C) = -1$, $F''(V_C) = 0$, and $F(V_R) = \max\{L, L^*\}$. That function characterizes the optimal contract in the unrestricted problem.

The agent's continuation utility evolves as (8) with $\beta(V_t)$ given by (12), and (8) has a (weak) solution. The agent is awarded consumption at $V > V_C$.¹²

If $L > L^*$, then the contract terminates at $V = V_R$. If $L \leq L^*$, then the contract never terminates and V_t moves upward at V_R .

The key feature of the solution is that principal can engage in full monitoring ($\beta(V) = 0$ for some V), but this event is coincident with a singularity ($F''(V) = -\infty$, from (12)).

¹²We would like to show here that the process (8) has a unique (weak) solution. This can be proven when $\beta(V)$ drops discretely to zero at V_R (see e.g. Harrison and Lemoine (1981) or Engelbert and Peskir (2014)), but we have been unable to find a proof for all cases in which $\beta(V)$ goes continuously to zero. If $\beta_\epsilon \rightarrow \beta$ uniformly locally, then Theorem 4.8 in Jacod and Shiryaev (2003) shows that weak solutions to (8) for the restricted problem converge to a weak solution for the unrestricted problem as $\epsilon \rightarrow 0$.

In contrast, when it is not optimal to engage in full monitoring ($\beta(V) > 0$ for all V), then the constraint on β_ϵ in the constrained problem is slack for some ϵ small, and the constrained and unconstrained problems are equivalent. Thus, we can exhibit a solution to the unrestricted contracting problem if we can solve the issue of the singularity. The key property of the unrestricted problem is that, as with the restricted problem, one can only have full monitoring at $V = V_R$:

Property 6 *In the unrestricted problem, the principal chooses $\beta(V) > 0$ for all $V \in (V_R, V_C]$. $\beta(V)$ may be zero only at $V = V_R$, and $\beta(V_R) = 0$ if and only if $F(V_R) \leq L^*$.*

Consider solutions to the HJB equation (13) that have a valid consumption boundary (right boundary, V_C) and that exhibit no termination. These solutions all have $\beta = 0$ somewhere, and thus V_t reflects upward at that point (from the law of motion, equation 8). For these solutions to the HJB equation, the particular value of L is irrelevant because termination never occurs. However, these solutions can be ordered, which we prove in the appendix (Lemma 8 and the proof of Theorem 4) and illustrate with dotted black lines in Figure 1.¹³

If we take an existing solution to the HJB equation with a valid consumption boundary (right boundary, V_C), we can find another solution with the consumption boundary shifted slightly to the left. This second solution will have higher values of $F(V)$ and $F''(V)$ everywhere both solutions exist, and therefore $\beta(V)$ will be higher everywhere as well (unless $\beta(V)$ is at the upper bound of 1). In addition, the value of V for which $\beta(V) = 0$ will shift to the left. This shift means that there is a single best no-termination solution, and it is the solution that has the lowest value of V_C and the highest $\beta(V)$ – thus, the best no-termination solution has $\beta(V > V_R) > 0$. We illustrate this solution as the solid black line in Figure 1.

This best no-termination solution generates $F(V_R) = L^*$, and it is the solution to the contracting problem for all $L \leq L^*$. If we continue to move V_C to the left, we find solutions

¹³Numerical solutions are generated using the discretization method in Kushner and Dupuis (2001) and implemented in MATLAB. We use a grid on $V \geq V_R$ with interval $dV = \frac{1}{200}$ and a grid on $\beta \in [0, 1]$ with interval $d\beta = \frac{1}{200}$. We allow the left boundary ($F(V_R)$) to exceed L for $\beta(V_R) = 0$, and the right boundary (V_C) is set by forcing $F'(V_C) = -1$ if $rF(V_C) \geq \mu - \gamma V_C$. We iterate the solution until the sum of the absolute value of all changes of $F(V)$ is less than 10^{-8} , or 25 million iterations, whichever comes first. In all cases, the absolute value of the discretized HJB equation at all points is less than 10^{-6} . When simulating the probability of termination in Section 4.2, we use the discretization $V_{t+dt} = V_t + \gamma V_t dt + \beta(V_t) \sigma \sqrt{dt} N(0, 1)$ with $dt = \frac{1}{100}$; then $V_{t+h} = \min\{V_{t+h}, V_C\}$. We count the fraction of trials that have $V_{t \leq T} \leq V_R$ for $T = 10$ or 5.




Figure 1: Solutions to the HJB equation as we change the proposed $\hat{V}_C$. This figure shows different solutions to the HJB equation (13) as we change $\hat{V}_C$, ignoring the left boundary condition at V_R . The dashed grey line plots $rF(V_C) = \mu - \gamma V_C$. The solid black line plots the solution with $\beta(V_R) = 0$ and it represents the solution to the unrestricted contracting problem when $L \leq L^*$. Solid black circles denote singularities where $\beta(V) = 0$ and $F''(V) = -\infty$. The dotted black lines plot solutions to (13) that contain singularities. The dashed black lines plot solutions to (13) with $\beta(V) > 0$, and they represent solutions to the unrestricted contracting problem when $L > L^*$. The curves are plotted using $r = .01$, $\gamma = .05$, $\sigma = 1$, $V_R = 1$, $V_F = 0$, $\theta = 1$, and $\mu = .5$.

with $F(V_R) > L^*$ that have $\beta > 0$, and so they are solutions for $L > L^*$ to both the unrestricted problem and the restricted problem for some ϵ . We illustrate these solutions as dashed black lines in Figure 1.

It will be optimal for the principal to fund the project if the expected proceeds are higher than the cost and if the optimal contract delivers the agent an expected lifetime utility of at least V_0 . This happens whenever $\max_{V \geq V_0} F(V) \geq I$.

4 Properties of the Optimal Contract

In this section we discuss the properties of the optimal contract. We focus on joint determination of monitoring, termination, and pay-for-performance incentives, as well as the efficiency gains from monitoring. We recall from Section 3.1 that monitoring-based incentives and pay-for-performance incentives must add up to the efficiency of diversion, which is 1. Thus, we will focus our analysis on $\beta(V)$ – the pay-for-performance incentives that are easier for empiricists to measure, remembering that $1 - \beta(V)$ is the level of monitoring-based incentives. Pay-for-performance incentives, monitoring-based incentives, and the direct cost of monitoring (m) are related through Property 3. We illustrate our discussion with a number of numerical examples.

4.1 The Timing of Monitoring and Pay-for-Performance Incentives

The optimal contract efficiently trades off monitoring and pay-for performance incentives to maximize the total surplus from the project. For interior (positive) levels of monitoring, the level of pay-for-performance incentives,

$$\beta_t = 1 - \psi_t m_t = \frac{\theta}{V_t - V_F} \times -\frac{1}{\sigma^2 F''(V_t)}, \quad (16)$$

is the product of the cost effectiveness of monitoring (the inverse of $\frac{\theta}{V_t - V_F}$) and the principal's aversion to volatility ($-\frac{1}{\sigma^2 F''(V_t)}$).

The optimal timing and intensity of monitoring is shaped by two competing forces: the *cost of termination* and the *inside stake*. First, the principal is effectively risk averse – his value function $F(V)$ is concave – because the volatility in the agent's continuation utility from pay-for-performance incentives implies that a sequence of bad luck will result in $V_t = V_R$ and contract termination. Since termination is costly to the principal, the principal is averse to volatility/risk. The principal can reduce the agent's performance-based incentives and thus reduce the risk of inefficient termination by substituting monitoring, but he must pay a cost to do so.

Second, the cost effectiveness of monitoring depends on how much the principal can punish the agent when the agent is caught diverting ($V_t - V_F$). When the agent has had a

more favorable history (higher V_t , so higher $\psi_t = V_t - V_F$), the agent's inside stake is higher, and the principal can take more away from the agent if he is caught shirking. This makes the monitoring technology more effective: monitoring-based incentives ($1 - \beta_t = m_t \psi_t$) can be created at a lower cost θm_t .

To proceed, we will establish and then explain several comparative statics on the intensity of monitoring:

Proposition 5 *The function $\beta(V)$ has the following properties:*

- (i) *As L increases, $\beta(V)$ increases (weakly) for all V .*
- (ii) *The boundary L^* , with termination if and only if $L > L^*$, is (weakly) decreasing in V_F and θ .*
- (iii) *If $\beta'(V) = 0$ and $\beta \in [0, 1)$, then $\beta''(V) > 0$. As one consequence, solutions without termination have $\beta'(V) > 0$ for V with $\beta(V) < 1$.*
- (iv) *Non-monotonic $\beta(V)$: If $V_R - V_F$ is small enough¹⁴, then there exists constants $L_a > L_b > L^*$ such that*
 - 1) $L_b > L \geq L^*$ *implies* $\beta'(V_R) > 0$
 - 2) $L_b = L$ *implies* $\beta'(V_R) = 0$
 - 3) $L_a > L > L_b$ *implies* $\beta'(V_R) < 0$
 - 4) $L \geq L_a$ *implies* $\beta(V_R) = 1$.

Because $\beta(V_C) = 1$, Proposition 5(iii) and $\beta'(V_R) > 0$ are sufficient to show that $\beta(V)$ is non-decreasing, while $\beta'(V_R) < 0$ is sufficient to show that $\beta(V)$ is non-monotonic. In addition, part (iv) can be extended to any $\hat{V} > V_R$ by looking at a new economy with $\{V_R, L\} = \{\hat{V}, F(\hat{V})\}$ (i.e. just looking at the high V portion of the contract, which is the full optimal contract for some economy with a larger value of V_R).

The first comparative static is that for any given continuation value, the optimal incentive effects of monitoring are lower (β is higher, pay for performance is higher) in relationships with higher liquidation values. So, when the principal has less to lose in termination, she provides fewer monitoring-based incentives. We have already shown an extreme version of this comparative static in Theorems 3 and 4: when $L < L^*$, we have $\beta(V_R) = 0$ and all

¹⁴The precise condition is given in the appendix (equation A-16).

incentives at V_R are monitoring-based incentives. What is new here is the general result that monitoring-based incentives ($1 - \beta(V)$) decline everywhere as L increases. The pattern is illustrated in Figure 2.




Figure 2: Comparative Statics with respect to L . This figure plots the principal's value function (F), monitoring policy (m), and performance-based incentives for the agent (β), all as a function of the agent's continuation value (V). The dashed black line displays the no-termination solution for which $L^* = 15.1 > 14 = L$, while the dashed and dotted black lines display cases in which there is termination in equilibrium ($L > L^* = 15.1$). The dashed grey line in the first panel plots the right boundary condition $F(V_C) = \frac{\mu}{r} - \frac{\gamma}{r} V_C$. The curves are plotted using $r = .01$, $\gamma = .05$, $\sigma = 1$, $V_R = 1$, $V_F = 0$, $\theta = 1$, and $\mu = .5$.

Further, Proposition 5(ii) shows that the cut-off value of L^* , with termination if and only if $L > L^*$, is (weakly) decreasing in V_F and θ . Intuitively, higher V_F and θ imply that the monitoring technology is less effective due to weaker punishment and higher monitoring costs. This change means monitoring-based incentives are more expensive: to obtain a similar level of $1 - \beta_t$, the cost of monitoring (θm_t) is higher. Thus, the range of liquidation values for which termination happens is larger.

