Optimal Contracts with Hidden Risk*

Rui Li
University of Massachusetts Boston

Noah Williams†
University of Wisconsin-Madison

August 17, 2018

Abstract

Several episodes in recent years have highlighted the problem of managers subjecting their firms to large risks. We develop a dynamic moral hazard model where a manager's diversion of funds is indistinguishable from random shocks. In addition, the manager takes unobservable actions which yield certain current payoffs, but expose the firm to large negative shocks. We show that standard pay-for-performance contracts, which are typically beneficial under moral hazard, may lead the manager to take on excess risk. We then characterize the optimal contract taking into account incentive provision and risk management. We solve two examples. One is explicitly solvable and in it the contract can be implemented with simple instruments. The owner gives the manager a constant salary payment and allows him to manage the ex-dividend assets of the firm, but imposes a "clawback" fee in the event of large negative shocks. The second example shows that it may sometimes be optimal for the owner to forgo risk management and allow the manager to take excess risk.

1 Introduction

"I accumulated this trading position and concealed it for the purpose of augmenting my reputation at Goldman and increasing my performance-based compensation."

Matthew Taylor, ex-trader for Goldman Sachs convicted of wire fraud, quoted in Bray and Baer [2013]

*Financial support from the National Science Foundation is gratefully acknowledged.

†Contact: `nmwilliams@wisc.edu`

In December 2007, after having lost most of his accumulated profits for the year and facing a significant reduction in his annual bonus, Matthew Taylor of Goldman Sachs placed an \$8.3 billion bet on S&P 500 futures. To conceal the trade from his supervisors, Mr. Taylor made false entries that appeared to take the opposite side of the position. The hidden trade went bad, and cost Goldman \$118.4 million to unwind (Bray and Baer [2013]). Mr. Taylor was not alone in this practice, as there have been many other episodes in recent years where managers and traders have engaged in actions that have subjected their firms to large potential losses. As emphasized by Rajan [2011], one prominent aspect of the 2008 financial crisis was that firms like AIG sold billions of dollars credit default swaps on mortgage bonds and other asset backed securities. While the housing market was healthy, the firms pocketed the fees from the sale of the swaps, but once the downturn began they incurred heavy losses. Although some of this activity was observable and known, at least part of it was unobservable and hidden from shareholders, regulators, and even upper-level management. As Rajan [2011] states, "After all the profits from such [risk taking] activities would look a lot healthier if no one knew the risks they were taking. Accordingly, Citibank's off-balance sheet conduits, holding an enormous quantity of asset-backed securities funded with short-term debt were hidden from all but the most careful analysis." Both regulators and shareholders were surprised once the actual extent of liabilities was uncovered.

In this paper we study the incentive to manage risk, and show how it interacts with performance-based contracts. Classic moral hazard problems in economics and corporate finance generally focus on hidden actions: managers can shirk and not put forth effort, or they may be able to divert resources from the firm for their own private benefit. Such agency frictions motivate performance-based pay contracts, which align the incentives of owners and managers. However in situations like those described above, managers may be able to take on excess risk which is unobservable to firm owners, for example by loading up on tail risk by making futures bets or selling insurance against unlikely events. The firm owners would observe the increased cash flows from the insurance or risk premiums, which they would interpret as good firm performance, worthy of reward for the managers. But the owners would be unaware of the scale of potential future losses. We develop a model which incorporates this tension between performance pay and risk taking. We show how owners can provide incentives to discourage excess risk, but also that it may sometimes be optimal for owners to forgo risk management and let managers gamble the firm's assets.

In our model an owner hires a manger to manage her firm's assets over an infinite time horizon, and both parties are risk averse. Moral hazard arises because the manager could divert the firm's resources for his own consumption, and this behavior is indistinguishable from random cash-flow shocks. In addition, the manager can take unobservable actions which yield risk-free current payoffs, but expose the firm's assets to large risks. We model this as the manager controlling the arrival rate of a jump process which has a negative impact on the firm's capital when it occurs. The firm's overall exposure to the shock has a zero mean, so the occasional large negative shock is offset by a positive flow income.[1]

Suppose that the manager were to increase the arrival rate of the shock, making large losses more likely, but this was unobservable to shareholders. Prior to arrival of a negative shock, the shareholders would observe the cash inflow resulting from the risk compensation. However the shareholders would attribute these flows to good firm performance, and thus would reward the manager. We illustrate the hazards of pay-for-performance contracts in the presence of hidden risk, a channel which was emphasized by Rajan [2011]. A contract which is optimal for the standard moral hazard friction of preventing diversion, but ignores the risk-taking incentives, will closely tie the manager's compensation to the firm's performance, and so may induce him to take on excess risk.

An optimal contract must balance the incentives to alleviate the moral hazard friction and to manage risk, and we show that in some cases the owner optimally gives up on risk management. It may be too costly for the owner to provide sufficient incentives for the manager to take prudent actions, so instead the owner focuses on preventing resource diversion. Rather than being a failure of corporate management, it thus may be a feature of optimal contracts that managers subject their firms to excess risks, an undesirable outcome in the absence of private information frictions.

After laying out the general model, we solve for the optimal contracts with moral hazard and risk management under two different parametric specifications. In the first case, we assume the owner and the manager have exponential preferences, which allows us to obtain explicit solutions for the optimal contracts. The risk management policy is time invariant in this case, with the contract implementing either high or low risk for all time, dependent on the parameters. We show that high risk is optimal if the firm's growth prospects are low, or if the owner is sufficiently less risk averse than the manager. We also show how the optimal contract can be implemented with some simple instruments. In particular, the owner gives the manager a constant salary payment and access to an account which represents the cumulation of his bonus or deferred compensation, and the balance on this account increases with the firm's assets. To provide incentives to manage risk, the contract also features a "clawback" provision in which the manager loses part of his deferred compensation in the event of a negative jump shock. Similar clawback provisions have been proposed by federal regulators and have been implemented in some financial institutions.

While exponential preferences are useful in obtaining explicit solutions, they imply con-

[1]Selling insurance would be more directly modeled as controlling the impact of a shock, rather than the arrival rate. However once a shock hits, the manager's actions would be revealed. Our formulation is mathematically similar, but preserves private information.

stant risk management incentives. Our second illustration studies risk dynamics in a tractable environment, assuming that the owner and the manager have logarithmic preferences. An important determinant of the optimal risk management policy in this case is the manager's "share" in the firm, which is roughly equal to the ratio of the net present value of the manager's compensation to the firm's total future cash flows. This share is the key endogenous variable whose evolution governs the dynamics of the contract. As his share increases, the manager's interests are more aligned with the firm's, the pay-performance sensitivity required to deter diversion decreases, and the tension between incentive provision and risk management is weakened. We show that the contract is divided into three regions. When the manager's share is low, the owner gives up risk management as the manager's incentive to divert resources is strong. For intermediate share levels, the contract balances incentive provision and risk management, and the manager is induced to take the low risk. When the share is sufficiently high, there is no tension between the two objectives, as with the appropriate levels of pay-performance sensitivity, the manager voluntarily chooses low risk. Under an optimal contract, the manager's share tends to decrease over time, pushing the firm into the high risk region. Smaller firms also tend to have higher risk, as the manager's share increases with the firm's assets.

In the previous literature, the most closely related paper is DeMarzo et al. [2014] which also studies a dynamic agency model with hidden cash-flow diversion and hidden risk taking. While our motivation and general approach is similar, there are some important differences. In their model, the firm is liquidated and the contract is terminated exogenously when a negative jump shock hits. By contrast, in our model jumps are recurrent. This allows us to study the endogenous consequences of negative jumps, with the punishments upon their occurrence being chosen as part of the contract. As we noted, this has important implications for the design of the "clawback" features of contracts, like those used in practice. In addition, in their model both the principal and the agent are risk neutral, as they are in the related papers by Rochet and Roger [2016] and Wong [2018].[2] By contrast, we consider potentially differing levels of risk aversion, and thus rationalize ongoing dividends and a salary payment as part of a pay package. Our model allows us to study the effect of risk preferences on the optimal risk taking policy, and also allows the accumulation of the assets which have persistent effects on firms' cash flows. Therefore our framework help us understand the interaction between firm dynamics and risk management.

More broadly, our paper is related to several strands of research. In static environments, several papers study risk-taking behavior with information frictions, including Jensen and

[2] Wong [2018] also highlights that excess risk may be optimal when the manager's stake in the firm is sufficiently small, which is similar to our results. Feng and Westerfield [2018] consider a related model where the agent has control over variability of cash flows in addition to moral hazard.

Meckling [1976], Diamond [1998], Palomino [2003], and Biais and Casamatta [1999]. An advantage of our dynamic model is the ability to study the time series properties of the optimal risk management policy and the implied firm dynamics, showing how incentives and outcomes under the contract evolve over time. In a dynamic setting, Biais et al. [2010] develop a similar model where a risk-neutral manager can take unobservable effort to prevent large losses. In their model, the moral hazard problem is one-dimensional, focusing on the manager's hidden effort, and thus doesn't feature the tension between incentives which is our focus. In terms of technique, this paper belongs to the growing literature on continuous time dynamic contract design which includes Sannikov [2008], DeMarzo and Sannikov [2006], Williams [2011, 2008], and our previous paper Li and Williams [2014], as well as those cited above. By casting the model in continuous time, we utilize powerful tools in stochastic control to characterize incentive compatible contracts in a relatively complex environment.

The remainder of the paper is arranged as follows: in Section 2 we introduce the basic setup; in Section 3 we characterize the optimal contracts in the general case; in Section 4 we consider the case exponential preferences which allow us to derive an explicit solution of the optimal contract and its implementation; in Section 5 we focus on the case with logarithmic preferences and risk management dynamics; Section 6 concludes.

2 The Model

At date zero, a firm's owner (the principal) hires a manager to oversee the assets y_t of the firm. Out of the asset stock, the owner pays himself a dividend d_t which he consumes, and he pays the manager c_t. The manager can secretly adjust his compensation by Δ_t by diverting the firm's assets for his own consumption. The assets y_t yield pre-consumption expected gains of $\mu(y_t)$ and are subject to continuous stochastic shocks following a Brownian motion W_t. In addition, the firm's assets are subject to occasional large negative shocks, which we depict using a jump process with arrival rate λ_t. To ensure that the jump shocks are mean zero, we assume that they are compensated at a rate λ_t each instant. We define the counting process $\{N_t\}$ indicating the arrivals of the shocks so that $dN_t = 1$ if a shock hits at t, and is zero otherwise. For simplicity we assume all large shocks are of unit size, and when a shock hits a deterministic amount $\phi(y_t)$ of assets are destroyed. We assume that the manager can control the arrival rate λ_t, and thus the shocks can represent potentially risky projects that the manager can undertake. The firm collects the premium $\lambda_t\phi(y_t)$ each instant, at the cost of possible losses of $\phi(y_t)$.[3] For simplicity we assume that the risk level λ_t can take on two

[3]Once a jump shock hits, the size of the loss is observable, so potential moral hazard in controlling losses could be eliminated with sufficient punishment. Fixing the loss sizes but controlling their likelihood preserves the private information.

values $\underline{\lambda} < \overline{\lambda}$. Thus the firm's assets follow:

$$dy_t = [\mu(y_t) - d_t - c_t - \Delta_t]dt + \sigma(y_t)dW_t - \phi(y_t)\left(-\lambda_t dt + dN_t\right).$$

The term $-\lambda_t dt + dN_t$ gives the impact of the compensated jump shocks, and is an increment of a martingale.