Proposition 5(iii) shows that the $\beta(V)$ functions in Figure 2 are somewhat generic: over

any given interval of V with strictly positive monitoring, the only allowable shapes for $\beta(V)$ are U-shaped and strictly increasing. This result also implies that if a solution has full monitoring ($\beta(V_R) = 0$), then $\beta(V)$ is strictly increasing in V until $\beta(V) = 1$. From Proposition 5(ii), we have full monitoring solutions with $\beta'(V) > 0$ when L is sufficiently low, or V_F and θ are sufficiently high.

Finally, Proposition 5(iv) allows us to categorize the behavior pay-for-performance incentives (or, one minus monitoring-based incentives) as a function of history. In particular, $\beta(V)$ is not monotonic: if V_F is high enough, meaning punishment is sufficiently ineffective, and L is high, meaning the loss to the principal in termination is low, then better performance leads to *more* monitoring-based incentives ($\beta'(V_R) < 0$); if instead punishment is sufficiently ineffective and the principal has a large loss in termination, better performance leads to *less* monitoring-based incentives. Figures 2, 4, and 5 all illustrate solutions with non-monotonic $\beta(V)$.

To understand the non-monotonicity of $\beta(V)$, let us first consider how we obtain $\beta'(V) > 0$. In relationships in which the agent's inside stake is always large (V_F is low, so $\psi = V - V_F$ is high), and/or the principal's loss in default is large, there is no termination. In these solutions, bad performance (and the resulting new, lower, inside stake) increases the principal's use of monitoring-based incentives. If the cost-effectiveness of monitoring is always high (because V_F is very low and a large punishment is always possible), the cost-effectiveness of monitoring is also relatively unchanging in V , so the "cost of termination" effect dominates and only $\beta'(V) > 0$ is possible. Similarly, if the agent's inside stake is not high, $\beta'(V) > 0$ still holds if the principal's liquidation value is very low – the cost of termination is so high that "cost of termination" effect always dominates. In either case, the principal prefers to pay for full monitoring to avoid a large loss in default.

To continue, (16) shows that $\beta'(V) > 0$ can only occur if the principal's effective risk aversion ($-F''(V)$) is higher as the distance to default becomes lower.¹⁵ Moreover, the change in the principal's effective risk aversion must be larger than the decrease in the cost-

¹⁵Alternately, if L is low enough, $F''(V)$ is monotonic. It is not true that F'' is always monotonic in V . To see this heuristically, consider the model when monitoring is impossible ($\theta = \infty$) and $V_R = 0$. In that case, the HJB (13) shows that $F''(V_R) = 2\frac{rL-\mu}{\sigma^2}$. So, for very large values of σ , the principal is effectively risk neutral near default. This is a common property of diffusions near fixed boundaries. However, our model features a volatility control, and so the principal's concavity is partially determined by the choice of monitoring intensity. Our model shows that when the agent's inside stake is high (or, the principal's liquidation value is low) the combined effect is to make the equilibrium outcome such that $-F''(V)$ is increasing as V declines.

effectiveness of monitoring from the declining inside stake. Proposition 5 shows that this intuition generalizes: all contracts share the feature that if the agent has a high enough inside stake ($V - V_F$ is high enough), increasing the inside stake induces the principal to rely more on pay-for-performance incentives and to monitor less.

However, if the principal's loss in liquidation is moderate and the cost-effectiveness of monitoring is not too high ($V - V_F$ is low), then if V declines too much, the cost of monitoring-based incentives ($\frac{\theta}{V - V_F}$) can go up enough that we see $\beta'(V) < 0$. In these solutions $\beta(V)$ is U-shaped, rather than strictly decreasing. Bad performance causes a decline in the agent's inside stake with a corresponding decline in the effectiveness of monitoring, and this effect dominates the change in the principal's effective risk aversion. The result is that for these relationships, the principal monitors less near default.

4.2 The Probability of Termination and Some Numerical Examples

Figure 3 shows the intensity of monitoring along with the probability of termination within 5 years and 10 years, all as a function of the agent's continuation value. We observe that as monitoring becomes more cost effective (θ declines), the probability of termination also declines. When monitoring is impossible ($\theta = \infty$), the probability of termination is high even for agents with the highest possible histories of success (8% over five years; 25% over ten years at $V = V_C$ and $\frac{V_C - V_R}{\sigma} = 4$). When monitoring is cost effective, however, the probability to termination within one or two years for successful agents drops close to 0%. Thus the magnitude of the effect of even some monitoring on the probability of termination can be very high.

We also observe that the probability of termination is convex in continuation value.¹⁶ Our finding, illustrated in Figure 3, is that as monitoring becomes easier (θ declines) the convexity of the probability of termination increases. Thus, agents with histories of success – high continuation values – have very little chance of termination, and the probability of termination is insensitive to small changes in the continuation value. In contrast, agents

¹⁶This property is general and applies as well to the case when monitoring is impossible or unused. If $P(V, T) = E[1_{\tau^q \leq T} | V_0 = V]$ is the probability that the agent will be terminated before T additional time has passed given a continuation value V_0 , and $L > L^*$, then we have boundary conditions $\frac{\partial}{\partial V} P(V = V_C, T) = 0$, $P(V = V_R, T) = 1$, and $P(V > V_R, T = 0) = 0$. From Ito's lemma, the ODE that describes P is $0 = -P_T + \gamma V P_V + \frac{1}{2} \beta(V)^2 \sigma^2 P_{VV}$. By inspection, $P_T > 0$ and $P_V < 0$, so we must have $P_{VV} > 0$.




Figure 3: The Probability of Termination. This figure plots the probability of termination $V_t = V_R$ within one year, within two years, and pay-for-performance incentives ($\beta(V)$), all as a function of the agent's continuation value (V). The solid black line displays the case in which monitoring is impossible ($\theta = \infty$), while the dashed and dotted black lines display cases in which monitoring is harder or easier. All probabilities are strictly positive. The curves are plotted using $r = .01$, $\gamma = .05$, $\sigma = 2$, $V_R = 1$, $V_F = 0$, $\mu = .5$, and $L = 10$.

with histories of failure – low continuation values – have a higher chance of termination, and the probability of termination is very sensitive to changes in the continuation value.

The transition from a region where the probability of termination responds relatively little to performance to a region where the response is much more dramatic is a result of the monitoring technology. When the agent has a history of success, moderate failure causes the principal to increase monitoring and reduce performance-based incentives. Since termination is a form of performance-based incentive, moderate failure after strong success cannot result in a high probability of termination. On the other hand, moving from moderate failure to large failure moves the principal past monitoring and back to performance-based incentives. This in turn implies that termination probability will substantially increase in this range. Overall, as Figure 3 illustrates, relative to an optimal contract without monitoring, active

use of monitoring lowers the chances of termination and increases the range of performance where termination probability is very insensitive to past performance.

Next, we examine some numerical comparative statics for θ and V_F . Given that decreasing L increases the incentives provided through monitoring for all V , one might think that θ and V_F would generate a similar comparative static. If θ or V_F is decreased, monitoring-based incentives are cheaper for any given continuation value, and so might monitoring be more intense everywhere? Increasing the cost-effectiveness of monitoring does *not* result in more monitoring for V , as illustrated in Figures 4 and 5. For both decreasing θ and V_F , there is an overall upward shift in the shape of the monitoring function, but there is also a leftward shift. Monitoring is used more intensely, but it is used more intensely nearer to termination. In other words, the principal responds to the more effective punishment by *delaying* monitoring until closer to termination. This makes monitoring more intense in those states, but less expensive overall because those states are less likely to be realized.

While we have not been able to prove the result analytically, simulations plotted in Figure 3 illustrate that the endogenous probability of termination is increasing in θ for a given V . This result holds true for all of the parameter values that we tried.

4.3 Model Robustness

In deriving the optimal contract, we made four assumptions that we might like to relax. The first is that we've assumed $s^r = s^b$, meaning that the monitoring technology's detection cutoff (3) is also the recommended level of diversion. The second is that possible diversion is unbounded. The third is that diversion has no exogenous efficiency loss. The fourth is $V_F < V_R$ rather than $V_F \leq V_R$.

In principle, the principal may want to set the technological input for the monitoring technology s^b to be different from the recommended diversion level s^r . However, any $s^b \neq s^r$ is sub-optimal. Consider first $s^b > s^r$: with this choice, monitoring is useless in detecting the agent's deviations around s^r , which means it is both costly and useless in deterring any local deviation.

If, instead, $s^b < s^r$, the principal uses a monitoring technology that detects shirking/diversion even when the agent follows the principal's recommended action. Worse, the principal cannot tell if the detection was driven by recommended or non-recommended diversion. Moreover, while setting $s^b < s^r$ makes detection technology less accurate, the linearity




Figure 4: Comparative statics with respect to θ . This figure plots the principal's value function (F), monitoring policy (m), and performance-based incentives for the agent (β), all as a function of the agent's continuation value (V). The solid black line displays the case in which monitoring is impossible ($\theta = \infty$), while the dashed and dotted black lines display cases in which monitoring is more or less costly. The dashed grey line in the first panel plots the right boundary condition $F(V_C) = \frac{\mu}{r} - \frac{2}{r}V_C$. The curves are plotted using $r = .01$, $\gamma = .05$, $\sigma = 1$, $V_R = 1$, $V_F = 0$, $\mu = .5$, and $L = 10$.

of the monitoring and diversion technologies implies that the incentive compatibility conditions will be the same as for $s^r = s^b$: the incremental loss from detection to the agent due to additional diversion will be the same in both cases. Thus, by accepting false positives, the principal adds noise to the agent's continuation value. To keep the agent's continuation utility a martingale, V_t must drift up at the rate $\psi_t m_t (s^r - s^b) dt$ to compensate the agent for the down-jumps associated with a false-positive detection event. (8) becomes

$$dV_t = \gamma V_t - dC_t - s^r dt + \beta_t \sigma dZ_t + \psi_t (m_t (s^r - s^b) dt - dN_t^{s^r}) \quad (17)$$

where the final term is a jump martingale. Because the principal is effectively risk averse – his value function is concave because of the loss from termination – the additional noise is




Figure 5: Comparative statics with respect to V_F . This figure plots the principal's value function (F), monitoring policy (m), and performance-based incentives for the agent (β), all as a function of the agent's continuation value (V). The solid black line displays the case in which the agent cannot be punished beyond simple termination ($V_F = V_R$), while the dashed and dotted black lines display cases in which punishment is more effective. The dashed grey line in the first panel plots the right boundary condition $F(V_C) = \frac{\mu}{r} - \frac{2}{r}V_C$. The curves are plotted using $r = .01$, $\gamma = .05$, $\sigma = 1$, $V_R = 1$, $\theta = 1$, $\mu = .5$, and $L = 10$.

undesirable by Jensen's inequality. Together, these arguments imply that setting $s^b = s^r$ is optimal, which we formalize in the following Lemma:

Lemma 6 *It is optimal to set $s^r = s^b$.*

We can also extend our model to one with bounded diversion and an exogenous efficiency loss. We can assume that $s_t \in [0, \bar{s}]$ and that diversion generates consumption equal to χs_t where $0 < \chi \leq 1$. This setup corresponds to a hidden effort/investment model where the agent can choose effort to generate cash flows equal to $(\mu - s_t)dt$, with a private benefit from diversion equal to χs_t . Our baseline model corresponds to the case in which $\bar{s} = \infty$ and $\chi = 1$. Results change if $\bar{s} < \infty$ or $\chi < 1$ and this case has been studied by DeMarzo and

Sannikov (2006, Proposition 8) in the setting without monitoring. Their result is that the principal will never optimally induce diversion if and only if

$$\frac{\mu - \bar{s}}{r} \leq \min_{V \geq V_R} F(V) + \frac{\gamma}{r} \left(\frac{\chi \bar{s}}{\gamma} - V \right) F'(V) \quad (18)$$

where F is the principal's value function. This condition is derived by comparing the value received without diversion ($F(V)$; $s^r = 0$) to the value received with full diversion (setting $s_t^b = s_t^r = \bar{s}$). However, the value from full diversion does not depend on the monitoring technology because monitoring is not used at any such full diversion time. Thus, the DeMarzo and Sannikov (2006) condition also applies in our model, except that the value function is the solution to our HJB equation (13) with monitoring and is thus higher than the value function in the model without monitoring.¹⁷

To analyze the case in which $V_F = V_R$, we will consider the limit of economies as $V_F \uparrow V_R$. As V_F increases, the principal is made (weakly) worse off because his set of feasible policies is made smaller: the principal can punish the agent after a detection event by reducing the agent's continuation value to $V \geq V_F$. At the same time, the principal's total utility is bounded from below by the economy with no monitoring. Thus, by the monotone convergence theorem, we have $F(V; V_F = V_R)$ exists and $F(V; V_F) \downarrow F(V; V_F = V_R)$ pointwise.¹⁸

¹⁷Equation (18) implies that if $\chi < 1$ and $\bar{s}$ exceeds some upper bound, diversion is never optimal. There is no such upper bound if $\chi = 1$. In one sense, this result might seem trivial: since diversion is a fully efficient way of granting consumption, it might be that one can simply replace consumption grants with recommended diversion. However, the continuation value at which the principal implements diversion, call it V_S , is lower than V_C in the no-diversion model because the principal is able to grant consumption while simultaneously relaxing incentives – an improvement. Instead of finding a unique V_C consistent with Property 1, one finds a unique V_S such that $0 = -rF(V) + \mu - \bar{s} + (\gamma V - \bar{s})F'(V)$ for $V \geq V_S$, the HJB equation (13) is satisfied for $V \leq V_S$, and where the value-matching, smooth-pasting, and super-contact conditions are met at V_S . Then V_t reflects downward at V_S , and the remainder of the contract form is unchanged. In particular, recommended diversion is not triggered by a decline in V when $\chi = 1$. See Figure 6 in DeMarzo and Sannikov (2006) of Figure 4 in Zhu (2013) for a graphical illustration.

¹⁸We observe here that our existing results apply when $V_R - V_F$ is made arbitrarily small, and so the economic distinction between $V_R = V_F$ and $V_F < V_R$ is likewise small.

5 Discussion

In this section, we explore several implications of our model. The first set of results helps resolve a puzzling empirical finding in the performance/pay-for-performance relationship. The second set of results helps explain the weak empirical performance/termination relationship. The third set of results can be used to study costly monitoring in the context of optimal security design and endogenous financing constraints.

The Performance / Pay-for-Performance Relation

The empirical literature has documented an apparently conflicting relationship between explicit pay-for-performance incentives given to managers and measures of a firm’s past performance. First, explicit pay-for-performance sensitivity for incumbent CEOs in financial distress has been found to be *larger* than for CEOs outside financial distress. In contrast, in broad samples, basic indications are that pay-for-performance sensitivity *increases* in the level of past performance.¹⁹

Our model helps resolve the apparent inconsistency by producing a U-shaped relationship between the level of past performance and future pay-for-performance sensitivity (Section 4.1). Since firms in distress are usually led by incumbent managers with a recent history of bad performance, we can start by examining the predictions of our model for managers with low continuation values ($V \approx V_R$). For these managers, the U-shaped relationship predicts that performance-based incentives should be strong.

Similarly, firms outside of distress are usually led by managers without a strong history of failure, which corresponds to managers with medium and high continuation values. This group represents most firms, and for these managers, our model predicts that pay-for-performance sensitivity increases in the level of past performance.

¹⁹See, for example, Gilson and Vetsuypens (1993) on the relationship between distress and pay-for performance incentives. This pattern appears to have been repeated during the recent recession and crisis, with the popular press questioning basic pay practices. (see, e.g., New York Times, April 3, 2010, “Striking Gold in Stock Options”). In contrast, to the best of our knowledge, there has been no direct empirical study of the broader question of the effect of performance on pay-for-performance sensitivity. However, Fernandes, Ferreira, Matos and Murphy (2010) show that recent stock returns are positively associated with incentive pay. Bushman, Dai, and Wang (2010) show that pay-for-performance sensitivity is decreasing in the probability of termination.

Termination and Performance

Our model produces a tight theoretical linkage between performance, monitoring, and termination, but it also explains why these variables might not appear to be tightly linked using standard empirical methods. For example, many empirical studies have found that while the relationship between performance and termination is negative, the explanatory power is low. Similarly, variables that represent monitoring ability or quality have a positive but small effect on termination probability.²⁰

Our model helps resolve the apparent lack of explanatory power exactly because the existence of a monitoring technology lowers the probability of termination unconditionally and increases the range in which the probability of future termination is insensitive to past performance. This is illustrated in Figure 3 of Section 4.2. In addition, as in Section 3, when termination cost is high (e.g. due to a high quality CEO-firm match; $\mu - rL - \gamma V_R$ high), monitoring can completely substitute for termination-based incentives. In this case, termination due to poor performance never occurs in equilibrium and the relationship between performance and termination completely disappears.

Security Design and Endogenous Financing Constraints

Finally, our model can be also used to study the implications of costly monitoring for optimal security design and endogenous financing constraints. The basic model without monitoring has a natural implementation as a long-term debt contract with cash balance (or credit commitment) as demonstrated by DeMarzo and Sannikov (2006) and Biais et al. (2007). We show in Appendix C that monitoring results in the addition of state-contingent short-term financing. This short-term financing acts as an insurance mechanism that allows some positive-NPV projects to survive negative cash flows, and there is liquidation in equilibrium (termination, default) only if the liquidation cost is low. As a result, active use of monitoring relaxes the endogenous financing constraint implied by the agency problem, allowing the creditor (the principal) to provide additional funds after negative cash flows to the borrower (the agent) without weakening incentives.

²⁰For studies of both the explanatory power of past performance, monitoring ability, and the interaction of the two, see, for example, Weisbach (1988), Jensen and Murphy (1990), Brickley (2003), Kaplan and Minton (2006), and Jenter and Lewellen (2014).

6 Conclusion

In this paper we provide a tractable characterization of the optimal long-term contract between an agent and a principal in a setting where the agent can privately divert resources for his own benefit, replacing the agent (contract termination) is costly, and the principal can monitor the manager at a cost. The important change relative to earlier work is the addition of a dynamic monitoring technology – a second signal in addition to project performance – that allows the principal to detect, at a cost, undesirable activity on the part of the agent.

First, we show that both optimal monitoring and the intensity of performance based compensation can be non-monotonic functions of the quality of past performance. In particular, agents with poor histories have smaller inside stakes, and further performance deterioration can lead to a shift away from monitoring and towards pay-for-performance incentives.

Second, our dynamic setting allows us to consider the effect of monitoring on early termination of the contract. We show that, when termination cost is high, monitoring can completely substitute for termination-based incentives. In this case, termination due to poor performance never occurs and the relationship between performance and termination completely disappears. When termination due to poor performance happens in equilibrium, active use of monitoring lowers the chances of termination and increases the range of performance where termination probability is very insensitive to past performance.

As we discussed in Section 5, these results may help explain some puzzling empirical findings on the relationship between performance history and future pay-performance sensitivity and on the linkage between termination, performance, and monitoring. Also, our model can be used to study the optimal dynamic relationship between costly monitoring and endogenous financing constraints implied by the agency problem.

We note that throughout the model we have assumed that the principal can commit to contract terms. However, the principal may not always want to punish the agent after bad news (negative cash flows or detecting shirking) because termination is costly for the principal. In addition, the principal may choose not to pay the monitoring cost if monitoring is not publicly observable. We discuss this further in Appendix B and show how the renegotiation proof contract in DeMarzo and Sannikov (2006) can be adapted to our setting.

Finally, one may view risk-neutrality and other details as “special” and wonder how important they are to the model’s qualitative results. In fact, our main insights are robust to various extensions of our setting. For example, one could apply our technology to the

setting of Sannikov (2008) in which the agent is risk-averse and there is a convex cost of effort. In that model, the marginal cost of effort is proportional to the volatility of the agent’s continuation value. With monitoring, the marginal cost of effort is proportional to the volatility of the agent’s continuation value plus the incentives provided by monitoring (i.e. one can replace the 1 in equation (7) with the marginal cost of effort). Time variation in the principal’s desired effort level will drive time variation in total incentives, while the “risk of termination” and “inside stake” effects in our model will continue to drive the composition of incentives.

A Proofs

A.1 Proof of Proposition 1

First²¹, we define the agent's total expected utility received under a contract ξ conditional on his information at time t as:

$$U_t = E^s \left[\int_0^\tau e^{-\gamma u} dC_u + \int_0^\tau e^{-\gamma u} s_u du + e^{-\gamma \tau} V_\tau | \mathcal{F}_t \right],$$

where $\tau = \min(\tau^d, \tau^f, \tau^q)$. 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 s , such that $dZ_t^s = \frac{1}{\sigma}(dY_t - \mu dt - s_t dt)$ is a Brownian motion and $\lambda(m_t, s_t - s_t^r)dt - dN_t^s$ is a martingale. Then, by the martingale representation theorem for Lévy processes, we have that there exists $\mathcal{F}_t$ -predictable processes $(\beta, \psi) = \{(\beta_t, \psi_t); 0 \leq t \leq \tau\}$ such that

$$U_t = U_0 + \int_0^t e^{-\gamma u} \beta_u \sigma dZ_u^s + \int_0^t e^{-\gamma u} \psi_u (\lambda(m_u, s_u - s_u^r) du - dN_u^s). \quad (\text{A-1})$$

Recall the agent's continuation value V_t^s defined in (4). Thus we have for $t \leq \tau$,

$$U_t = \int_0^t e^{-\gamma u} dC_u + \int_0^t e^{-\gamma u} s_u du + e^{-\gamma t} V_t^s. \quad (\text{A-2})$$

Differentiating (A-1) and (A-2), we obtain

$$\begin{aligned} dU_t &= e^{-\gamma t} \beta_t \sigma dZ_t^s + e^{-\gamma t} \psi_t (\lambda(m_t, s_t - s_t^r) dt - dN_t^s) \\ &= e^{-\gamma t} dC_t + e^{-\gamma t} s_t dt - \gamma e^{-\gamma t} V_t^s dt + e^{-\gamma t} dV_t^s, \end{aligned}$$

and so

$$dV_t^s = \gamma V_t^s dt - dC_t - s_t dt + \beta_t \sigma dZ_t^s + \psi_t (\lambda(m_t, s_t - s_t^r) dt - dN_t^s).$$

This equation also implies the evolution of promised value given in (6) for $s = s^r$.

Next, define the payoff to strategy s^* consisting of following an arbitrary strategy until time $t < \tau$ and then s^r thereafter to be $\hat{U}_t$, so that

$$\hat{U}_t = \int_0^t e^{-\gamma u} dC_u + \int_0^t e^{-\gamma u} s_u^* du + e^{-\gamma t} V_t^{s^r}.$$

²¹This proof is based on a similar proof in Sannikov (2008), extended to include jump processes and monitoring.