The level of the assets y_t and the cumulative realizations of the large shocks N_t are publicly observable, but the manager's diversion Δ_t and chosen risk level λ_t are private information. Thus moral hazard arises and the owner needs to offer a contract which induces the manager to not divert and to take appropriate risk. Much of the contracting literature focuses on the diversion friction, which is also similar to hidden effort. We show how this standard friction interacts with the hidden risk choice. At date zero, the owner offers the manager a contract which, for each date t, pays him c_t, provides a dividend d_t, suggests that $\Delta_t \equiv 0$, and suggests a risk level λ_t. Given a contract, the manager solves:

$$\max_{\hat{a}\in A} E^{\hat{a},p}\left[\int_0^\infty e^{-\rho t}u(c_t + \hat{\Delta}_t)dt\right]. \tag{1}$$

Here, u is the manager's utility function, $\rho > 0$ is the rate of time preference, $p_t = (d_t, c_t)$ is the principal's choice vector with p the entire process $\{p_t\}$, and $\hat{a}_t = (\hat{\Delta}_t, \hat{\lambda}_t)$ is the manager's choice vector with $\hat{a}$ the entire process $\{\hat{a}_t\}$. A is the set of all feasible choice processes $\hat{a}$ and $E^{\hat{a},p}$ is the expectation operator induced by the processes $\hat{a}$ and p. Let $\{\mathcal{F}_t\}_{t\in[0,\infty)}$ be the suitable augmentation of the filtration generated by the processes $\{y_t, N_t\}$, which represents the public history. We require that $\hat{a}_t$ be $\mathcal{F}_t$–predictable, meaning that the manager makes his decisions about Δ_t and λ_t before the realization of the shocks at t.

Given the monitoring structure of our model, the manager is able to adjust his consumption from the suggested level to any level by diverting the firm's resources. So, for any contract inducing the manager's choice $\hat{a} = \left(\hat{\Delta}, \hat{\lambda}\right)$ and paying a compensation plan $\hat{c}$, we could define an equivalent contract without diversion by choosing compensation plan c such that $c_t = \hat{c}_t + \hat{\Delta}_t$ and the manager's choice a such that $a_t = \left(0, \hat{\lambda}_t\right)$. Hence, we can focus on the contracts under which $\Delta_t = 0$ for all t without loss of generality. Therefore, in our model, a contract is *incentive compatible* if $\hat{a} = a = (0, \lambda)$ is indeed the manager's optimal choice.

Given the owner's policy p and manager's action a, we define the manager's promised utility as:

$$q_t \equiv E^{a,p}\left[\int_t^\infty e^{-\rho t}u(c_t + \Delta_t)dt|\mathcal{F}_t\right] \text{ for } t \in [0,\infty].$$

As in Sannikov [2008] and Li and Williams [2014], a martingale representation theorem gives the following law of motion of q.

Proposition 1 *Given the policy p and action a, there exist two $\{\mathcal{F}_t\}$-predictable and square-integrable process, $\{\gamma_t^W, \gamma_t^N\}$ satisfying:*

$$dq_t = [\rho q_t - u(c_t + \Delta_t)]dt + \gamma_t^W \sigma(y_t) dW_t^a - \gamma_t^N dM_t^\lambda \tag{2}$$

where:

$$\begin{aligned} W_t^a &\equiv \int_0^t \frac{1}{\sigma(y_s)} \left[dy_s - (\mu(y_s) - d_s - c_s - \Delta_s + \phi(y_s)\lambda_s) ds - \phi(y_s) dN_s\right] \\ M_s^\lambda &\equiv \int_0^t \left[-\lambda_s ds + dN_s\right], \end{aligned}$$

which are martingales under the policies a and p.

Proof. See Appendix A. ■

Here γ_t^W denotes the sensitivity of promised utility to the Brownian motion shocks and γ_t^N measures the sensitivity to the jumps. To gain insight on manager's incentives, note that we can re-write (2) as:

$$\begin{aligned} dq_t &= [\rho q_t - u(c_t + \Delta_t)]dt + \gamma_t^W \left[dy_t - (\mu(y_t) - d_t - c_t - \Delta_t + \phi(y_t)\lambda_t) dt - \phi(y_t) dN_t\right] \\ &\quad - \gamma_t^N [-\lambda_t dt + dN_t]. \end{aligned}$$

The following proposition shows that incentive compatible contracts are characterized by two local optimality conditions, which come from optimizing the drift of promised utility with respect to Δ_t and λ_t, and evaluating the result at $\Delta_t = 0$. This is simply another way of stating that given the contract, the incentive constraints capture the agent's instantaneous first order conditions.

Proposition 2 [Incentive compatibility] *A contract (a, p) is incentive compatible if and only if for any $t \in [0, \infty)$:*

$$u'(c_t) = \gamma_t^W, \tag{3}$$

and:

$$\lambda_t = \begin{cases} \underline{\lambda} & \text{if } \phi(y_t)\gamma_t^W \leq \gamma_t^N \\ \bar{\lambda} & \text{if } \phi(y_t)\gamma_t^W > \gamma_t^N. \end{cases} \tag{4}$$

Proof. See Appendix B. ■

The first incentive constraint, which is analogous to those in Sannikov [2008], Williams [2008, 2011], links the manager's consumption to variations in his promised utility. Since $u' > 0$ we have $\gamma_t^W > 0$, so promised utility increases with the Brownian increments dW_t. This is the usual pay-for-performance channel to provide incentives under moral hazard. The second incentive constraint is similar to that in Li and Williams [2014], and determines the choice of the risk level. Notice that γ_t^N is the size of the punishment in utility terms when a negative jump shock is realized. While choosing $\bar{\lambda}$ increases the frequency of such a punishment, it also makes the manager's performance appear better because the jump risk is compensated. Thus the incentive constraint (4) captures the tradeoff in the risk choice between greater asset flows due to the jump compensation and more frequent punishment with negative shocks.

The more closely the manager's pay is tied to firm performance through γ_t^W, the greater will be the manager's gains from them jump-risk compensation, and so the stronger would be his incentives to take riskier actions. This helps explain why standard performance-based contracts, which do not account for the manager's control over risk, provide incentives for riskier actions. Incentive compatibility thus requires that greater pay-performance sensitivity must be accompanied by a larger punishment when a jump shock hits.

3 Optimal Contracts

We now turn to the owner's choice of an optimal contract. We suppose that the owner is risk averse, with instantaneous utility function $v(d)$ and the same discount rate ρ as the agent. Thus she chooses a contract (a, p) to maximize:

$$\max_{a,p} E^{a,p}\left[\int_0^\infty e^{-\rho t} v(d_t)dt\right], \tag{5}$$

subject to delivering a specified level q_0 of utility to the manager (a participation constraint), and to the contract (a, p) being incentive compatible.

We successively consider three different information structures. First, we study full information, where the manager's consumption and risk choices are observable. This provides an efficient benchmark to evaluate the other cases. Then we consider moral hazard, where the manager's consumption is unobservable, but the owner views the risk level as fixed and not under the control of the manager. This isolates the effect of moral hazard in diversion, a friction commonly studied in the contracting literature. Here we also show that if the owner does not account for the manager's risk choice, then the manager may take on excess risk. Finally, we turn to moral hazard in diversion and risk management.

3.1 Full Information

Under full information, the owner can directly control diversion and risk, so we only require that the contract deliver the utility promise q_0. By varying this level we can trace out the Pareto frontier between the owner and the manager. Thus the owner's problem is:

$$\max_{\{c_t, d_t, \lambda_t, \gamma_t^W, \gamma_t^N\}} E^{a,p}\left[\int_0^\infty e^{-\rho t} v(d_t)dt\right],$$

with (q_0, y_0) given and:

$$\begin{aligned} dy_t &= [\mu(y_t) - d_t - c_t + \lambda_t \phi(y_t)]dt + \sigma(y_t)dW_t^a - \phi(y_t)dN_t, \\ dq_t &= \left[\rho q_t - u(c_t) + \lambda_t \gamma_t^N\right] dt + \gamma_t^W \sigma(y_t) dW_t^a - \gamma_t^N dN_t. \end{aligned}$$

We denote the owner's value function J^F under full information, J^M under moral hazard in diversion, and J^R under moral hazard in effort and risk. The owner's Hamilton-Jacobi-Bellman (HJB) equation under full information can be written:

$$\rho J^F(y,q) = \max_{c,d,\lambda,\gamma^N,\gamma^W} \Bigg\{ v(d) + J_y^F(y,q)\left[\mu(y) - c - d + \lambda\phi(y)\right] + J_q^F(y,q)\left[\rho q - u(c) + \lambda\gamma^N\right]$$

$$+ \frac{1}{2}\sigma(y)^2\left[J_{yy}^F(y,q) + 2\gamma^W J_{yq}^F(y,q) + (\gamma^W)^2 J_{qq}^F(y,q)\right] + \lambda[J^F(y-\phi(y), q-\gamma^N) - J^F(y,q)]\Bigg\}$$

Note that the objective is linear in λ and the terms multiplying it are:

$$J_y^F(y,q)\phi(y) + J_q^F(y,q)\gamma^N + J^F(y-\phi(y), q-\gamma^N) - J^F(y,q).$$

This sum is negative when J^F is concave, in which case the owner will want to implement the lower risk level, $\lambda = \underline{\lambda}$. The first-order conditions for the other choice variables are:

$$\begin{aligned} v'(d) &= J_y^F(y,q), \\ J_q^F(y,q)u'(c) &= -J_y^F(y,q), \\ \gamma^W &= -J_{yq}^F(y,q)/J_{qq}^F(y,q), \\ J_q^F(y,q) &= J_q^F(y-\phi(y), q-\gamma^N). \end{aligned}$$

Solving these equations and the HJB equation gives the optimal contract. Below we show how to do so under two different parametric specifications.

To solve the PDE comprising the HJB equation, we need boundary conditions which are provided by Pareto frontier considerations. First, note that for any level of assets y there is a maximal amount of promised utility $\bar{q}(y)$ that the owner can provide the manager. This can be found by setting $d_t \equiv 0$ and maximizing the manager's utility:

$$\rho\bar{q}(y) \;=\; \max_c \left\{ u(c) + \bar{q}'(y)\left[\mu(y) - c + \underline{\lambda}\phi(y)\right] + \frac{1}{2}\sigma(y)^2\bar{q}''(y) + \underline{\lambda}[\bar{q}(y-\phi(y)) - \bar{q}(y)] \right\}.$$

This provides the owner with his minimal value $\underline{J} = v(0)/\rho$, so we have the boundary condition $J^F(y, \bar{q}(y)) = \underline{J}$. Similarly, if the owner gave no consumption to the manager, he would deliver the minimal utility $\underline{q} = u(0)/\rho$ and obtain the value $\bar{J}(y)$ which solves:

$$\rho\bar{J}(y) = \max_d \left\{ v(d) + \bar{J}'(y)\left[\mu(y) - d + \underline{\lambda}\phi(y)\right] + \frac{1}{2}\sigma(y)^2\bar{J}''(y) + \underline{\lambda}[\bar{J}(y-\phi(y)) - \bar{J}(y)] \right\}.$$

Thus we have the other boundary condition $J^F(y, \underline{q}) = \bar{J}(y)$.

3.2 Moral Hazard with Uncontrolled Risk

Now we suppose that the owner cannot observe the manager's consumption and treats the risk level as fixed at λ, outside the control of the manager. This is a standard moral hazard problem, where by structuring payments appropriately the owner can ensure that the manager does not divert resources. This requires that the owner tie the manager's compensation to his promised utility (and the firm's assets), via the incentive constraint (3). Imposing this, the owner's HJB equation can be written:

$$\begin{aligned}\rho J^M(y,q) = \max_{c,d,\gamma^N} \quad & \left\{v(d) + J_y^M(y,q)[\mu(y) - c - d + \lambda\phi(y)] + J_q^M(y,q)\left[\rho q - u(c) + \lambda\gamma^N\right]\right. \\ & +\frac{1}{2}\sigma(y)^2\left[J_{yy}^M(y,q) + 2u'(c)J_{yq}^M(y,q) + (u'(c))^2 J_{qq}^M(y,q)\right] \\ & \left.+\lambda[J^M(y-\phi(y), q-\gamma^N) - J^M(y,q)]\right\}. \end{aligned} \tag{6}$$

The first-order conditions are now:

$$\begin{aligned} v'(d) &= J_y^M(y,q), \qquad (7) \\ J_q^M(y,q)u'(c) &= -J_y^M(y,q) + \sigma(y)^2 u''(c)[J_{yq}^M(y,q) + u'(c)J_{qq}^M(y,q)] \\ J_q^M(y,q) &= J_q^M(y-\phi(y), q-\gamma^N). \end{aligned}$$

Relative to full information, we now see that private information makes the asset volatility $\sigma(y)$ affect the manager's consumption. This reflects the need to tie pay to performance to provide incentives. Furthermore, even though the jump shock is observable and the owner does not consider the manager's risk taking incentives, under the optimal contract the manager's consumption will still respond to the jumps. This reflect both risk sharing and "paying for luck" as in Li [2016]. The contract responds to observable shocks to lessen the cost of providing incentives and to share risk between the risk averse manager and owner. We also have one of the same boundary conditions as above, since when the manager effectively owns the assets and attains his maximal utility, the private information is of no consequence. Thus $J^M(y,\bar{q}(y)) = \underline{J}$. The other boundary condition no longer holds, as the manager's maximal utility is not attainable with an incentive compatible contract.