Let us identify the drift of the process $\hat{U}_t$ under the strategy s^* . After differentiating and combining terms,

$$\begin{aligned} e^{\gamma t} d\hat{U}_t &= (s_t^* - s_t^r)dt + \beta_t \sigma dZ_t^{s^r} + \psi_t (\lambda(m_t, 0)dt - dN_t^{s^r}) \\ &= (s_t^* - s_t^r)dt - \beta_t (s_t^* - s_t^r)dt - \psi_t m_t \max((s_t^* - s_t^r), 0)dt \\ &\quad + \beta_t \sigma dZ_t^{s^*} + \psi_t (\lambda(m_t, s_t^* - s_t^r)dt - dN_t^{s^*}) \end{aligned}$$

where the second line reflects a change in the probability measure from the one induced by $s = s^r$ to the one induced by s^* , and $\lambda(m_t, s_t^* - s_t^r)$ has been replaced by $m_t \max(s_t^* - s_t^r, 0)$ in the third term. If (7) does not hold on a set of positive measure, then one could chose s^* so that the drift of $\hat{U}$ is nonnegative and positive on a set of positive measure and so

$$E^{s^*} [\hat{U}_t] > \hat{U}_0 = V_0^{s^r},$$

and so the strategy s_t^r would not be optimal for the agent. If (7) does hold for the strategy s^r then $\hat{U}_t$ is a super-martingale (under measure induced by s^*) for any strategy s^* and so

$$E^{s^*} [\hat{U}_t] \leq \hat{U}_0 = V_0^{s^r}.$$

But this implies that choosing $s = s^r$ is optimal for the agent iff (7) does hold for the strategy s^r . (Note that $E^{s^*} [\hat{U}_\infty] \leq \hat{U}_0 = V_0^{s^r}$ as well as $V_t^{s^r}$ is bounded from below.)

To address the quitting time, observe that $\tau^d \leq \inf\{t \geq 0 : V_t^{s^r} = V_R\}$ implies that if the agent does not quit early ($\tau^q \geq \tau^d$) he has continuation utility of at least V_R , which is what he would obtain by quitting. $\square$

A.2 Proof of Proposition 2

We show that any incentive compatible contract such that $s_t^r > 0$ for $t \in \mathcal{T}$ can be weakly improved upon by a contract that sets $s_t^r = 0$ for all $t \in \mathcal{T}$.

We remember that for any recommendation s^r , the agent can always choose some $s_t > s_t^r$ at any $t < \tau = \min(\tau^d, \tau^f, \tau^q)$ – diverting more resources than recommended is always feasible for the agent. From Proposition 1 and equation 8, for $t \in \mathcal{T}$ such that $t \leq \tau$, the agent's continuation value evolves as

$$dV_t^{s^r} = \gamma V_t^{s^r} dt - dC_t - s_t^r dt + \beta_t \sigma dZ_t, \quad (\text{A-3})$$

for incentive compatible contracts. In addition, the contract is incentive compatible if and only if (7) holds and $\tau^d \leq \inf\{t \geq 0 : V_t^{s^r} = V_R\}$. This means that the incentive compatible level of β_t is not changed if we change the recommended level of diversion.

Now consider the contract that sets $dC'_t = dC_t + s_t^r dt$, $s_t^{r'} = 0$ for all $t \in \mathcal{T}$, and $dC'_t = dC_t$,

$s_t^{r'} = s_t^r = 0$ for all other t . Let the functions C, m, τ, β, ψ depend on Y as if the realized output were $Y_t = Y_t' - \int_0^t (s_t^r - s_t^{r'}) dt$ for any $t \leq \tau$, where Y_t' is the output under the new contract. Under this new contract the agent's continuation value evolves for all $t \in \mathcal{T}, t \leq \tau$ and $s = s^{r'}$ as

$$dV_t^{s^{r'}} = \gamma V_t^{s^r} dt - dC_t' + \beta_t (dY_t' - \mu dt),$$

which can be rewritten as

$$dV_t^{s^{r'}} = \gamma V_t^{s^{r'}} dt - dC_t - s_t^r + \beta_t \sigma dZ_t. \quad (\text{A-4})$$

Under the new contract the functions C, m, τ, β, ψ depend identically on realized Z as under the previous contract. Since equations (A-3) and (A-4) are identical and since outside the set $\mathcal{T}$ the new contract coincides with the old one, it follows that not only is the new contract incentive compatible but it also delivers the agent the same utility as the old contract. At the same time the expected utility of the principal is unchanged. Thus, any incentive compatible contract such that $s_t^r > 0$ for $t \in \mathcal{T}$ can be weakly improved upon by a contract that sets $s_t^r = 0$ for all $t \in \mathcal{T}$. $\square$

A.3 Algebra and Bounding for Property 4

To exclude a jump from $\beta_\epsilon \in (\epsilon, 1)$ to $\beta_\epsilon = 0$ on $V \in (V_R, V_{\epsilon, C})$, we compare the top and bottom parts of the maximization in (14). Equality in F_ϵ and F'_ϵ implies that the positive value of β_ϵ must be such that $0 = \frac{\theta}{V - V_F} \beta_\epsilon + \frac{1}{2} \sigma^2 \beta_\epsilon^2 F''_\epsilon$. The factor of $\frac{1}{2}$ makes this expression inconsistent with the first-order condition (12), and thus excludes any interior value of β_ϵ . One can similarly exclude a jump from $\beta_\epsilon = 1$ to $\beta_\epsilon = 0$.

It remains to exclude a jump from $\beta_\epsilon = \epsilon$ to $\beta_\epsilon = 0$ on $V \in (V_R, V_{\epsilon, C})$. Taking the derivative of the top expression in (14) to obtain a value for F'_ϵ and then substituting $F'_\epsilon(V)$ out of the original expression, we obtain

$$\mu - r F_\epsilon(V) = \frac{\theta}{V - V_F} + \frac{\gamma \theta}{\gamma - r} \frac{V}{(V - V_F)^2} + V^2 \frac{\gamma^2}{\gamma - r} F''_\epsilon(V)$$

Since the principal cannot do better than the first-best, the left hand side must be positive. Since the right hand side must then also be positive, and we have $V > V_R > V_F$, the above expression implies a lower bound on F''_ϵ that does not depend on ϵ . In contrast, inspection of (15) shows that if $\beta_\epsilon = \epsilon$, then F''_ϵ is unboundedly negative for sufficiently small ϵ . Thus, the necessary super-contact condition (continuous $F''_\epsilon(V)$) cannot be met for small ϵ .

A.4 Proof of Theorem 3

The proof will proceed as follows. First, assuming the existence of a solution to (14) we will prove two useful lemmas characterizing such solutions, which we will use for both the restricted and unrestricted problems. Then we will show the existence and uniqueness of the solution to (14) given the initial conditions. Finally, we will verify that the postulated solution indeed satisfies the required boundary conditions and solves the restricted contracting problem.

We start with two preliminary results that we will use for both the restricted and unrestricted problems.

Lemma 7 *Let $\hat{F}$ solve either (14) for some $\epsilon > 0$ or (13). Assume that the boundary conditions in Property 1 are met at some $\hat{V}_C$, but ignore the boundary condition at V_R . Then $\hat{F}''(V) < 0$ for all $V < \hat{V}_C$ on which $\hat{F}$ is defined.*

Proof. Because concavity is a necessary condition for $\beta < 1$ from Property 3, we need only show the result when $\beta = 1$. Combining either (13) or (14) with the first derivative with respect to V implies

$$0 = -r\hat{F}(V) + \mu + \gamma V\hat{F}'(V) + \frac{1}{2}\sigma^2\hat{F}''(V) \quad (\text{A-5})$$

$$0 = (\gamma - r)\hat{F}'(V) + \gamma V\hat{F}''(V) + \frac{1}{2}\sigma^2\hat{F}'''(V) \quad (\text{A-6})$$

Since the principal cannot do better than the first-best, we have $\mu > r\hat{F}(V)$. If $\hat{F}''(V) = 0$, then (A-5) and $\mu > r\hat{F}(V)$ imply $\hat{F}'(V) < 0$. Then, $\hat{F}''(V) = 0$, $\hat{F}'(V) < 0$, and (A-6) together imply $\hat{F}'''(V) > 0$. Thus $\hat{F}''(V)$ can only cross zero once, and only from below. Since $\hat{F}''(V_C) = 0$, we must have $\hat{F}''(V < V_C) < 0$. ■

Lemma 8 *Let $\hat{F}$ and $\tilde{F}$ both solve either (14) for some $\epsilon > 0$ or (13). Assume that the boundary conditions in Property 1 are met at $\hat{V}_C$ and $\tilde{V}_C$, respectively, but ignore the boundary conditions at V_R . Then the following four statements are equivalent:*

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

Proof. We will show that $\hat{V}_C > \tilde{V}_C$ implies the other three conditions. Then, one can use the same arguments to show that $\hat{V}_C < \tilde{V}_C$ implies opposite conditions, and thus we

have “if and only if”. Note that plugging boundary conditions at $\hat{V}_C$ and $\tilde{V}_C$ into the HJB equation and using concavity (Lemma 7), $\hat{V}_C > \tilde{V}_C$ if and only if $\hat{F}(\tilde{V}_C) < \tilde{F}(\tilde{V}_C)$ and $\hat{F}'(\tilde{V}_C) > \tilde{F}'(\tilde{V}_C)$.

First, if $\hat{F}(V) < \tilde{F}(V)$ for any $V \leq \tilde{V}_C$, then $\hat{F} < \tilde{F}$ for all $V \leq \tilde{V}_C$ in the shared domain. Assume to obtain a contradiction those two solutions cross at some $\check{V}$ s.t. $\hat{F}(\check{V}) = \tilde{F}(\check{V})$ and $\hat{F}'(\check{V}) \neq \tilde{F}'(\check{V})$. Consider the largest such $\check{V}$; since $\hat{F}(\tilde{V}_C) < \tilde{F}(\tilde{V}_C)$, we must have $\hat{F}'(\check{V}) < \tilde{F}'(\check{V})$.

There can be two types of crossing: $\exists V' \in (\check{V}, \tilde{V}_C)$ s.t. $\hat{F}'(V') = \tilde{F}'(V')$, or there does not exist such a V' . If not, then we have $\hat{F}'(\tilde{V}_C) < \tilde{F}'(\tilde{V}_C)$ from continuity, which contradicts the boundary conditions. Thus we must have $\exists V' > \check{V}$ s.t. $\hat{F}'(V') = \tilde{F}'(V')$, and we consider the smallest V' such that this is the case. Then, this construction implies that $\hat{F}'(V) < \tilde{F}'(V)$ for $V \in (\check{V}, V')$ and that $\hat{F}(V') < \tilde{F}(V')$.

Notice that the right hand side of the HJB equation (13 or 14) is decreasing in $\hat{F}(V')$ and increasing in $\hat{F}''(V')$: $\max \left[-\frac{\theta}{V' - V_F}(1 - \beta) + \frac{1}{2}\beta^2\sigma^2\hat{F}''(V') \right]$ attains a higher value for higher $\hat{F}''(V')$. Thus, $\hat{F}(V') < \tilde{F}(V')$ and $\hat{F}'(V') = \tilde{F}'(V')$ imply $\hat{F}''(V') < \tilde{F}''(V')$. But, $\hat{F}'(V') = \tilde{F}'(V')$ and $\hat{F}''(V') < \tilde{F}''(V')$ imply $\exists \delta > 0$ s.t. $\hat{F}'(V' - \delta) > \tilde{F}'(V' - \delta)$, which contradicts $\hat{F}'(V) < \tilde{F}'(V)$ for $V \in (\check{V}, V')$.

Thus, $\hat{F}(V) = \tilde{F}(V)$ and the $\tilde{V}_C$ boundary condition require $\hat{F}'(V) = \tilde{F}'(V)$, and, by the HJB equation, equality in the second derivative; solutions cannot cross.

Second, we will extend that argument to show that if $\hat{F}'(V) > \tilde{F}'(V)$ for any V , then $\hat{F}' > \tilde{F}'$ for all V in the shared domain. Assume to obtain a contradiction that the two derivatives cross at some $\check{V}$ s.t. $\hat{F}'(\check{V}) = \tilde{F}'(\check{V})$ and $\hat{F}''(\check{V}) \neq \tilde{F}''(\check{V})$, and that $\check{V}$ is the largest V with such a crossing.

We note from the HJB that $\hat{F}''(\check{V}) > \tilde{F}''(\check{V})$ implies $\hat{F}(\check{V}) > \tilde{F}(\check{V})$ (since, as before, the right hand side of the HJB equation (13 or 14) is increasing in $\hat{F}''(\check{V})$). With $\hat{F}(\check{V}) > \tilde{F}(\check{V})$, the first part of the lemma shows that we must have $\hat{F}(V) > \tilde{F}(V)$ for all V , and so we must have $\hat{V}_C < \tilde{V}_C$, contradicting our original assumption. Thus, any derivative crossing must have $\hat{F}''(\check{V}) < \tilde{F}''(\check{V})$, and thus $\hat{F}(V) < \tilde{F}(V)$ for all V .

Then, at $\tilde{V}_C$, the boundary conditions imply $\hat{F}'(\tilde{V}_C) > \tilde{F}'(\tilde{V}_C)$ and $\hat{F}''(\tilde{V}_C) < \tilde{F}''(\tilde{V}_C)$. This generates a contradiction: if $\hat{F}'(\check{V}) = \tilde{F}'(\check{V})$ and $\hat{F}''(\check{V}) < \tilde{F}''(\check{V})$, then to obtain $\hat{F}'(\tilde{V}_C) > \tilde{F}'(\tilde{V}_C)$, there must exist a $V' > \check{V}$ such that $\hat{F}'(V') = \tilde{F}'(V')$, which contradicts the definition of $\check{V}$ as the largest crossing.

Thus, $\hat{F}'(V) = \tilde{F}'(V)$ implies $\hat{F}''(V) = \tilde{F}''(V)$ (and, by the HJB equation, equality in value), and derivatives cannot cross.

From the HJB equation, if $\hat{F} > \tilde{F}$ everywhere and $\hat{F}' < \tilde{F}'$ everywhere, then we have $\hat{F}'' > \tilde{F}''$. ■

We now proceed to analyze the HJB equation.²² A necessary condition for optimality (Property 4) is that $\beta_\epsilon(V) > 0$ on $V \in (V_R, V_{\epsilon,C}]$, so we will consider that region first. Re-writing (14), we have

$$F''_\epsilon(V) = \min_{\beta_\epsilon \in [\epsilon, 1]} \frac{rF_\epsilon(V) - \mu + \frac{\theta}{V-V_F}(1 - \beta_\epsilon) - \gamma VF'_\epsilon(V)}{\frac{1}{2}\beta_\epsilon^2 \sigma^2} \quad (\text{A-7})$$

The right-hand-side can be written as the function $H_{\epsilon,\beta}(V, F_\epsilon(V), F'_\epsilon(V))$, which is differentiable in all of its arguments. Since β_ϵ is bounded, $H_{\epsilon,\beta}$ has uniformly bounded derivatives in F'_ϵ and F_ϵ , and $H_{\epsilon,\beta}$ is Lipschitz continuous in F_ϵ and F'_ϵ . Notice also that the growth conditions are satisfied. It follows that solutions to (A-7) exist globally, are $\mathcal{C}^2$ and are unique and continuous in initial conditions.²³

To complete our solution, we need to show that our remaining condition can be met: that for any L , either there exists a $V_{\epsilon,C}$ such that $F_\epsilon(V_R; V_{\epsilon,C}) = L$ and $\beta_\epsilon(V_R) > 0$, or there exists a $V_{\epsilon,C}$ such that $F_\epsilon(V_R; V_{\epsilon,C}) \geq L$ and $\beta_\epsilon(V_R) = 0$. We need only consider $\beta = 0$ at V_R from Property 4.

Lemma 8 shows that proposed solutions to (14) that obey the right boundary condition (at $V_{\epsilon,C}$) can be ranked by the proposed value $\hat{V}_{\epsilon,C}$. For $\hat{V}_{\epsilon,C} = V_R$, we have the highest solution with $rF_\epsilon(V_R = \hat{V}_{\epsilon,C}) = \mu - \gamma V_R$. As we increase $\hat{V}_{\epsilon,C}$, $F_\epsilon(V; \hat{V}_{\epsilon,C})$ declines for all V . In particular, $\lim_{V \rightarrow V_R} F_\epsilon(V; \hat{V}_{\epsilon,C}) = F_\epsilon(V_R; \hat{V}_{\epsilon,C})$ declines as $\hat{V}_{\epsilon,C}$ increases.

Within the set of solutions that obey the right boundary condition, we consider the subset with $\beta_\epsilon(V_R; \hat{V}_{\epsilon,C}) > 0$. Label that infimum value of $F_\epsilon(V_R; \hat{V}_{\epsilon,C})$ as L_ϵ^* and the corresponding value of $\hat{V}_{\epsilon,C}$ as $V_{\epsilon,C}^*$. Any solution with $\beta_\epsilon(V_R; \hat{V}_{\epsilon,C}) > 0$ must be better than the principal's value with $\beta_\epsilon = 0$ and immediate payment, which is $\frac{\mu}{r} - \frac{1}{r} \frac{\theta}{V_R - V_F} - \frac{\gamma}{r} V_R$. Thus, L_ϵ^* and $V_{\epsilon,C}^*$ must be finite.

Because solutions are continuous in initial conditions, and both F and F'' are ordered, L_ϵ^* is also the maximum value of $F_\epsilon(V_R; \hat{V}_{\epsilon,C})$ among all proposed solutions with $\beta_\epsilon(V_R; \hat{V}_{\epsilon,C}) = 0$. Then, consider the proposed solution with $V_{\epsilon,C}^*$ such that $\beta_\epsilon(V_R; V_{\epsilon,C}^*) = 0$ and $F_\epsilon(V_R; V_{\epsilon,C}^*) = L_\epsilon^*$. This solution can be implemented for any value of L : because $\beta_\epsilon(V_R; V_{\epsilon,C}^*) = 0$, there is no termination and L is never realized. This solution is preferred to all other no-termination solutions by construction.

If, instead, $L > L_\epsilon^*$, then the principal prefers to allow termination, with $\beta_\epsilon(V_R) > 0$.

²²This 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.

²³To see these conditions directly, plug (12) into (13) to observe that for regions in which $\beta \in (\epsilon, 1)$ we have $rF_\epsilon(V) = \mu + \frac{\theta}{V-V_F}(\frac{1}{2}\beta_\epsilon(V) - 1) + \gamma VF'_\epsilon(V)$. This first-order ODE can be solved through direct integration, taking β_ϵ to be an unknown bounded function. Similarly, in regions in which $\beta_\epsilon = \epsilon$ or $\beta_\epsilon = 1$, we have constant or linear coefficients and standard results imply existence and uniqueness there (see e.g. Zhu (2013)). We use the more powerful Lipschitz continuity to show the solution is $\mathcal{C}^2$.

Since L is greater than the infimum value of $F(V_R; \hat{V}_{\epsilon,C})$ among all proposed solutions with $\beta_\epsilon(V_R; \hat{V}_{\epsilon,C}) > 0$, and solutions to the ODE are continuous in initial conditions and ordered (Lemma 8), there exists exactly one solution with $\hat{V}_{\epsilon,C} < V_{\epsilon,C}^*$ such that $F(V_R; \hat{V}_{\epsilon,C}) = L$.

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

Standard existence and uniqueness results (see e.g. Karatzas and Shreve (1991)), 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 (8) has a unique solution and that V_t is a Sticky Brownian Motion near V_R if $\beta_\epsilon(V_R) = 0$.

The remainder of the proof proceeds as any standard verification theorem. For any incentive compatible contract, we define

$$G_{\epsilon,t} = \int_0^t e^{-rs} (dY_s - \theta m_{\epsilon,t} dt - dC_s) + e^{-rt} F_\epsilon(V_t),$$

where V_t evolves according to (8). The process G is such that $G_{\epsilon,t}$ is measurable and

$$e^{rt} dG_{\epsilon,t} = dY_t - \theta m_{\epsilon,t} dt - dC_t + dF_\epsilon(V_t) - rF_\epsilon(V_t).$$

From Ito's lemma using (8) we have

$$dF_\epsilon(V_t) = (\gamma V_t dt - dC_t) F'_\epsilon(V_t) + \frac{1}{2} \beta_{\epsilon,t}^2 \sigma^2 F''_\epsilon(V_t) dt + \beta_{\epsilon,t} \sigma F'_\epsilon(V_t) dZ_t,$$

and $dY_t = \mu dt + \sigma dZ_t$. Thus,

$$\begin{aligned} e^{rt} dG_{\epsilon,t} &= \left[\mu - \theta m_{\epsilon,t} + \gamma V_t F'_\epsilon(V_t) + \frac{1}{2} \beta_{\epsilon,t}^2 \sigma^2 F''_\epsilon(V_t) \right] dt \\ &\quad - rF_\epsilon(V_t) - (1 + F'_\epsilon(V_t)) dC_t + (\sigma + \beta_{\epsilon,t} \sigma F'_\epsilon(V_t)) dZ_t \end{aligned}$$

where we remember that the incentive compatibility is equivalent to $\beta_{\epsilon,t} \geq 1 - (V_t - V_F) m_{\epsilon,t}$.

Substituting from (13) for $rF_\epsilon(V_t)$ in the above yields

$$\begin{aligned} e^{rt} dG_{\epsilon,t} &= \left[\theta m_\epsilon(V_t) - \theta m_{\epsilon,t} + \frac{1}{2} \sigma^2 (\beta_{\epsilon,t}^2 - [1 - (V_t - V_F) m_\epsilon(V_t)]^2) F''_\epsilon(V_t) \right] dt \\ &\quad - (1 + F'_\epsilon(V_t)) dC_t + (\sigma + \beta_{\epsilon,t} \sigma F'_\epsilon(V_t)) dZ_t, \end{aligned}$$

whenever $V \in (V_R, V_{\epsilon,C})$.