Given the choice, the manager may now choose excess risk. Suppose that the owner treats risk as being fixed at $\underline{\lambda}$ and designs a contract taking into account the moral hazard friction. But the owner does not observe the actual level of risk chosen by the manager. Facing this contract, the manager will in fact choose the high risk $\overline{\lambda}$ if:

$$\phi(y)\gamma^W - \gamma^N = \phi(y)u'(c) - \gamma^N > 0. \tag{8}$$

If the contract does not sufficiently punish the manager be cutting his utility promise when the jump shock hits, with $\gamma^N \leq \phi(y)u'(c)$, then the manager will take on more risk than the contract recommends and will choose $\lambda = \overline{\lambda}$. This is a bad outcome for the owner, whose preferences over risk are the same as under full information, as there is no interaction between the manager's consumption and the risk level. Thus the owner (typically) prefers low risk, but the manager may have an incentive to choose high risk.

There is also a tension between the pay-for-performance aspect of the contract and the incentive to take on risk. In order to provide incentives for the manager to not divert resources, the owner makes the manager's promised utility co-vary positively with the firm's performance as $\gamma^W = u'(c) > 0$. But a larger γ^W makes (8) more likely to hold, so the manager will choose the high risk. Greater rewards for good performance give the manager the incentive to increase the risk that the firm faces. The risk is compensated, so it is a gamble of assets with zero expected return. But the owner will view the risk compensation as reflecting better performance of the firm, and so will reward the manager. This is similar to the dynamic described by Rajan [2011], where prior to the financial crisis pay-for-performance contracts exacerbated the incentive for managers to load up on tail risk. The immediate risk compensation meant higher managerial rewards, but increased the risk of large firm losses. We next describe how to manage such risks.

3.3 Moral Hazard in Risk Management

We next suppose that the owner cannot observe the manager's consumption, but now he recognizes that the risk level λ is a hidden action of the manager. Above we showed that under full information it was always optimal for the owner (as long as his value function was concave) to implement the low risk level. However that need not be true with private information, as the owner must trade off the costs of providing incentives to discourage diversion with those to manage risk. Even though the owner would prefer lower risk if he could impose it directly, private information may make it too costly, and so he may let the manager choose the higher risk level.

Based on these considerations, we split the state space into a region $\mathcal{L}$ where the low-risk choice $\underline{\lambda}$ is implemented and $\mathcal{H}$ where $\overline{\lambda}$ is. For $(y, q) \in \mathcal{H}$, the owner's value function J^R satisfies the same HJB equation (6) and related optimality conditions with $\lambda = \overline{\lambda}$. For $(y, q) \in \mathcal{L}$, the owner must satisfy the second incentive constraint in (4) above. We now (tentatively) assume that this constraint binds for $\underline{\lambda}$, which implies:

$$\gamma^N = \phi(y)\gamma^W = \phi(y)u'(c).$$

Thus the choice of γ^N is now constrained and the owner's HJB equation can be written:

$$\begin{aligned}\rho J^R(y,q) = \max_{c,d} \quad & \big\{ v(d) + J^R_y(y,q)[\mu(y) - c - d + \underline{\lambda}\phi(y)] + J^R_q(y,q)\left[\rho q - u(c) + \underline{\lambda}\phi(y)u'(c)\right] + \\ & + \frac{1}{2}\sigma(y)^2\left[J^R_{yy}(y,q) + 2u'(c)J^R_{yq}(y,q) + (u'(c))^2 J^R_{qq}(y,q)\right] \\ & + \underline{\lambda}[J^R(y-\phi(y), q-\phi(y)u'(c)) - J^R(y,q)]\big\} \end{aligned} \tag{9}$$

The first-order conditions are now:

$$\begin{aligned} v'(d) &= J^R_y(y,q), \\ J^R_q(y,q)u'(c) &= -J^R_y(y,q) + \sigma(y)^2 u''(c)[J^R_{yq}(y,q) + u'(c)J^R_{qq}(y,q)] \\ &\quad + \underline{\lambda}\phi(y)u''(c)[J^R_q(y,q) - J^R_q(y-\phi(y), q-\phi(y)u'(c))]. \end{aligned} \tag{10}$$

Relative to uncontrolled risk, the risk management incentives now make the risk level λ affect the manager's consumption. Because of this, the owner's objective is no longer linear in λ, and simple concavity considerations do not determine the owner's optimal risk choice. Instead the risk level influences the manager's consumption, which through the moral hazard channel affects the volatility of his promised utility. This additional risk dependence may increase the owner's costs, and in some regions of the state space may lead the owner to prefer allowing the manager to choose the high risk level. When it becomes too costly to deter both diversion and risk taking, the owner may let the manager gamble.

In addition to the owner's preferences, we need to consider the manager's incentives. Clearly if the manager has no incentive to take excess risk, then the incentive constraint is slack and the owner need not alter the moral hazard contract with risk fixed at the low level. Therefore we define the region $\mathcal{R}$ as that in which facing the contract which treats the risk level as fixed at $\underline{\lambda}$ the manager would have an incentive to choose the risky action $\overline{\lambda}$, and $\mathcal{S}$ where instead the manager would choose the safe action $\underline{\lambda}$. That is, $\mathcal{R}$ is the region where as in (8) the manager's incentive constraint is violated for the contract which presumes that the risk is fixed at a low level. Then over the state space we have: $J^R(y,q)$ satisfies (6) with $\underline{\lambda}$ for $(y,q) \in \mathcal{S}$; satisfies (9) for $(y,q) \in \mathcal{L} \cap \mathcal{R}$; and satisfies (6) with $\bar{\lambda}$ for $(y,q) \in \mathcal{H} \cap \mathcal{R}$.

4 An Explicitly Solvable Case

We now turn to our first example. In this section we explicitly solve for the optimal contract when both the owner and the manager have exponential preferences, and the assets follow a linear evolution with constant diffusion and constant losses when the jump shock hits. That is, we assume the following functional forms:

$$u(c) = -\exp(-\theta_A c), \;\; v(d) = -\exp(-\theta_P d), \quad \mu(y) = \mu y, \quad \sigma(y) = \sigma, \quad \phi(y) = \phi.$$

Our results here are similar to those in Williams [2011] and Li and Williams [2014]. Under this specification, for each information structure the owner's value function has the form:

$$J^i(y,q) = J^i \exp(-\theta_P \mu y)(-q)^{-\theta_P/\theta_A}, \quad i = F, M, R,$$

where J^i is a constant which differs in each case. The policy functions also have the same form in each case. Note that $J^i(y,q)$ is concave, so absent risk management considerations, the owner will implement the low risk level.

4.1 Full Information

Using the functional forms and $J^F(y,q)$, the optimal policies under full information are:

$$\begin{aligned}
d^F &= -\frac{\log(-\mu J^F)}{\theta_P} + \mu y + \frac{1}{\theta_A}\log(-q) \\
c^F &= -\frac{1}{\theta_A}\log(-\mu q) \\
\gamma^{W,F} &= -\frac{\theta_A \theta_P}{\theta_A + \theta_P}\mu q \equiv \bar{\gamma}^{W,F} q \\
\gamma^{N,F} &= \left[1 - \exp\left(\frac{\theta_A \theta_P}{\theta_A + \theta_P}\mu \phi\right)\right] q \equiv \bar{\gamma}^{N,F} q,
\end{aligned}$$

where we use the superscript F to distinguish these policies from those below. The consumption of both the owner and manager are linear in the log of promised utility, and both the utility adjustment terms are proportional to the level of promised utility. These utility terms here represent risk-sharing, so they naturally depend on both risk aversion levels θ_A and θ_P. From the HJB equation, the constant term J^F solves:

$$\begin{aligned}
\mu \log\left(-\mu J^F\right) &= \mu - \rho + \frac{\theta_P}{\theta_A}\left(-\rho + \mu - \mu \log \mu - \underline{\lambda}\bar{\gamma}^{N,F}\right) - \theta_P \mu \underline{\lambda}\phi \\
&\quad + \frac{1}{2}\sigma^2 \theta_P^2 \mu^2 \frac{\theta_A}{\theta_A + \theta_P} + \underline{\lambda}\left[\exp(\theta_P \mu \phi)\left(1 - \bar{\gamma}^{N,F}\right)^{-\theta_P/\theta_A} - 1\right].
\end{aligned}$$

Under the contract, the firm's assets y_t and the manager's promised utility q_t follow jump-diffusion processes with constant coefficients (arithmetic and geometric, respectively):

$$\begin{aligned}
dy_t &= \left[\frac{\log(-\mu J^F)}{\theta_P} + \underline{\lambda}\phi\right] dt + \sigma dW_t - \phi dN_t, \\
dq_t / q_t &= \left[\rho - \mu + \underline{\lambda}\bar{\gamma}^{N,F}\right] dt + \sigma \bar{\gamma}^{W,F} dW_t - \bar{\gamma}^{N,F} dN_t.
\end{aligned}$$

The optimal policies are similar in the cases below, resulting in a similar asset and promised utility evolution as well.

4.2 Moral Hazard with Uncontrolled Risk

Under moral hazard with uncontrolled risk, the optimal policy functions for d and γ^N are the same as under full information, apart from changing the constant in d to J^M:

$$\begin{aligned} d^M &= -\frac{\log(-\mu J^M)}{\theta_P} + \mu y + \frac{1}{\theta_A}\log(-q) \\ \gamma^{N,M} &= \left[1 - \exp\left(\frac{\theta_A \theta_P}{\theta_A + \theta_P}\mu\phi\right)\right] q \equiv \bar{\gamma}^{N,M} q \end{aligned}$$

To find the optimal policy for the agent's consumption c^M, it is useful to define $\bar{c}^M$ by:

$$-\exp(-\theta_A c^M) = \bar{c}^M q.$$

Thus $\bar{c}^M$ gives the proportion of promised utility q that is delivered by current consumption c^M. Since $q < 0$ increases in the consumption rate $\bar{c}^M$ are associated with reductions in the level of consumption c^M. Under full information the consumption rate is μ, but under moral hazard $\bar{c}^M$ is determined implicitly by:

$$\bar{c}^M = \mu + \sigma^2 \theta_A \bar{c}^M (\theta_P \mu - (\theta_P + \theta_A)\bar{c}^M), \tag{11}$$

and we have the following result.

Lemma 1 *There exists a unique optimal $\bar{c}^M$ which is strictly positive and is given by:*

$$\bar{c}^M = \frac{-1 + \sigma^2 \theta_A \theta_P \mu + \sqrt{(\sigma^2 \theta_A \theta_P \mu - 1)^2 + 4\mu\sigma^2\theta_A(\theta_P + \theta_A)}}{2\sigma^2 \theta_A (\theta_P + \theta_A)}$$

In addition $\gamma^{W,M} \geq \gamma^{W,F}$.

Proof. See Appendix C. ■

From (11), notice that as $\sigma \to 0$ the information friction vanishes, and $\bar{c}^M \to \mu$ as under full information. However moral hazard adds two additional terms in (11) which are proportional to the asset volatility. As $\theta_P \to 0$ the owner becomes closer to risk neutral and the first of these term vanishes, so we're left with:

$$\bar{c}^M = \mu - \sigma^2 \theta_A^2 \left(\bar{c}^M\right)^2.$$

To provide incentives, the manager must face a more volatile utility process. Thus to deliver a given utility promise q requires greater consumption (lower $\bar{c}^M$) as compensation. The manager's risk aversion θ_A governs how much his consumption rate is reduced with private information. For $\theta_P > 0$, the additional terms in (11) capture the co-insurance component of the contract, linking the utilities of the owner and the manager.