Observe that for any policies, we have

$$\left[\theta m_\epsilon(V_t) - \theta m_{\epsilon,t} + \frac{1}{2} \sigma^2 (\beta_{\epsilon,t}^2 - [1 - (V_t - V_F) m_\epsilon(V_t)]^2) F_\epsilon''(V_t) \right] \leq 0$$

and

$$-(1 + F_\epsilon'(V_t)) dC_t \leq 0$$

with equality for the optimal policies. The first equation comes from the specification of $m_\epsilon(V_t)$ and the second comes from $F_\epsilon' \geq -1$ and $dC_t \geq 0$. Thus, we have

$$F_\epsilon(V_0) = G_{\epsilon,0} \geq E [G_{\epsilon, \min(\tau^d, \tau^f)}]$$

where $G_{\epsilon,\infty} = \int_0^\infty e^{-rs} (dY_s - \theta m_{\epsilon,t} dt - dC_s)$ since $\lim_{t \rightarrow \infty} e^{-rt} F_\epsilon(V_t) = 0$. Since we have equality for the desired optimal policies, any other policies deliver less in gains to the principal.

A.5 Proof of Theorem 4

We continue the analysis from Theorem 3. For a proposed $\hat{V}_C$, both the unrestricted and restricted HJB equations (13) and (14) have $\beta(\hat{V}_C) = 1$. Then, the arguments in the proof of Theorem 3 are sufficient to show that there is a unique, $\mathcal{C}^2$, and continuous in initial conditions solution to the HJB equation (13) near $\hat{V}_C$. In particular, Lemmas 7 and 8 apply to solutions to the unrestricted problem as well as the restricted problem.

For a fixed $\hat{V}_C$ and the corresponding solution to the *restricted* ODE (14, for small $\epsilon > 0$), define $Q_\epsilon(\hat{V}_C) = \max\{V | \beta_\epsilon(V; \hat{V}_C) = \epsilon\}$ if there exists a V such that $\beta_\epsilon(V; \hat{V}_C) = \epsilon$, and $Q_\epsilon(\hat{V}_C) = V_R$ if there does not exist such a V . We can use max rather than sup from the continuity of $\beta_\epsilon(V > V_R)$ with respect to initial conditions. Intuitively, $Q_\epsilon(\hat{V}_C)$ separates V into two regions, and the constraint $\beta_\epsilon \geq \epsilon$ is not binding for the region $V > Q_\epsilon(\hat{V}_C)$. The important properties of $Q_\epsilon(\hat{V}_C)$ are

1. $Q_\epsilon(\hat{V}_C)$ is continuous in $\hat{V}_C$ because solutions to the ϵ -restricted problems (and therefore the choice of β_ϵ) are each continuous in initial conditions
2. If $Q_\epsilon(\hat{V}_C) > V_R$, then $Q_\epsilon(\hat{V}_C)$ is strictly increasing in $\hat{V}_C$, from the result on concavity in Lemma 8 and the optimal choice of β_ϵ (15).
3. There exist $\hat{V}_C$ such that $Q_\epsilon(\hat{V}_C) = V_R$, such as $\hat{V}_C = V_R$. The supremum of such $\hat{V}_C$ is finite.

Next, define $Q(\hat{V}_C) = \inf_\epsilon \{Q_\epsilon(\hat{V}_C)\}$. Intuitively, if $Q(\hat{V}_C) > V_R$, then $Q(\hat{V}_C)$ is the highest value of V for which the *unrestricted* problem has a singularity ($\beta = 0$ or $F'' = -\infty$). Notice

that for any $\hat{V}_C$ and $\epsilon > 0$, the unrestricted and restricted HJB equations are the same for $(Q_\epsilon(\hat{V}_C), \hat{V}_C]$. In particular, for any $\hat{V} > Q(\hat{V}_C)$, our construction implies there exists an $\epsilon > 0$ such that $\hat{V} > Q_\epsilon(\hat{V}_C)$. Thus, for any $\hat{V}_C$, there is a unique, $\mathcal{C}^2$, and continuous in initial conditions solution to the unrestricted HJB equation (13) for $V \in (Q(\hat{V}_C), \hat{V}_C]$.

Label as V_C^* the infimum value of $\hat{V}_C$ such that $Q(\hat{V}_C) > V_R$. By continuity, V_C^* is also the maximum value of $\hat{V}_C$ such that $Q(\hat{V}_C) = V_R$. Then, V_C^* is associated with a solution to the HJB equation that has $\beta(V_R; V_C^*) = 0$ and $\beta(V > V_R; V_C^*) > 0$.

Further, a solution to the HJB equation with $V_C = V_C^*$ is $\mathcal{C}^2$ over $V \in (V_R, V_C^*]$. This solution is $\mathcal{C}^1$ at V_R because $F(V; V_C^*)$ is $\mathcal{C}^2$ over $V \in (V_R, V_C^*]$ and (12) implies $\lim_{V \rightarrow V_R} \beta(V)^2 F''(V) = 0$. Thus, direct evaluation of the HJB equation shows that a finite and continuous F and $V_F < V_R$ imply a finite and continuous F' .

$F_\epsilon(V_R)$ is increasing as $\epsilon \downarrow 0$ because the set of allowable policies expands, and $F_\epsilon(V_R)$ is bounded above by the first-best. Then, from the monotone convergence theorem, we must have $\lim_{\epsilon \rightarrow 0} L_\epsilon^* = L^* = F(V_R) = \lim_{\epsilon \rightarrow 0} F_\epsilon(V_R)$. Because $Q(V_C^*) = V_R$, the solution to the HJB (13) using V_C^* and evaluated at V_R is the same as the limit to the monotone sequence $F_\epsilon(V_R)$. We note that because $\beta(V_R; V_C^*) = 0$, V_t reflects upward at $V_t = V_R$ so that there is no termination, and this contract is always feasible.

Thus, any solution with $F(V_R) < L^*$ is impossible because it would be strictly inferior to the solution with $F(V_R; V_C^*) = L^*$ and $\beta(V_R; V_C^*) = 0$. The existence and uniqueness of $\{V_C^*, L^*\}$ shows the existence and uniqueness of solutions to the HJB with $L < L^*$. Any solution with $L > L^*$, by construction, has $\beta(V) > 0$, which means it is the same as the solution to at least one of the ϵ -restricted problems and preferred to the $F(V_R) = L^*$ solution. Standard arguments from the proof of Theorem 3 show the existence and uniqueness of solutions to the HJB with $L > L^*$.

Existence of a (weak) solution to (8) is standard (see, e.g. Theorem 4.22 in Karatzas and Shreve (1991)). The verification argument for the unrestricted problem closely follows the argument presented in the proof of Theorem 3.

A.6 Proof of Proposition 5

The first statement of the proposition follows directly from Lemma 8 and (12).

The second statement of the proposition follows from the fact that $F(V)$ must be weakly decreasing in V_F and θ (from inspection of the problem). In particular, this must be true for $F(V_R)$ for no-termination solutions (the change in $F(V_R)$ is zero for termination solutions).

Third, we show that $\beta'(V) = 0$ and $\beta(V) \in [0, 1)$ implies $\beta''(V) > 0$. Combining (12) and (13) yields

$$\beta(V) = -\frac{2}{\theta}(V - V_F) \left(\mu - rF(V) + \frac{\theta}{V - V_F} + \gamma VF'(V) \right). \quad (\text{A-8})$$

Taking the derivative of (A-8), setting it equal to zero, and then substituting in (12) and re-arranging yields

$$F'(V) = \frac{-(\mu - rF(V)) + \frac{V\gamma\theta}{\sigma^2\beta(V)}}{\gamma V + (\gamma - r)(V - V_F)} \quad (\text{A-9})$$

Taking the second derivative of (A-8) and substituting in (12), (A-9), and $\beta'(V) = 0$, we obtain

$$\beta''(V) = 4 \frac{(\gamma - r)(\mu - rF(V))}{\theta(\gamma V + (\gamma - r)(V - V_F))} + 2 \frac{\gamma^2 V(2V - V_F) + (\gamma - r)(2\gamma - r)(V - V_F)^2}{(V - V_F)\sigma^2\beta(V)(\gamma V + (\gamma - r)(V - V_F))}$$

which is strictly positive (since the principal cannot do better than the first best, we have $\mu - rF(V) > 0$).

Fourth, we show that we can sign $\beta'(V)$. We start with the HJB (13) and take its derivative (using the envelope theorem and the super-contact condition) to find

$$0 = (\gamma - r)F'(V) + \gamma VF''(V) - \frac{\theta}{(V - V_F)^2}(\beta - 1) + \frac{1}{2}\beta^2\sigma^2 F'''(V). \quad (\text{A-10})$$

Then, we combine (13) and (A-10) to eliminate terms in $F'(V)$:

$$\begin{aligned} 0 = & \theta \frac{(1 - \beta)(\gamma V + (\gamma - r)(V - V_F))}{\gamma(V - V_F)^2 V} + (rF(V) - \mu) \frac{\gamma - r}{\gamma V} \\ & + F''(V) \left(\gamma V - \frac{1}{2}\beta^2\sigma^2 \frac{\gamma - r}{\gamma V} \right) + \frac{1}{2}\beta^2\sigma^2 F'''(V) \end{aligned} \quad (\text{A-11})$$

Next, we take the derivative of the optimal value of β (from 12) to find $\beta' > 0$ if

$$0 < F''(V) + (V - V_F)F'''(V) \quad (\text{A-12})$$

Combining (12) and (A-12), we see that if $1 > \beta > 0$, then $\beta' > 0$ if and only if

$$F'''(V) > \frac{\theta}{\beta(V - V_F)^2\sigma^2} \quad (\text{A-13})$$

To continue, combining (A-11) and (A-13), we see that if $1 > \beta > 0$, then $\beta' > 0$ if and only if

$$F(V) < \frac{\mu}{r} + \frac{\theta}{r} \frac{\beta(\beta - 2)\sigma^2(\gamma V + (V - V_F)(\gamma - r)) + 2\gamma^2 V^2(V - V_F)}{2\beta\sigma^2(\gamma - r)(V - V_F)^2} \quad (\text{A-14})$$

This expression is continuous in β .

Finally, we fix $V = V_R$. Then (A-14) with equality describes a collection of points

in (β, F) space in which $\beta'(V_R) = 0$. It is a continuous curve from $(\beta = 0, F = \infty)$ to $(\beta = 1, F = \alpha)$, where α equals (A-14) evaluated at $\beta = 1$. This curve is decreasing, except possibly for $F < \alpha$.

Separately, we know from Theorem 4, Lemma 8, and (12), that for a fixed V_R , the true value of β increases from 0 to 1 as F increases from L^* to $\frac{\mu}{r}$. This is accomplished by varying L from L^* to $\frac{\mu}{r}$ and collecting points in (β, F) space as we change economies. Thus, we describe a second curve (which is also weakly monotonic and continuous except possibly at $\beta = 1$) for which (β, F) are parts of solutions to the HJB.

We illustrate the two curves in Figure 6.




Figure 6: Illustration for the proof of Proposition 5(iv). The dashed black line depicts points $\{\beta(V_R), F(V_R)\}$ solving A-14 which have the property that $\beta'(V_R) = 0$. The solid black line depicts $\{\beta(V_R), F(V_R)\}$ that are solutions to the HJB equation (13) as we vary L from 20 to 35. The solid black line is strictly increasing for $0 < \beta < 1$. The existence of the intersection $\{\beta^*, \alpha^*\}$ shows the existence of an L such that $\beta'(V_R) = 0$. The sufficient condition for an intersection in the proof of the proposition implies that the value labeled α is less L^* . The curves are plotted using $r = .01$, $\gamma = .05$, $\sigma = 1$, $V_R = 2.25$, $V_F = 0$, $\theta = 1$, and $\mu = .5$.