From the incentive constraint, we have:

$$\gamma^{W,M} = -\theta_A \bar{c}^M q,$$

while under full information we have:

$$\gamma^{W,F} = -\theta_A \frac{\theta_P}{\theta_A + \theta_P} \mu q.$$

So under moral hazard, the risk sharing term $\frac{\theta_P}{\theta_A+\theta_P}$ in the utility adjustment vanishes, as the incentive constraint ties the manager's promised utility to only his own consumption. This pushes toward increasing the sensitivity $\gamma^{W,M}$ of promised utility to shocks under moral hazard. However, the agent's consumption rate is also reduced to $\bar{c}^M$ rather than μ, which tends to work in the opposite direction. Lemma 1 shows that the first term dominates, so that the agent's promised utility is more responsive to shocks under moral hazard than under full information. This is as we would expect: in order to provide incentives to not divert assets for consumption, the owner must make the manager's promised utility more responsive to firm performance. Finally, the constant term J^M solves:

$$\begin{aligned}
\mu \log(-\mu J^M) \quad &= \quad \mu - \rho + \frac{\theta_P}{\theta_A} \left(-\rho + \bar{c}^M - \mu \log \bar{c}^M - \underline{\lambda} \bar{\gamma}^{N,M} \right) - \theta_P \mu \underline{\lambda} \phi \\
&\quad + \frac{1}{2} \sigma^2 \theta_P \left(\theta_P \mu^2 - 2\theta_P \bar{c}^M \mu + (\theta_P + \theta_A) \left(\bar{c}^M \right)^2 \right) \\
&\quad + \underline{\lambda} \left[\exp(\theta_P \mu \phi)(1 - \bar{\gamma}^{N,M})^{-\theta_P/\theta_A} - 1 \right]
\end{aligned}$$

As discussed above, the manager may have an incentive to take high risk actions when faced with the moral hazard contract which treats risk as fixed. The conditions when this is optimal are described in the following lemma.

Lemma 2 *The manager has an incentive to choose the higher level of risk under the optimal contract with moral hazard and uncontrolled risk if and only if:*

$$\bar{c}^M \geq \frac{-1}{\phi \theta_A} \bar{\gamma}^{N,M} = \frac{1}{\phi \theta_A} \left[\exp \left(\frac{\theta_A \theta_P}{\theta_P + \theta_A} \mu \phi \right) - 1 \right] \equiv \bar{c}_0^M. \tag{12}$$

This result defines the region $\mathcal{R}$ where the manager has the incentive to take the riskier action. Note that the condition depends only on parameters, so the region $\mathcal{R}$ either covers the whole (y, q) space or is empty. For example, when the owner's risk aversion θ_P is small enough, the condition holds and the manager takes the riskier action. This is detrimental to the owner, as he prefers the low risk action, as formally shown in the following proposition.

Proposition 3 *If risk-taking is not controlled by the manager, the owner always prefers the low risk action.*

Proof. From the HJB equation, if the owner were allowed to choose the optimal risk level he would solve:

$$\begin{aligned}\lambda^M &= \arg\min_{\tilde\lambda\in\{\underline{\lambda},\bar\lambda\}} \tilde\lambda\left[-\frac{\theta_P}{\theta_A}\bar\gamma^{N,M} - \theta_P\mu\phi + \exp\left(\theta_P\mu\phi\right)\left(1-\bar\gamma^{N,M}\right)^{-\frac{\theta_P}{\theta_A}} - 1\right] \\ &= \arg\min_{\tilde\lambda\in\{\underline{\lambda},\bar\lambda\}} \tilde\lambda\left[\exp\left(\frac{\theta_P\theta_A}{\theta_A+\theta_P}\mu\phi\right) - \left(1+\frac{\theta_P\theta_A}{\theta_A+\theta_P}\mu\phi\right)\right].\end{aligned}$$

Since $\frac{\theta_P\theta_A}{\theta_A+\theta_P}\mu\phi > 0$, we have:

$$\exp\left(\frac{\theta_P\theta_A}{\theta_A+\theta_P}\mu\phi\right) - \left(1+\frac{\theta_P\theta_A}{\theta_A+\theta_P}\mu\phi\right) > 0.$$

Therefore the owner prefers the low risk level and $\lambda^M = \underline{\lambda}$. ■

These results provide parametric examples of the conditions discussed in Section 3 above. When θ_P is small, the owner is close to risk neutral. Therefore he will absorb more of the jump risk, and thus $\bar\gamma^{N,M}$ will be close to zero. Since the manager then is relatively insured against losses, he benefits from choosing the higher risk $\overline{\lambda}$ and earning the greater risk compensation, as Lemma 2 shows. The owner would interpret this as better firm performance, and would reward the manager. As Proposition 3 shows, the higher risk and increased likelihood of large losses is excessive from the owner's viewpoint.

4.3 Moral Hazard in Risk Management

We now turn to the risk management case where the owner accounts for the manager's endogenous choice of risk. On $\mathcal{H}$, the owner's value function is the same as in the uncontrolled case above with $\lambda = \bar\lambda$. Even though the owner would prefer the low risk level if he were able to control it, the information frictions make it too costly to do so. On $\mathcal{L}\cap\mathcal{R}$ the owner implements the low risk and γ^N is incentive-constrained. The form of the optimal policy for d is again unchanged, apart from the constant which is now J^R:

$$d^R = -\frac{\log(-\mu J^R)}{\theta_P} + \mu y + \frac{1}{\theta_A}\log(-q).$$

We again define the consumption rate so $\bar c^R q = -\exp(-\theta_A c^R)$, where now $\bar c^R$ solves:

$$\bar c^R = \mu + \sigma^2\theta_A\bar c^R(\theta_P\mu - (\theta_P+\theta_A)\bar c^R) - \underline{\lambda}\theta_A\phi\bar c^R\left[1-\exp\left(\theta_P\mu\phi\right)\left(1+\phi\theta_A\bar c^R\right)^{-(\theta_P+\theta_A)/\theta_A}\right]. \quad (13)$$

As above, the second term in this expression captures the moral hazard and risk sharing components associated with the Brownian shocks. The additional terms following $\underline{\lambda}$ tie the manager's consumption to the likelihood of jump shocks, capturing the incentive to manage risk.

The sensitivity of promised utility is again determined by the incentive constraint:

$$\gamma^{W,R} = -\theta_A \bar{c}^R q.$$

For risk management, we normalize $\gamma^N = \bar{\gamma}^{N,R} q$ and the incentive constraint gives:

$$\bar{\gamma}^{N,R} = -\phi\theta_A \bar{c}^R.$$

Then note that as $\theta_P \to 0$ from (13) we have:

$$\bar{c}^R = \mu + \theta_A^2 \left(\bar{c}^R\right)^2 \left[\sigma^2 - \underline{\lambda}\frac{\phi}{1-\bar{\gamma}^{N,R}}\right].$$

Relative to moral hazard alone, this adds the final term reflecting risk management, which compares the ratio of the loss from the negative shock (ϕ) to the promised utility remaining after the shock ($1-\bar{\gamma}^{N,R}$). The constant term J^R then solves:

$$\begin{aligned}
\mu \log(-\mu J^R) \quad = \quad & \mu - \rho + \frac{\theta_P}{\theta_A}\left(-\rho + \bar{c}^R - \mu \log \bar{c}^R - \underline{\lambda}\bar{\gamma}^{N,R}\right) - \theta_P \mu \underline{\lambda}\phi \\
& + \frac{1}{2}\sigma^2\theta_P\left(\theta_P\mu^2 - 2\theta_P \bar{c}^R \mu + (\theta_P + \theta_A)\left(\bar{c}^R\right)^2\right) \\
& + \underline{\lambda}\left[\exp(\theta_P\mu\phi)(1-\bar{\gamma}^{N,R})^{-\theta_P/\theta_A} - 1\right].
\end{aligned}$$

The following result compares the contracts under moral hazard and risk management, noting that on the safe region $\mathcal{S}$, the policies are the same.

Proposition 4 *If* (12) *is satisfied then* $(y,q) \in \mathcal{R}$ *and solutions of* $\bar{c}^R$ *and* $\bar{\gamma}^{N,R}$ *exist with* $\bar{c}^R \leq \bar{c}^M$ *and* $\bar{\gamma}^{N,R} \leq \bar{\gamma}^{N,M}$*; otherwise* $(y,q) \in \mathcal{S}$ *and* $\bar{c}^R = \bar{c}^M$ *and* $\bar{\gamma}^{N,R} = \bar{\gamma}^{N,M}$.

Proof. See Appendix D. ■

When the risk management constraint binds, the manager's consumption rate is lower than under moral hazard alone, which in turn is lower than under full information. Recalling that the level of consumption varies negatively with the consumption rate, we see that consumption increases with more severe information frictions. To provide incentives, these frictions require that the manager's pay respond more to shocks. To deliver the same utility q with a more volatile pay stream thus requires higher average pay. The promised utility responses also differ depending on the source of the information friction. Recalling that $q < 0$, under risk management the manager's promised utility is more sensitive to jump shocks ($\bar{\gamma}^{N,R}q \geq \bar{\gamma}^{N,M}q$) and less sensitive to the Brownian shocks ($\gamma^{W,R} \leq \gamma^{W,M}$) than under moral hazard (with uncontrolled risk).

These results agree with our discussion above. The usual pay-for-performance channel in the contract comes through γ^W, and larger sensitivities increase the incentive for the manager to take more risk to gain the risk compensation. So lowering γ^W under risk management

Parameter	Value	Parameter	Value
μ	0.10	θ_A	1.00
σ	0.16	θ_P	1.00
ϕ	0.50	ρ	0.04
$\bar{\lambda}$	0.90	$\underline{\lambda}$	0.10

Table 1: Benchmark parameter values

lessens this temptation. At the same time, increasing the sensitivity to the jump shock γ^N ensures that the manager will face a larger utility loss if a negative jump shock hits. This also gives the manager less incentive to take high risk actions.

We now consider whether it is worthwhile for the owner to provide incentives for the manager to implement low risk. We showed above that the regions $\mathcal{R}$ and $\mathcal{S}$ summarizing the manager's incentives cover the whole state space for (y, q), dependent only on parameters. Similarly, the regions $\mathcal{L}$ and $\mathcal{H}$ summarizing the owner's preferences are also only functions of parameters. For example, $\mathcal{L}$ covers the entire state space if:

$$\begin{aligned} J^R_{\underline{\lambda}}(y,q) = J^R \exp(-\theta_P \mu y)(-q)^{-\theta_P/\theta_A} &\geq J^M_{\bar{\lambda}}(y,q) = J^M_{\bar{\lambda}} \exp(-\theta_P \mu y)(-q)^{-\theta_P/\theta_A} \\ \Rightarrow J^R &\geq J^M_{\bar{\lambda}}, \end{aligned}$$

where J^R and $J^M_{\bar{\lambda}}$ are the constants defined above. $\mathcal{H}$ covers the whole space if the reverse inequality holds. Thus optimal contracts will implement the same choice of high or low risk at every date, dependent on the parameters. We illustrate below how the key parameters determine which risk level is implemented, which bears on the relationship between firm characteristics and volatility.

4.4 Illustrations of the Optimal Contracts

In this subsection, we illustrate how the features of the optimal contract depend on the parameters. We focus on variations in μ, σ, θ_A, θ_P, and ϕ which determine the production technology, monitoring structure, the preferences, and the size of the jump shock. The benchmark parameter values are listed in Table 1.

We consider the cases of full information with low risk, moral hazard with low and high uncontrollable risks, and risk management. Table 2 summarizes the constants which characterize the optimal contracts. The constants from the normalized firm value are listed in the first row. The firm value is highest under full information, and decreases as the moral hazard problem arises and the level of the risk increases. With the benchmark parameter values, moral hazard with low risk is not incentive compatible. So risk-management is optimal as it alleviates the efficiency loss due to the manager's ability to secretly take on

	First best($\underline{\lambda}$)	Moral Hazard ($\underline{\lambda}$)	Moral Hazard ($\bar{\lambda}$)	Risk Management
J^i	-332.43	-332.65	-334.33	-332.84
$\bar{c}^i$	0.10	0.0997	0.0997	0.0993
$-\bar{\gamma}^{W,i}$	0.0500	0.0997	0.0997	0.0993
$-\bar{\gamma}^{N,i}$	0.02531	0.02532	0.02532	0.04964
$\bar{d}^i$	-3.5038	-3.5045	-3.5095	-3.5050

Table 2: Characterization of the optimal contracts

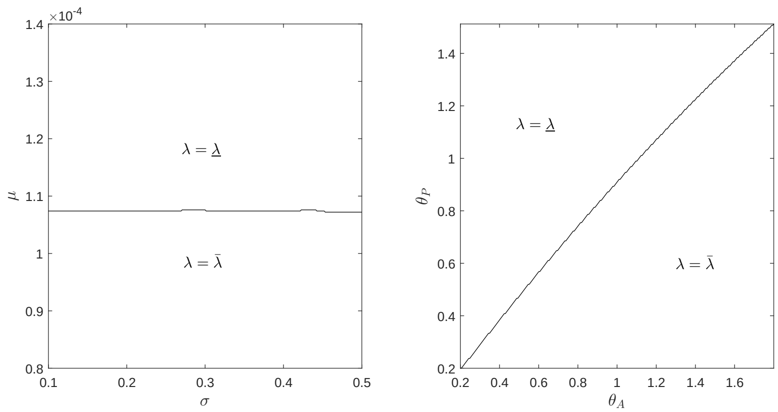

Figure 1: Division of the subsets of the parameter space into the high ($\mathcal{H}$) and low ($\mathcal{L}$) risk regions.

excess risk. The consumption rate $\bar{c}$ is listed on the second row, which recall varies inversely with the level of consumption. The manager's compensation is the lowest in the first best case because his utility process is the least risky. This is clear from the corresponding value of $-\bar{\gamma}^W$ which indicates the sensitivity of q with respect to the cash flow shocks. With risk management, this sensitivity is lower, since making pay less sensitive to performance reduces the manager's incentive to take high risk. The sensitivity of utility to the jump shocks $-\bar{\gamma}^N$ is nearly the same under full information and moral hazard, but is much larger under risk management in order to deter risk. Dividends increase with $\bar{d}^i$, the constant term in the d^i function, which follows the same pattern as the firm value J.

We next illustrate the effect of some key parameters on the risk incentives. We fix the values of other parameters and show how the regions $\mathcal{L}$ and $\mathcal{H}$ are divided over the σ-μ and θ_A-θ_P spaces in Figure 4.4. Risk management is always optimal unless μ is very low, approximately 1.07×10^{-4}, or unless the ratio of the owner's risk-aversion coefficient to the manager's is low. In other words, if the firm's profitability is relatively low or the owner is relatively less risk averse, then allowing the manager to take high risk is optimal. The optimal risk management policy is insensitive to σ, the volatility of the cash flows, which

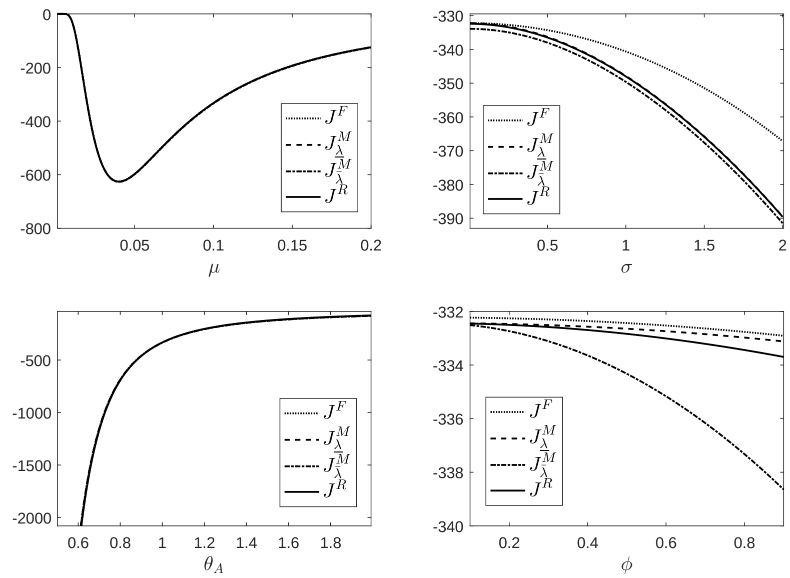

Figure 2: Comparative statics of the firm value, $\bar{J}$.

also governs the severity of the moral hazard friction. Under this specification of the model, high risk is only optimal for a subset of parameter values. Our example below considers a different specification where high risk eventually prevails for most firms.

We now analyze how the features of optimal contracts vary with key parameters. Figure 2 shows how the firm value changes in each case. The upper-left panel shows that there is a non-monotonic relationship between the value and μ which is nearly the same across the different cases. Specifically, the value decreases with μ up to a threshold level (which is close to ρ), and increases with μ above that point.[4] The firm value decreases with σ because both the owner and the manager are risk averse. In addition, higher σ makes monitoring more difficult, so the moral hazard friction is worse and the efficiency of the private information contracts further decrease. The firm value increases with the manager's risk aversion coefficient θ_A, according to lower-left panel. A lower θ_A means a smaller risk premium needs to be paid to deliver the same promised utility level. A larger ϕ means larger losses from the jump shock, so J decreases in all cases. Choosing lower risk and imposing risk management greatly alleviates the efficiency loss, as J^R is much closer to J^F than is $J^M_{\underline{\lambda}}$. We don't show it, but the divided rate $\bar{d}$ has similar comparative statics to the firm value J.

Figure 3 shows how the negative of the consumption rate, $-\bar{c}$ which varies positively

[4] Without uncertainty the threshold level is ρ. To understand this, assume that there is no uncertainty in the asset return or moral hazard. If $\mu < \rho$, the owner would front load her dividend payment and a higher μ raises the opportunity cost of doing so due to the intertemporal substitution effect. Symmetrically, if $\mu > \rho$, she would delay the dividend payment, and a higher μ raises the benefit of doing so.

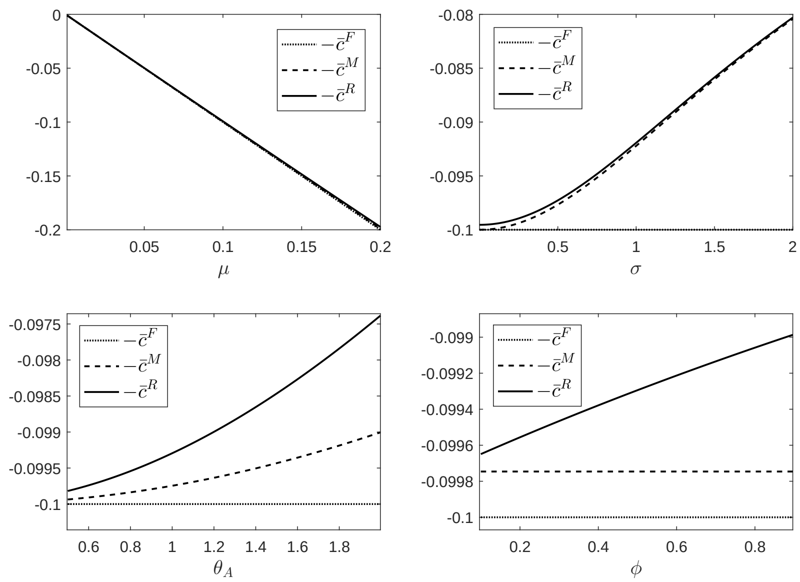

Figure 3: Comparative statics of the negative of the manager's consumption rate $-\bar{c}$.

with the level of consumption c, depends on the parameters.[5] A higher μ indicates better expected growth for the firm's assets, so a lower compensation level is needed to deliver the promised utility, as shown in the upper-left panel. With moral hazard, higher levels of σ and θ_A increase the manager's required compensation for bearing risk due to incentive provision. Therefore $-\bar{c}$ increases, as shown in the upper-right and lower-left panels. Increases in θ_A also increase the gaps between consumption rates in the different information structures. More severe information frictions require the manager to bear more risk, which is more costly when he is more risk averse. The lower-right panel shows that only compensation with risk management varies with ϕ. Since the risk is compensated, when the potential losses ϕ are larger so is the hidden benefit of high risk. To deter excess risk, the manager must then bear more overall risk, requiring more compensation.

4.5 Implementation of the Optimal Contracts

So far we have focused on a direct implementation of the optimal contracts, where the owner provides a history-dependent payment to the manager which provides appropriate incentives. We now show how the optimal contracts in the moral hazard and risk management cases can be implemented with relatively simple instruments. In particular, the manager receives a constant salary payment $\bar{s}$ each period, but also can save or borrow in an account we

[5] Notice that $-\bar{c}$ in the case of moral hazard with uncontrollable risk is independent of the arrival rate of the disasters, so we plot $\bar{c}^M = \bar{c}^M_{\underline{\lambda}} = \bar{c}^M_{\bar{\lambda}}$.

denote x_t. We interpret this account as representing the cumulation of the manager's bonus or deferred compensation. This account has a constant interest rate r and requires the manager to pay a flat fee b. Out of the firm's assets, the owner withdraws a dividend d_t for himself, and pays the manager his salary plus the interest on his account balance. Then the owner gives the manager the remaining assets y_t, whose increments accrue to the manager's account. Giving the manager the entire remainder of the firm's assets would make the manager experience the full effect ϕ of losses due to negative jump shocks. While this may deter the manager from taking excess risk, it is in general too strong a response to optimally balance risk sharing and risk management incentives. Thus the contract also provides insurance α which partially compensates the manager for jump risks. In the event of a negative jump shock, the manager then loses $\phi - \alpha$ of his deferred compensation. This is a form of a "clawback" provision, similar to what has been proposed by federal regulators and implemented in some financial institutions. Facing this wealth evolution, the manager then chooses the risk level and whether to divert resources from the firm for his own consumption, as well as whether to make additional withdrawals from his account. We show that this combination of salary payment, deferred compensation, shares in the firm, and clawback provisions implements the optimal contracts.

In particular, we choose the instruments so that the agent's optimal choices agree with those in the direct implementations of the contracts. To do so, we show that the promised utility q_t can be mapped into the manager's account balance x_t via the relationship:

$$x_t = -\frac{1}{\theta_A r}\log(q_t). \tag{14}$$

Using this transformation, the optimal dividend policies d^F, d^M, and d^R under the different information structures derived above can all be written:

$$\begin{aligned} d_t^i &= \bar{d}^i + \mu y_t - \frac{1}{\theta_A}\log(-q_t) \\ &= \bar{d}^i + \mu y_t - r x_t \end{aligned}$$

where $\bar{d}^i = -\log(-\mu J^i)/\theta_P$ is a constant with $i = F, M, R$. We now must distinguish among the payment the manager receives from the owner, which we denote s_t, his diversion of resources Δ_t, and his total consumption c_t. With wealth x_t in the account, the payment consists of the interest payment plus the constant flow salary:

$$s_t = r x_t + \bar{s}.$$

Then the evolution of the firm's assets can be written:

$$\begin{aligned} dy_t &= [\mu y_t - s_t - d_t + \lambda_t\phi - \Delta_t]dt + \sigma dW_t - \phi dN_t \\ &= [\mu y_t - (r x_t + \bar{s}) - (\bar{d} + \mu y_t - r x_t) + \lambda_t\phi - \Delta_t]dt + \sigma dW_t - \phi dN_t \\ &= [-(\bar{s} + \bar{d}) + \lambda_t\phi - \Delta_t]dt + \sigma dW_t - \phi dN_t. \end{aligned}$$

At any date t, the manager begins with assets x_t in the account, which earn the constant interest rate r. In addition, the account receives an inflow due to his salary $\bar{s}$ and an outflow for the flat fee b. The manager also receives the gain on the firm's assets after payments, and gets the additional insurance α in the event of a negative jump shock. The manager's total consumption c_t consists of any diversion Δ_t from the firm's assets plus whatever additional amount he withdraws from the account. We write the withdrawal as $c_t - \Delta_t$. Altogether, the manager's wealth evolves as:

$$\begin{aligned} dx_t &= [rx_t + \bar{s} - b - (c_t - \Delta_t)]dt + dy_t + \alpha dN_t \\ &= [rx_t + \bar{s} - b - (c_t - \Delta_t) - (\bar{s} + \bar{d}) + \lambda_t \phi - \Delta_t]dt + \sigma dW_t - (\phi - \alpha)dN_t \\ &= [rx_t - c_t + \lambda_t \phi - \bar{b}]dt + \sigma dW_t - \beta dN_t \end{aligned} \tag{15}$$

where $\beta = \phi - \alpha$ and $\bar{b} = \bar{d} + b$. Since the manager retains the residual assets of the firm, the diversion Δ_t drops out of his wealth evolution. He only cares about his total consumption, not whether it is financed by diverting assets from the firm or by withdrawing from his account. Thus without loss of generality we set $\Delta_t \equiv 0$, but note this does not yet imply that the manager consumes the amount suggested by the contract. We refer to β, the manager's net losses due to a negative jump shock, as the "clawback".