If $\alpha < L^*$, then there is one intersection of the two curves, and that intersection takes place at some $1 > \beta > 0$. At that intersection, we find a value of $(\beta(V_R), F(V_R))$ such that the pair is a solution to the HJB that sets $\beta'(V_R) = 0$. Label the value of $F(V_R)$ at the

intersection α^* and the value of $\beta(V_R)$ at the intersection β^* . From the inequality in (A-14), if the solution to the HJB has $F(V_R) < \alpha^*$, then $\beta'(V_R) > 0$.

To conclude, we need to show that if V_F is large enough, then $\alpha < L^*$. The principal can always offer the agent a (sub-optimal) contract with full monitoring and immediate consumption, in which case the principal obtains $\frac{\mu}{r} - \frac{\theta}{r} \frac{1}{V_R - V_F} - \frac{\gamma}{r} V$, and this value must be less than L^* . Substituting $\beta = 1$ into (A-14), a sufficient condition for $\alpha < L^*$ is

$$\frac{\mu}{r} - \frac{\theta}{r} \frac{1}{V_R - V_F} - \frac{\gamma}{r} V_R > \frac{\mu}{r} + \frac{\theta - \sigma^2(\gamma V_R + (V_R - V_F)(\gamma - r)) + 2\gamma^2 V_R^2 (V_R - V_F)}{2\sigma^2(\gamma - r)(V_R - V_F)^2} \quad (\text{A-15})$$

which can be re-arranged to find

$$\sigma^2(\gamma - r)\theta(V_R - V_F) + 2\sigma^2(\gamma - r)\gamma V_R(V_R - V_F)^2 - \theta\sigma^2\gamma V_R + 2\theta\gamma^2 V_R^2(V_R - V_F) < 0 \quad (\text{A-16})$$

The quadratic coefficient for $V_R - V_F$ is positive and the constant coefficient is negative, so the condition in (A-16) is met if $V_R - V_F$ is in a neighborhood of zero.

Thus, if $V_R - V_F$ is sufficiently small, there exists α^* with $F(V_R) = \alpha^*$ and $\beta'(V_R) = 0$. In the statement of the proposition, we label α^* as L_b . L_a is the infimum value of L such that $\beta(V_R) = 1$, and the ordering $L_a > L_b > L^*$ follows from the arguments above.

A.7 Proof of Lemma 6

The arguments of Proposition 1 follow, with the replacement of $\lambda(m_t, s_t - s_t^r)$ by $\lambda(m_t, s_t - s_t^b)$. Thus, (6) becomes

$$dV_t^{s^r} = \gamma V_t^{s^r} dt - dC_t - s_t^r dt + \beta_t (dY_t - (\mu - s_t^r)dt) + \psi_t (\lambda(m_t, s_t^r - s_t^b)dt - dN_t^{s^r}),$$

and (7) retains only the $s_t^r > 0$ case (since $s^b \geq 0$).

To keep the agent's continuation utility a martingale, V_t must drift up at the rate $\psi_t \lambda(m_t, s_t^r - s_t^b)dt$ to compensate the agent for the risk of a false-positive detection event. By accepting false positives, the principal effectively creates an additional random noise in the agent's continuation value.

We can re-write the HJB equation (9) as

$$\begin{aligned} rF(V) = & \max_{\substack{\beta \geq 1 - \psi m \\ m \geq 0 \\ \psi \leq V - V_F \\ s^b, s^r \geq 0}} \left[\mu - \theta m - s^r + (\gamma V_t - s^r)F'(V) + \frac{1}{2}\beta^2 \sigma^2 F''(V) \right. \\ & \left. + \lambda(m, s^r - s^b) (\psi F'(V) + F(V - \psi) - F(V)) \right]. \end{aligned}$$

The important term is the last one: $\lambda(m, s^r - s^b) (\psi F'(V) + F(V - \psi) - F(V))$. The principal's value function changes because of the additional drift in V and the jump of size ψ and probability $\lambda(m_t, s_t^r - s_t^b)$. We note that Theorem 4 implies that with $s^r = s^b$ the solution to the above equation exists, is concave with $\psi > 0$. Because the solution is strictly concave (except at the single point $V = V_C$), we have $F(V - \psi) < F(V) - \psi F'(V)$ for any $\psi > 0$. Thus, it is optimal to keep $s^r = s^b$ to make $\lambda(m, s^r - s^b) = 0$. This result is an application of Jensen's inequality to the non-informative noise that $s^r > s^b$ adds to the agent's continuation value.

In addition, $F'(W) \geq -1$ shows that the optimal value of s^r is zero, which is a second proof of Proposition 2.

B Commitment

In our setting, as with many principal-agent models, the principal commits to a sequence of policies that induce the agent to take a desired action. However, the principal may want to renegotiate the contract once she sees the outcome. There are three types of renegotiation we may be concerned with. First, the principal may want to deviate from the optimal monitoring policy given that shirking is an out-of-equilibrium event. Second, the principal may not want to impose the full penalty $\psi_t = V_t - V_F$ after catching the agent shirking. Third, the principal may not want to reduce the agent's continuation value as much as the contract specifies after bad cash flow news. This is the case if $F'(V) > 0$: welfare could be improved with a one-time surprise increase in the agent's continuation value.

The second and third types of renegotiation can be made sub-optimal by putting constraints on the contract. The first type – that the principal may want to choose not to monitor given that shirking is an out-of-equilibrium event – must be dealt with outside the model. We can assume that the principal is able to make a verifiable announcement of the level of monitoring, for example by hiring an external auditor at known available terms.²⁴

To make the contract proof against the second and third type of renegotiation, we must solve the original optimization problem under the constraint that $F(V)$ is non-increasing. This ensures both that the principal's welfare cannot be improved with a one-time increase in the agent's continuation value and also that the principal's value is maximized by imposing the harshest possible punishment. DeMarzo and Sannikov (2006) demonstrate how to do this in the setting without monitoring, and the same solution applies in our setting:

Proposition 9 *Solve (13) for boundary conditions $F'(V_C) = -1$ and $F''(V_C) = 0$ where the maximum of the resulting function is L with $F(V^*) = L$. Assume that the agent's*

²⁴In practice, all U.S. public firms are required by law and regulation to have some of their financial statements audited by a publicly named third-party auditor, and firms can choose to undergo additional audit-based certification as well.

continuation value evolves as

$$dV_t = \gamma V_t dt - dC_t + \beta_t \sigma dZ_t + dP_t \quad (\text{B-17})$$

where the process P ($P_0 = 0$) reflects V_t at $V_* > V_R$. The contract is terminated randomly at $V_t = V_*$, with the probability that the project continues equal to $\exp\left(\frac{-P_t}{V_* - V_R}\right)$. If such a contract exists, then it is the optimal renegotiation-proof contract.

Proof. See DeMarzo and Sannikov (2006). ■

Making the contract renegotiation-proof induces qualitative changes in one of our results: while partial monitoring can still occur in equilibrium (i.e. $\beta(V) < 1$ for some V), it is the case that all optimal contracts have termination in equilibrium (i.e. $\beta(V) > 0$ for all V). In order for punishment to be credible, the principal must receive more in the event of punishment than otherwise. Thus, for full monitoring at some $\hat{V}$ to be renegotiation-proof, we require $L \geq F(\hat{V})$. In addition, the principal must pay to monitor, so full monitoring will be undertaken only when avoiding termination is strictly preferred to default: $L < F(\hat{V})$. These conditions are mutually exclusive, and so there is no such $\hat{V}$.

It is still the case that *partial* monitoring occurs in equilibrium. To see an example of such a solution, consider any solution to the baseline (non-renegotiation-proof) problem that has some $\hat{V}$ with $\beta(\hat{V}) < 1$ and $F'(\hat{V}) < 0$. Many of the solutions illustrated in Figures 2, 4, or 5 have this property. If we restrict the domain of the solution to $V \in [V^* = \arg \max_V F(V), V_C]$ (i.e. we cut off the $F'(V) > 0$ region), then this new restricted solution is a renegotiation-proof contract for the economy with $L = \max_V F(V) = F(V^*)$.

C Securities Implementation

Having determined the optimal contract as a function of the history of cash flows, we now show how it can be implemented using financial arrangements. We could think of our setting as representing a situation where an entrepreneur (the agent) with limited liability needs outside funds from a lender/creditor (the principal) to finance an investment opportunity.

The optimal contract without monitoring has a natural implementation as a long-term debt contract with cash balance (or credit commitment) as demonstrated by DeMarzo and Sannikov (2006) and Biais et al. (2007). As we discuss below, monitoring results in an inclusion of additional state-contingent short-term financing in such implementation. This short-term financing acts as an insurance mechanism that allows some positive-NPV projects to survive negative cash flows, and there is liquidation in equilibrium (termination, default) only if the liquidation cost is low. As a result, active use of monitoring relaxes the endogenous financing constraint implied by the agency problem, allowing the creditor to provide additional funds to the entrepreneur after negative cash flows without weakening incentives.

In an implementation of the optimal contract with monitoring, the long-term debt contract has the lender deliver a fixed payment at time 0 in exchange for a series of payments in the future. The payment consists of funds needed to start the project as well as additional cash M_0 . The entrepreneur agrees to hold cash (M_t) at account with the lender, which can be used to finance cash-flow shortfalls from the project or the agent's consumption.²⁵ Any outstanding balance on this account acts as additional collateral for the loan.

The lender receives a coupon payment dD_t as specified in the debt contract. The lender also agrees to provide the entrepreneur with additional short-term state-contingent financing. We denote the cumulative short-term financing as SF_t , so $dSF_t > 0$ denotes a transfer from the lender to the cash account, and $dSF_t < 0$ denotes repayment out of the cash account.

Default occurs either if a coupon is missed or if the lender discovers the borrower diverting funds inappropriately. In the event of default at time t , the lender is able to recover the assets of the project with value L . The lender additionally takes ownership of any cash balance the firm has accumulated (M_t). The entrepreneur receives a reservation value of V_R if default happens because of missed payments. If default happens because the entrepreneur was caught diverting the project's resources for his own benefit, the entrepreneur receives V_F (since we restrict our attention to financing arrangements that are incentive compatible, this is always an out-of-equilibrium event).

The total cash-on-reserve balance evolves as

$$dM_t = rM_t dt + dY_t - dC_t - dD_t + dSF_t \quad (\text{C-18})$$

where M_t is the value of the cash balance and r is the interest rate. Over each dt the entrepreneur deposits $dY_t - dC_t$ from the project, using dC_t for consumption. We assume that the entrepreneur's consumption choices are observable but not contractible: the creditor holds the cash balance and so can observe the project's cash flows, so she can infer the agent's consumption choice. This provides the entrepreneur an additional degree of freedom by allowing him to choose the timing of his own consumption, as opposed to the optimal contract (Section 3) in which the principal (creditor) chooses the timing of the agent's (entrepreneur's) consumption. This allows the entrepreneur to potentially over- or under-consume.

This financing arrangement implements the optimal contract as follows:

Proposition 10 *Assume the debt contract is as described above with debt coupons*

$$dD_t = [\mu - \gamma V_R - (\gamma - r)M_t] dt. \quad (\text{C-19})$$

Then it is optimal for the lender to monitor with intensity $m_t = m(M_t + V_R)$, given in

²⁵These assumptions are consistent with empirical work, such as Santikian (2014) documenting the extent and value of cross-selling of banking services to loan customers.