Facing the wealth evolution (15) with given initial wealth x, the manager chooses his consumption and the risk level in order to maximize his utility. The HJB equation for this problem with value $V(x)$ can be written:

$$\rho V(x) = \max_{c,\lambda} \left\{ -\exp(-\theta_A c) + V'(x)[rx - c + \lambda\phi - \bar{b}] + \frac{1}{2}V''(x)\sigma^2 + \lambda[V(x-\beta) - V(x)] \right\}.$$

The value function has a form similar to our J^i functions above, $V(x) = v\exp(-r\theta_A x)$, where v is a constant. This leads to the optimal policy function:

$$c(x) = -\frac{1}{\theta_A}\log(-rv) + rx,$$

and the risk choice is $\lambda_t = \underline{\lambda}$ if:

$$r\theta_A\beta \geq \log(1 + r\theta_A\phi),$$

which provides a lower bound on the clawback β, and $\lambda_t = \overline{\lambda}$ otherwise. Given λ, the constant v solves:

$$r\log(-rv) = r - \rho - \theta_A r(\lambda\phi - \bar{b}) + \frac{1}{2}\sigma^2\theta_A^2 r^2 + \lambda[\exp(\beta\theta_A r) - 1]. \tag{16}$$

If we let μ_x denote the drift of x_t from (15) above, we have:

$$\begin{aligned} \mu_x &= rx - c(x) + \lambda\phi - \bar{b} \\ &= -\frac{1}{\theta_A}\log(-rv) + \lambda\phi - \bar{b} \\ &= \frac{1}{\theta_A r}\left(r - \rho + \frac{1}{2}\sigma^2\theta_A^2 r^2 + \lambda[\exp(\beta\theta_A r) - 1] \right), \end{aligned}$$

where the last equality uses (16).

We now show that we can choose the parameters $(r, b, \alpha, \bar{s})$ to implement the optimal contracts under moral hazard and risk management. We have already conjectured the link between q_t and x_t in (14) above. Given the form of $V(x)$, this implies that we must have $v = -1$. Under this relationship, the implementation already delivers the optimal dividend d^i . Moreover, under the optimal contract,

$$c^i = -\frac{1}{\theta_A}\log\left(\bar{c}^i\right) - \frac{1}{\theta_A}\log\left(-q\right), \tag{17}$$

while under the implementation:

$$c^i = -\frac{1}{\theta_A}\log\left(-rv\right) + rx = -\frac{1}{\theta_A}\log\left(r\right) - \frac{1}{\theta_A}\log\left(-q\right) \text{ for } i = F, M, R. \tag{18}$$

Therefore (17) and (18) imply that the implementation must set the interest rate the manager faces equal to his consumption rate:

$$r = \bar{c}^i.$$

The total payments from the owner, $rx + \bar{s}$, must also match total consumption, so:

$$\bar{s} = -\frac{1}{\theta_A}\log\left(\bar{c}^i\right) \text{ for } i = F, M, R.$$

We then need to establish the validity of the conjectured relationship between x and q from equation (14) above. To do so, recall that in both the moral hazard and risk management cases the incentive constraint for consumption holds:

$$\gamma^{W,i} = -\theta_A\bar{c}^i q,$$

while the loading on the jump shock satisfies $\gamma^{N,i} = \bar{\gamma}^{N,i}q$. Using these, along with the consumption and dividend policies, we find that promised utility evolves as:

$$dq_t = [\rho - \bar{c}^i + \lambda\bar{\gamma}^{N,i}]q_t dt - \sigma\theta_A\bar{c}^i q_t dW_t - \bar{\gamma}^{N,i}q_t dN_t.$$

Then using the relation (14) and applying a generalized version of Ito's lemma[6], we have:

$$dx_t = \frac{1}{\theta_A r}[\bar{c}^i - \rho + \frac{1}{2}\sigma^2\theta_A^2(\bar{c}^i)^2 - \lambda\bar{\gamma}^{N,i}]dt + \sigma dW_t - \frac{\log(1-\bar{\gamma}^{N,i})}{\theta_A r}dN_t.$$

Comparing with (15), the conjectured relationship between x_t and q_t holds as long as:

$$\beta = \frac{\log(1-\bar{\gamma}^{N,i})}{\theta_A r},$$

so $\alpha = \phi - \beta$. This setting ensures the loading on the jump shock N_t matches (15), and also equates the drift with μ_x defined above. Finally, given the other parameters of the implementation, we set the flat fee b so that $v = -1$, which from (16) requires:

$$b = \frac{1}{\theta_A r}\left(r\log r + \rho - r + \lambda\bar{\gamma}^{N,i}\right) + \lambda\phi - \bar{d} - \frac{1}{2}\sigma^2\theta_A r.$$

	First best($\underline{\lambda}$)	Moral Hazard ($\underline{\lambda}$)	Risk Management
r	0.10	0.0997	0.0993
b	0.6247	0.8257	0.5970
β	0.2500	0.2506	0.4880
$\bar{s}$	2.3026	2.3051	2.3097

Table 3: Implementation of the optimal contracts.

In Table 3, we illustrate the implementation of the optimal contracts based on the benchmark parameter values from above. Recall that we have $r = \bar{c}$, so increases in the interest rate imply reductions in the level of consumption. Focusing on the comparison between the moral hazard and risk management cases, we see that in accounting for the manager's risk management incentives, the owner slightly reduces the interest rate r and increases the salary payment $\bar{s}$. The owner also cuts the flat fee b the manager must pay on the account. All of these changes are associated with higher consumption by the manager, which reflects compensation for the additional risk he must bear. This risk is seen most clearly in the significant increase in the clawback term β under risk management, which translates to a smaller amount of insurance α. Recalling that $\phi = 0.5$, under full information the manager and owner fully share the losses of negative jump shocks, as the manager faces a loss of $\beta = 0.25$. Under moral hazard alone, the clawback increases slightly, as reductions in x_t make it easier for the owner to provide future incentives to deter diversion. But under risk management, the clawback β increases substantially. To provide incentives for prudent behavior, the manager faces a much larger share of potential losses. The clawback provision is thus the main channel for providing risk management incentives.

While our analysis allows for comparative statics of any aspect of the implementation, we focus here on the clawback term which is the key for risk management. Figure 4 plots the clawback term β versus some of the parameters under moral hazard (treating risk as uncontrollable) and risk management. As expected, the clawbacks are substantially larger for all parameters under risk management than moral hazard. Increases in the agent's risk aversion θ_A mean that under full information the agent should face a smaller share of the risk, which is also reflected by a reduction in the clawback under moral hazard. Also under moral hazard, increases in the diffusion risk σ effectively worsen the information friction governing diversion, so the clawback increases to lessen the cost of providing incentives. The clawbacks under risk management decrease slightly with faster firm growth μ and the agent's risk aversion θ_A, but are otherwise relatively insensitive to the parameters. The clawbacks in both contracts increase roughly proportionately with the size of the loss ϕ. Under moral

[6] See Theorem 1.14 and Example 1.15 in Øksendal and Sulem [2005]

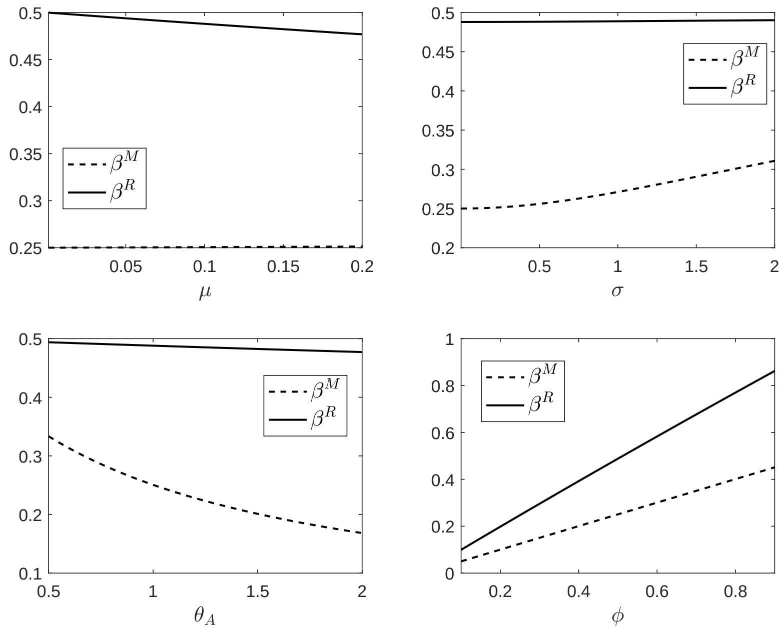

Figure 4: Comparative statics of the "clawback" term β in the event of a negative jump.

hazard, the slope is roughly one-half as the risk is shared nearly evenly, while under risk management the slope is roughly unity as the manager bears the bulk of the jump risk.

5 A Case with Variable Risk Management

While the previous example allowed us to gain insight into the nature of the contracts, it was very special. The risk management incentives of both parties were constant, so that there was no time variation in the jump risk. We now consider a specification with variable risk management, where the manager's incentives and the owner's preferences to implement high or low risk vary over time. Our results suggest that rather than high risk actions representing a failure of corporate management, in some cases high risk may be optimal.

5.1 The Specification

Time variation in risk management is a generic feature of our general model, which we illustrate with a tractable but plausible specification. We now assume that both the owner and the manager have logarithmic preferences, and the firm's assets follow a linear evolution with constant expected returns, proportional diffusion risk, and proportional losses from the jump shock. That is, we now assume the following functional forms:

$$u(c) = \log(c),\ v(d) = \log(d),\quad \mu(y) = \mu y,\quad \sigma(y) = \sigma y,\quad \phi(y) = \phi y,$$

where we recycle some notation from above. As above, under each information structure the owner's value function a similar form, but here the private information cases cannot be solved explicitly. Nonetheless, the homogeneity of this specification simplifies the solutions.

Reconsidering the boundaries of the Pareto frontier, the upper bound on the manager's utility now takes the form:

$$\rho\bar{q}(y) \;=\; \max_c \left\{ \log c \,+\, \bar{q}'(y)[\mu y \,-\, c \,+\, \underline{\lambda}\phi y] \,+\, \frac{1}{2}\sigma^2 y^2 \bar{q}''(y) \,+\, \underline{\lambda}[\bar{q}((1 \,-\, \phi)y) \,-\, \bar{q}(y)] \right\}.$$

This has the solution:

$$\bar{q}(y) = \log(y)/\rho + z_p$$

where the constant z_p is given by:

$$z_p = \log(\rho)/\rho + [\mu - \rho - 1/2\sigma^2 + \underline{\lambda}(\phi + \log(1-\phi))]/\rho^2.$$

At this bound, the manager consumes a constant fraction of assets, $c = \rho y$ and the owner's utility diverges as his consumption tends to zero, giving the boundary condition $\lim_{q\to\bar{q}(y)} J^F(y,q) = -\infty$. The upper bound on the owner's utility, which arises in the limit as $q \to -\infty$, can be characterized in the same way:

$$\bar{J}(y) = \log(y)/\rho + z_p,$$

with the same constant z_p and consumption function $d = \rho y$.

We now define a variable $z = q - \log(y)/\rho$ which effectively indexes the location on the Pareto frontier between the extremes of $\bar{q}$ and $\bar{J}$. Intuitively, we think of z as the manager's "share" in the firm. A large z means that, in expectation, a larger fraction of the firm's future cash flows will be paid as the manager's compensation. Thus a higher z implies that the manager's incentives are more aligned with the firm's interests.[7] Under each information structure, the owner's value function can be written:

$$J^i(y,q) = \log(y)/\rho + g_i(q - \log(y)/\rho),\quad i = F, M, R,$$

where $g_i(z)$ is a function which differs in each case. The boundary conditions imply:

$$\lim_{z\to z_p} g_i(z) = -\infty \text{ for } i = F, M, R, \text{ and } \quad \lim_{z\to-\infty} g_F(z) = z_p.$$

To provide the manager with the maximal utility, the owner must consume nothing. Conversely, providing the manager with no consumption only delivers the owner his maximal utility under full information, as there are losses from private information.