Property 3, and to provide short-term financing

$$dSF_t = -\psi_t m_t (dY_t - \mu dt) \quad (\text{C-20})$$

where $\psi_t = M_t + V_R - V_F$. The entrepreneur does not divert cash from the project and consumes dC_t so that M_t reflects from below at $\bar{M} = V_C - V_R$. The entrepreneur has continuation value $V_t = M_t + V_R$, and the optimal contract is implemented.²⁶

This security design uses default, short-term financing, and monitoring to provide variable levels of incentives to the entrepreneur. Consider first the baseline case in which there is no monitoring or short-term financing. In this case our implementation reduces to the case considered by Biais et al. (2007). The key link is that the entrepreneur's inside stake and the project's financial slack are equivalent ($M_t = V_t - V_R$). Since the entrepreneur can divert project cash flows or overconsume, the debt contract must award the entrepreneur the same value (in the form of continuation utility) to prevent both types of diversion. Since 1) the entrepreneur can consume directly out of the cash reserve, and 2) all project cash flows are added to the cash reserve, then the entrepreneur receives enough from cash flows to incentivize him not to divert funds. Default disciplines the entrepreneur by setting a lower bound on negative performance after which he receives no more consumption. Because the entrepreneur's accumulated inside value – his continuation utility – is equal to the level of cash reserve plus the outside option, the entrepreneur is willing to choose the timing of consumption to maximize project value (and hence implement the optimal contract). Thus, the proposition shows that, as long as the incentive-compatibility conditions on project diversion are met, security design will ensure that incentive-compatibility conditions on overconsumption are met as well.

We now describe how short-term financing and incentives combine. Putting the securities together, the law of motion for M_t , which represents both the cash held on account with the lender and the entrepreneur's inside stake (Proposition 10), is

$$dM_t = \underbrace{\gamma (M_t + V_R) dt}_{\text{Level Effect of Cash Balance}} + \underbrace{(dY_t - \mu dt)}_{\text{Volatility from Project Cash Flows}} - \underbrace{\psi_t m_t (dY_t - \mu dt)}_{\text{Offsetting Volatility from Short-Term Financing}} - dC_t \quad (\text{C-21})$$

The short-term financing has three salient characteristics. First, new financing is proportional to monitoring: just as agency concerns introduce endogenous financing constraints, monitoring is necessary to allow the creditor to provide additional funds in distress without weakening the entrepreneur's incentives. Indeed, the sensitivity of the cash balance to project cash flows is $1 - \psi_t m_t$ instead of 1, and so more negative cash flow innovations are required before default triggers liquidation.

Second, short-term financing is conditional: additional funds are provided when cash

²⁶Proof is available upon request.

flow innovations are negative and withdrawn when cash flow innovations are positive. Thus, financing is provided as the firm enters distress and withdrawn as the firm exits distress. This additional financing acts as insurance, and its provision depends explicitly on the dynamics of cash flows.

Third, the expected discounted value of the short-term financing cash flows is zero. Thus, any risk-neutral party would be willing to provide it for even a nominal fee. Because short-term financing (as insurance) enhances the value of long-term debt, the lender has an explicit reason to provide it.

The entrepreneur must still receive sufficient project value to prevent diversion, but that value can come in the form of expected future consumption or in the form of avoided losses. Monitoring creates a potential loss to the entrepreneur if he is caught diverting. This means that the change in the entrepreneur's future consumption from cash flow innovations can be lower, with more effective monitoring leading to a larger *cash flow wedge*: $\psi_t m_t = 1 - \beta_t$. This cash flow wedge created by monitoring allows investors to offset part of the cash flow innovations: providing funds after a loss and withdrawing them after a gain. Thus, new short-term financing covers a fraction of negative cash flows, with the fraction equal to the effectiveness of monitoring: $\psi_t m_t$.

In sum, monitoring allows relaxation of firm's endogenous financing constraint by making the provision of additional funds in distress possible without eroding entrepreneurial incentives. This provision is value improving because short-term financing coupled with monitoring is a cost-effective way to reduce the risk of default.

The short-term financing can be provided in a number of ways. First, it can be provided directly by a lender who is actively involved in the project's operations: the lender monitors the entrepreneur and provides extra funds when cash flows are negative and withdraws funds when cash flows are positive. Alternatively, loan commitment contracts may serve a similar economic role. Such contracts have state-contingent rates and fees, are conditioned on changing borrower-specific variables, and the amount of credit can be increased or decreased dynamically. Importantly, these arrangements are often used as a facility for managing liquidity, thus acting as insurance mechanisms (see Shockley and Thakor (1997)).

We also note that the debt coupons depend on the project's performance. As in Biais et al. (2007), the coupon decreases with the project performance record. Specifically, debt covenants could state that as cash reserves increase and default risk decreases, the coupon rate is reduced.²⁷

Finally, we note that since the project's cash flows in the implementation match the project's cash flows in the contract and the agent's cash flows are the same in both, the lender's net cash flows must match those of the principal. Thus, the debt holder faces the same optimization problem as the principal, chooses the value-maximizing level of monitoring

²⁷As pointed out by Biais et al. (2007), this property of implementation is consistent with observed features of venture capital contracts and with the increased reliance on step-up bonds or performance pricing loans.

and short-term financing, and receives the same level of expected utility.

To initialize the financing arrangements, transfers must be made from the creditor to the entrepreneur. For simplicity, we can assume that debt markets are competitive, and so the creditor receive no rents. The initial value of the cash balance is M_0 , where M_0 is the largest solution to $I + M_0 = F(M_0) + M_0$, meaning that the total value of the lending relationship to the creditor is exactly equal to the capital necessary to fund the project. If $M_0 > \bar{M}$, the entrepreneur consumes the difference at time 0.

References

- Biais, B., T. Mariotti, G. Plantin, and J.-Ch. Rochet., 2007. Dynamic security design: Convergence to continuous time and asset pricing implications. *Rev. Econ. Stud.* 74: 345-390.
- Biais, B., T. Mariotti, J.-Ch. Rochet, and S. Villeneuve., 2010. Large risks, limited liability and dynamic moral hazard. 2010. *Econometrica* 78: 73-118.
- Brickley, J. A., 2003. Empirical research on CEO turnover and firm-performance: A discussion. *J. Account. Econ.* 36: 227-233.
- Bushman, R. M, Z. Dai, and X. Wang., 2010. Risk and CEO turnover. *J. Financ. Econ.* 96: 381-398.
- Clementi, G. L., and H. A. Hopenhayn., 2006. A theory of financing constraints and firm dynamics. *Q. J. Econ.* 121: 229-265.
- DeMarzo, P. M., and M. J. Fishman., 2007a. Optimal long-term financial contracting. *Rev. Financ. Stud.* 20: 2079-2128.
- DeMarzo, P. M., and M. J. Fishman., 2007b. Agency and optimal investment dynamics. *Rev. Financ. Stud.* 20: 151-188.
- DeMarzo, P. M., and Y. Sannikov., 2006. A continuous-time agency model of optimal contracting and capital structure. *J. Finance* 61: 2681-2724.
- DeMarzo, P. M., M. J. Fishman, Z. He, and N. Wang., 2012. Dynamic agency and the q theory of investment. *J. Finance* 67:2295-2340.
- Diamond, D., 1984. Financial intermediation and delegated monitoring. *Rev. Econ. Stud.* 51: 393-414.
- Dumas, B., 1991. Super contact and related optimality conditions. *J. Econ. Dyn. Control* 15:675-685.
- Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov., 2012. Dynamic CEO compensation. *J. Finance* 67: 1603-1647.
- Edmans, A., and X. Gabaix., 2011. Tractability in incentive contracting. *Rev. Financ. Stud.* 24: 2865-2894.
- Engelbert, H. J., and G. Peskir., 2014. Stochastic differential equations for sticky Brownian motion. *Stochastics* 86: 993-1021.
- Fernandes, N., M. A. Ferreira, P. Matos, and K. J. Murphy., 2010. The pay divide: (Why) are US top executives paid more?. Working paper, USC Marshall.

- Gale, D., and M. Hellwig., 1985. Incentive compatible debt contracts: The one period problem. *Rev. Econ. Stud.* 52: 647-663.
- Gilson S. C., and M. R. Vetsuypens., 1993. CEO compensation in financially distressed firms: An empirical analysis. *J. Finance* 48: 425-458.
- Harrison, J. M., and A. J. Lemoine., 1981. Sticky Brownian motion as the limit of storage processes. *J. Appl. Probab.* 18: 216-226.
- He, Z., 2009. Optimal executive compensation when firm size follows geometric Brownian motion. *Rev. Financ. Stud.* 22: 859-892.
- He, Z., 2012. Dynamic compensation contracts with private savings. *Rev. Financ. Stud.* 25:1494-1549.
- Jacod, J., and A. N. Shiryaev., 2003. Limit theorems for stochastic processes. Springer, 2nd edition.
- Jenter, D., and K. Lewellen., 2014. Performance-induced CEO turnover. Working paper, Stanford GSB.
- Jensen, M. C., and K. J. Murphy., 1990. Performance pay and top-management incentives. *J. Polit. Econ.* 98: 225-64.
- Karatzas, I., and S. E. Shreve., 1991. Brownian motion and stochastic calculus. Springer, 2nd edition.
- Kaplan, S. N., and B. Minton., 2006. How has CEO turnover changed? Increasingly performance sensitive boards and increasingly uneasy CEOs. NBER Working Paper No. 12465.
- Kushner, H. J., and P. Dupuis., 2001. Numerical methods for stochastic control problems in continuous time. Springer, 2nd edition.
- Malenko, A., 2012. Optimal design of internal capital markets. Working paper, MIT Sloan.
- Philippon, T., and Y. Sannikov., 2007. Real options in a dynamic agency model, with applications to financial development, IPOs, and business risk. NBER Working Paper No. 13584.
- Piskorski, T., and A. Tchisty., 2010. Optimal mortgage design. *Rev. Financ. Stud.* 23: 3098-3140.
- Sannikov, Y., 2008. A continuous-time version of the principal-agent problem. *Rev. Econ. Stud.* 75: 957-984.
- Sannikov, Y., 2007. Agency problems, screening and increasing credit lines, Working paper, UC Berkeley.

- Sannikov, Y., and A. Skrzypacz., 2010. The role of information in repeated games with frequent actions. *Econometrica* 78: 847-882.
- Santikian, L., 2014. The ties that bind: Bank relationships and small business lending. *J. Financ. Intermed.* 23:177-213.
- Shockley R., and A. Thakor., 1997. Bank loan commitment contracts: Data, theory, and tests. *J. Money Credit Bank.* 29: 517-534.
- Tchistyi, A., 2013. Security design with correlated hidden cash cows: The optimality of performance pricing. Working paper, UC Berkeley.
- Townsend, R. M., 1979. Optimal contracts and competitive markets with costly state verification. *J. Econ. Theory* 21: 265-293.
- Wang, C., 2005. Dynamic costly state verification. *Econ. Theory* 25: 887-916.
- Williams, N., 2011. Persistent private information. *Econometrica* 79:1233-1275.
- Weisbach, M. S., 1988. Outside directors and CEO turnover. *J. Financ. Econ.* 20: 431-460.
- Zhang, Y., 2009. Dynamic contracting with persistent shocks. *J. Econ. Theory* 144: 635-675.
- Zhu, J., 2013. Optimal contracting with shirking. *Rev. Econ. Stud.*, 80:812-839.