[7] See Ai and Li [2015], Ai et al. [2016], and DeMarzo et al. [2012] for similar normalizations and interpretations with different preferences.

5.2 Full Information

With the form of the value function, under full information the optimality conditions imply:

$$
\begin{aligned}
d^F &= \frac{1}{1-g_f'(z)}\rho y, \\
c^F &= \frac{-g_f'(z)}{1-g_f'(z)}\rho y \\
\gamma^{W,F} &= \frac{1}{\rho y} \\
\gamma^{N,F} &= -\frac{1}{\rho}\log(1-\phi).
\end{aligned}
$$

Substituting these into the HJB equation, we find that $g_F(z)$ solves:

$$\rho g_F(z) = (1-g_f'(z))[\rho z_p - \log(1-g_f'(z))] + g_f'(z)[\rho z - \log(-g_f'(z))].$$

The solution of this first-order ODE is:

$$g_F(z) = z_p + \log(1-\exp(\rho(z-z_p)))/\rho,$$

which satisfies the boundary conditions. Then we can write the value function as:

$$J^F(y,q) = \log\left(y\exp(\rho z_p) - \exp(\rho q)\right)/\rho,$$

and the policy functions simplify to:

$$
\begin{aligned}
d^F &= (1-\exp(\rho(z-z_p)))\rho y = \rho y - \frac{\rho\exp(\rho q)}{\exp(\rho z_p)} \\
c^F &= \exp(\rho(z-z_p))\rho y = \frac{\rho\exp(\rho q)}{\exp(\rho z_p)}.
\end{aligned}
$$

As in the exponential case, the manager's consumption c^F is only a (direct) function of his promised utility q and not the firm's assets y. The owner's consumption is set at his maximally desired level ρy minus the consumption he provides the manager, so the total amount of consumption is the efficient level: $c^F + d^F = \rho y$. The current promised utility thus determines how the efficient total consumption is divided between the owner and manager.

5.3 Moral Hazard with Uncontrolled Risk

Under moral hazard with uncontrolled risk, the optimality conditions for the dividend d^M and weighting on the jump shock $\gamma^{N,M}$ have the same form as above, simply substituting the function $g_M(z)$ for $g_F(z)$. We now normalize consumption as $c^M = \bar{c}^M(z)y$, a form similar to the full information case, so the incentive constraint for diversion (3) gives:

$$\gamma^{W,M} = \frac{1}{c^M} = \frac{1}{\bar{c}^M(z)y}. \tag{19}$$

The optimality condition for consumption from (7) then implies:

$$g'_M(z)\bar{c}^M(z)^2 + \frac{1}{\rho}(1-g'_M(z))\bar{c}^M(z)^3 + \sigma^2 g''_M(z)\left(1 - \frac{\bar{c}^M(z)}{\rho}\right) = 0.$$

The manager's consumption rate $\bar{c}^M(z)$ is the maximizing root of this cubic equation. Substituting the choices into the HJB equation, the function $g_M(z)$ then solves the following second order ODE:

$$\begin{aligned}\rho g_M(z) &= -\log(1-g'_M(z)) + \log\rho - 1 + \frac{1-g'_M(z)}{\rho}[\mu - \bar{c}^M(z) + \lambda(\phi + \log(1-\phi))] \\ &\quad + g'_M(z)(\rho z - \log\bar{c}^M(z)) - \frac{\sigma^2}{2\rho}\left[1 - g'_M(z) - g''_M(z)\left(\frac{1}{\rho} - \frac{2}{\bar{c}^M(z)} + \frac{\rho}{\bar{c}^M(z)^2}\right)\right].\end{aligned}$$

This ODE has the boundary conditions $\lim_{z\to z_p} g_M(z) = -\infty$ and $\lim_{z\to-\infty} g'_M(z) = 0$. Unlike the previous examples, we must solve this ODE numerically.

5.4 Moral Hazard in Risk Management

Under risk management, we again normalize consumption $c^R = \bar{c}^R(z)y$. The optimality condition for dividends and the incentive constraint for diversion are then:

$$\begin{aligned}d^R &= \frac{1}{1-g'_R(z)}\rho y \equiv \bar{d}^R(z)y, \\ \gamma^{W,R} &= \frac{1}{c^R} = \frac{1}{\bar{c}^R(z)y}.\end{aligned}$$

On $\mathcal{H}$, risk management does not add any constraint and we have:

$$\gamma^{N,R} = \gamma^{N,F} = -\frac{1}{\rho}\log(1-\phi),$$

along with the other conditions from the previous section. But on $\mathcal{L}$ the contract now must satisfy the incentive constraint for risk:

$$\gamma^{N,R} = \frac{\phi y}{c^R} = \frac{\phi}{\bar{c}^R(z)},$$

The optimality condition for consumption (10) then implies:

$$\begin{aligned}0 &= g'_R(z)\bar{c}^R(z)^2 + \frac{1}{\rho}\bar{c}^R(z)^3(1-g'_R(z)) + \sigma^2 g''_R(z) \\ &\quad + \underline{\lambda}\phi\bar{c}^R(z)\left[g'_R(z) - g'_R\left(z - \phi\frac{1}{\bar{c}^R(z)} - \frac{1}{\rho}\log(1-\phi)\right)\right].\end{aligned}$$

Substituting the choices into the HJB equation, the function $g_R(z)$ solves:

$$\begin{aligned}\rho g_R(z) &= \log\rho - \log\left(1-g_R'(z)\right) - 1 + \frac{1-g_R'(z)}{\rho}\left[\mu - \bar{c}^R(z) + \underline{\lambda}\phi\right] \\ &+ g_R'(z)\left[\rho z - \log\bar{c}^R(z) + \underline{\lambda}\phi\frac{1}{\bar{c}^R(z)}\right] \\ &- \frac{\sigma^2}{2\rho}\left[\left(1-g_R'(z)\right) - g_R''(z)\left(\frac{1}{\rho} - \frac{2}{\bar{c}^R(z)} + \frac{\rho}{\bar{c}^R(z)^2}\right)\right] \\ &+ \underline{\lambda}\left[g_R\left(z - \phi\frac{1}{\bar{c}^R(z)} - \frac{1}{\rho}\log(1-\phi)\right) - g_R(z)\right],\end{aligned}$$

with the same boundary conditions as above.

In the exponential case, the contract provided constant risk management incentives. But now the regions $\mathcal{R}$ and $\mathcal{S}$, where the manager's risk incentive constraint does or does not bind, as well as $\mathcal{L}$ and $\mathcal{H}$, where the owner implements the low or high risk levels, are functions of the endogenous share variable z. Under the optimal contract, the share z_t follows:

$$dz_t = \mu_z\left(z_t\right)dt + \left(\frac{1}{\bar{c}^R(z_t)} - \frac{1}{\rho}\right)\sigma dW_t - \left[\gamma^{N,R}(z_t) + \frac{1}{\rho}\log\left(1-\phi\right)\right]dN_t$$

with

$$\mu_z\left(z\right) = \rho z - \ln\left(\bar{c}^R(z)\right) + \lambda(z)\gamma^{N,R}(z) - \frac{1}{\rho}\left(\mu - \bar{d}^R(z) - \bar{c}^R(z) + \lambda(z)\phi + \frac{1}{2}\sigma^2\right).$$

As z_t evolves over time, the contract may now transit between the different regions, and thus switch between the high and low risk levels, as we show next.

5.5 Characterizations of the Optimal Contract

We now illustrate the dynamics of the optimal contract, using the parameter values summarized in Table 4. Since the policies depend on the share variable z, we focus on that dimension of the state space. The optimal risk policy and the related incentive constraint in the risk management case are shown in Figure 5. The upper panel shows that the higher risk level is optimal if and only if z is lower than a threshold level z^0. That is $\mathcal{H} = \{(y,q) : z < z^0\}$. When z increases, the manager's share of future cash flows increases and his interest is more aligned with the firm's. When z is low, a stronger pay-performance sensitivity is then needed to deter diversion. However, this has the side-effect of increasing the manager's willingness to take high risk, and so makes risk management more costly. In other words, the incentives for diversion and risk are more in conflict in the low z region, so it is optimal to give up risk management. This can be seen more clearly in the lower panel of the figure. Recall that the manager chooses the low risk if and only if $\gamma^{N,R} \geq \phi(y)\gamma^{W,R}$. On $\mathcal{H}$, the contract

Parameter	Value	Parameter	Value
μ	0.22	$\bar{\lambda}$	0.50
σ	0.16	$\underline{\lambda}$	0.20
ϕ	0.50	ρ	0.04

Table 4: The parameter values for the logarithmic utility case.

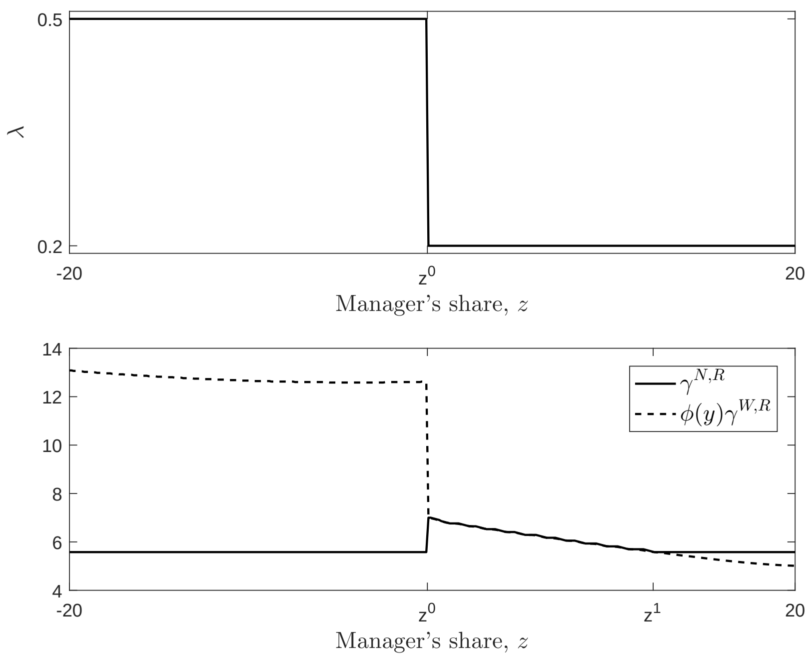

Figure 5: Risk management policy and the incentive compatibility conditions.

imposes very high level of the performance sensitivity $\phi(y)\gamma^{W,R}$, while $\gamma^{N,R} = \gamma^{N,F}$, the full information level without taking account of the manager's risk choice. To the right of z^0, risk management is optimal and less costly. Specifically, figure shows that the incentive constraint for risk management is binding over the interval $[z^0, z^1]$ (the region $\mathcal{L} \cup \mathcal{R}$) and relaxed if $z > z^1$ (the region $\mathcal{S}$). On $\mathcal{S}$, the manager's incentive to divert is too weak to influence his risk taking behavior.

The upper panel of Figure 6 plots the $g_i(z)$ functions from the firm's value in the different cases. Compared with g_M with $\bar{\lambda}$, the advantage of risk management increases with z, and g_R converges to g_M with $\underline{\lambda}$ as z converges to z_p. In the lower panel, we show the compensation rates in the different cases. Consumption in the full information and moral hazard cases are as expected, with the manager receiving the most compensation when he bears the most risk. Consumption under risk management $\bar{c}^R$ differs substantially from the other cases, jumping up when z surpasses z^0. To understand this, notice that the manager's incentives

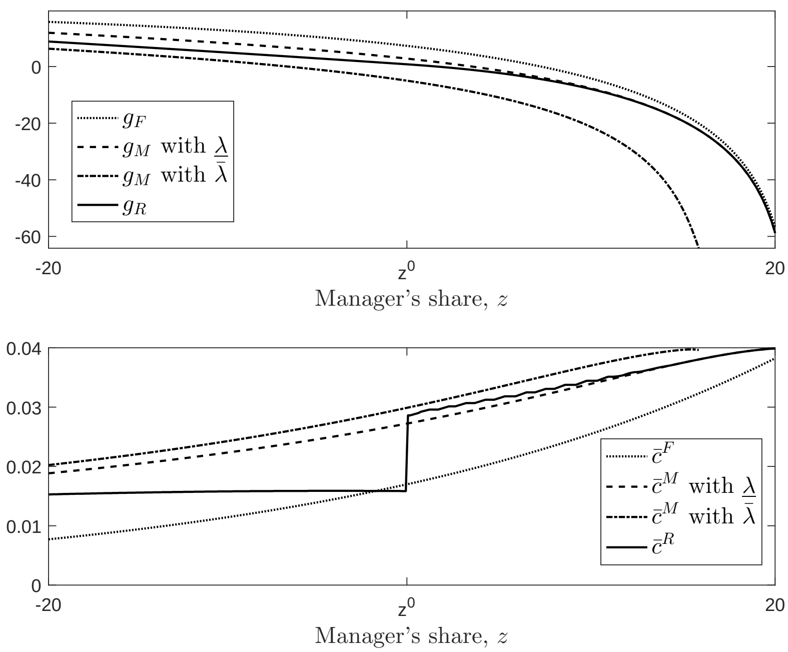

Figure 6: The key component of the firm value and the optimal consumption rate.

to divert and take high risk decrease with the current compensation level, as shown in the incentive constraints (19) and (8). Therefore, shifting compensation to the region with risk management lowers the cost.

To better understand the dynamics of the risk management contract, Figure 7 plots the drift and diffusion functions for the share variable z. The drift is negative for all z, so over time the contract will tend toward the low z region where the firm value and the level of risk are both higher. The diffusion rate is significantly higher in the region $\mathcal{H}$ (to the left of z^0), where compensation is significantly lower and stronger pay-performance-sensitivity is need to deter diversion. So the dynamics of z suggest that, as the contract evolves over time, a firm will ultimately move out of the region with risk management. Reductions in the firm's assets, whether due to a sequence of negative cash flow shocks or the large negative shocks, will make this process happen more quickly.

The dynamics of the risk management contract are further illustrated in Figure 8, which plots quantiles of the simulated distribution of firms over the z-space across time. For this figure, we simulate a panel of two million firms, each of which start with z equal to zero. At each date the figure shows the location of the quantiles of the cross-sectional distribution (in z) of firms. We see that over time the mass of the distribution tends toward low z values, where the moral hazard friction dominates and owners allow high risk. However there is an upper tail that remains above the threshold level z^0, and continues to impose low risk.

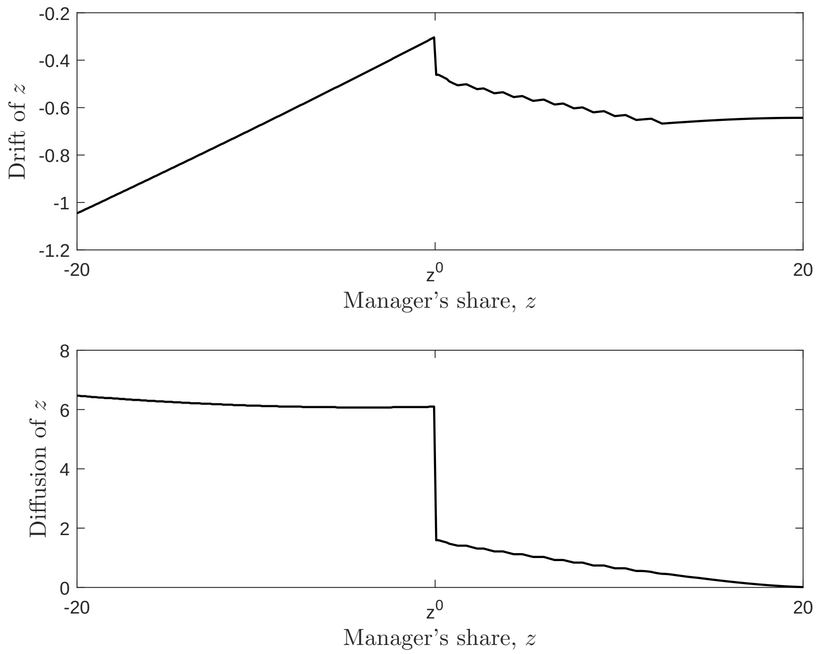

Figure 7: Dynamics of the share variable z.

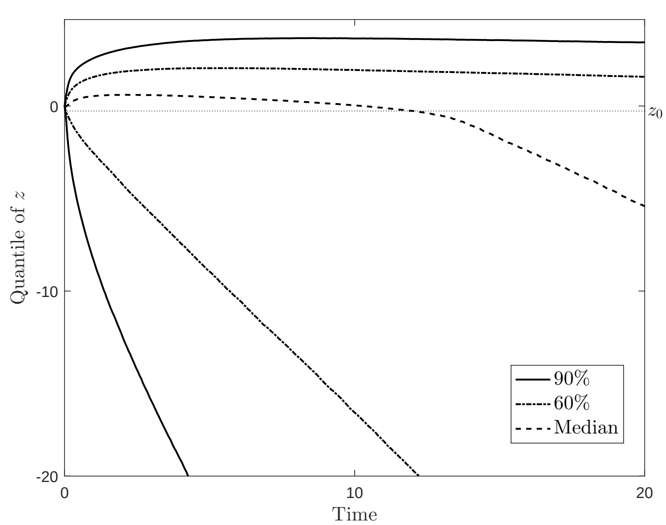

Figure 8: A simulated distribution of firms over the share variable z across time.

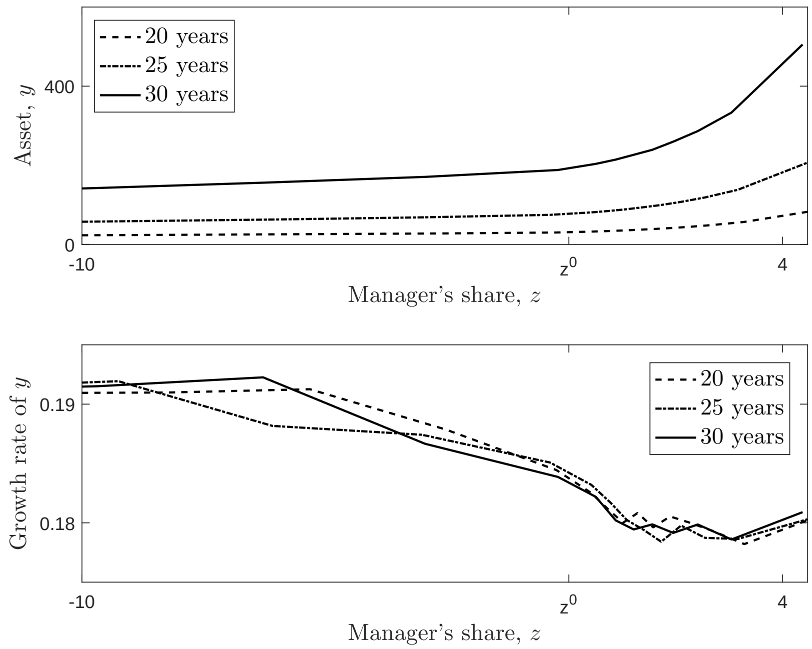

Figure 9: The relationship between the share variable z and firm assets (top panel) and growth rates (bottom panel) at different dates from a simulated panel of firms.

To gain further insight on the dynamics of firms under the risk management contract, Figure 9 shows the relationship between the firm assets y_t and the share variable z_t at different points in time. The contract is nonstationary, with firm assets growing and the share variable declining over time. To summarize the relationship between these variables, we take the simulated panel of firms and collect the cross-sections at years 20, 25, and 30. For each year, we divide the samples into 25 groups according to the level of z with an equal number of samples in each group. The top panel of the figure plots the average size of the assets y_t in each group at each date as a function of average share z_t, while the bottom panel plots the average growth rate of assets (the drift of $\log y_t$). The different years all show the same relationships, with later dates having a greater average size y_t for any z_t. Firm size increases with the share variable z_t, while the growth rate decreases with it. Thus the upper tail in Figure 8 implementing low risk corresponds to larger firms, with smaller expected growth rates. Our model also replicates the empirically observed negative relationship between firm size and growth, which here arises due to incentive provision.

6 Conclusion

In this paper, we have studied a dynamic agency model in which the manager can take two hidden actions: divert resources for his own consumption, and raise his perceived performance by exposing the firm to larger risk. We showed that the incentives to deter diversion and that to discourage risk can conflict with each other, and optimal incentive compatible contracts balance these objectives. With exponential preferences, we showed how the optimal risk policy depends on the parameters, and developed an implementation of the contract using some simple instruments. In particular, we showed that risk management concerns rationalize a clawback provision in contracts, where a manager is docked accumulated pay in the event his firm suffers a large loss. With logarithmic preferences, we showed that risk management varies over time, with most firms eventually tending toward giving up on risk management in favor of deterring diversion.

Our results suggest that the episodes documented in recent years where managers have exposed their firms to large losses could have two different interpretations. First, as Rajan [2011] and others have emphasized, firms may have adopted pay-for-performance contracts to provide incentives for managers, without recognizing that these contracts would also encourage hidden risk taking behavior. Alternatively, firms may well have considered the risk incentives of managers, but been in a situation where deterring risk was too costly. Our results suggest that may be a typical situation, as over time moral hazard incentives tend to dominate risk management concerns. Rather than being a failure of corporate management, high risk and large losses by firms may be optimal.

References

H. Ai, D. Kiku, and R. Li. Firm Dynamics under Limited Commitment. *Working Paper, University of Minnesota*, 2016.

Hengjie Ai and Rui Li. Investment and CEO Compensation under Limited Commitment. *Journal of Financial Economics*, 116:452–472, 2015.

B. Biais and C. Casamatta. Optimal Leverage and Aggregate Investment. *Journal of Finance*, 54:1291–1323, 1999.

B Biais, T Mariotti, Jean-Charles Rochet, and Stephane Villeneuve. Large Risks, Limited Liability, and Dynamic Moral Hazard. *Econometrica*, 78:73–118, 2010.

C Bray and J Baer. Ex-Trader Admits to Fraud. *Wall Street Journal*, April 3:C3, 2013.

P. DeMarzo, D. Livdan, and A. Tchistyi. Risking Other People's Money: Gambling, Limited Liability, and Optimal Incentives. *Working Paper, Stanford University*, 2014.

Peter DeMarzo, Michael Fishman, Zhiguo He, and Neng Wang. Dynamic Agency and q Theory of Investment. *Journal of Finance*, 67:2295–2340, 2012.

Peter M. DeMarzo and Yuliy Sannikov. Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model. *Journal of Finance*, 61:2681–2724, 2006.

P. Diamond. Managerial Incentives: On the Near Linearity of Optimal Compensation. *Journal of Political Economy*, 106:931–957, 1998.

Felix Ziyu Feng and Mark Westerfield. Dynamic Asset Allocation with Hidden Volatility. Working paper, University of Notre Dame, 2018.

J. Jacod and A. Shiryaev. *Limit Theorems for Stochastic Processes*. Springer, 2002.

M. Jensen and W. Meckling. Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure. *Journal of Financial Economics*, 3:305–360, 1976.

R. Li and N. Williams. Optimal Unemployment Insurance and Cyclical Fluctuations. *Working Paper, University of Wisconsin-Madison*, 2014.

Rui Li. Dynamic Agency with Persistent Observable Shocks. *Working paper, University of Massachusetts Boston*, 2016.

B. Øksendal and A. Sulem. *Applied Stochastic Control of Jump Diffusions*. Springer, New York, 2005.

F. Palomino. Risk Taking and Optimal Contracts for Money Managers. *Rand Journal of Economics*, 34:113–137, 2003.

R Rajan. *Fault Lines: How Hidden Fractures Still Threaten the World Economy*. Princeton University Press, 2011.

Jean-Charles Rochet and Guillaume Roger. Risky utilities. *Economic Theory*, Jun 2016.

Y. Sannikov. A Continuous-Time Version of the Principal-Agent Problem. *Review of Economic Studies*, 75:957–984, 2008.

N. Williams. Persistent Private Information. *Econometrica*, 79:1233–1274, 2011.

Noah Williams. On Dynamic Principal-Agent Models in Continuous Time. Working paper, University of Wisconsin, 2008.

Tak-Yuen Wong. Dynamic Agency and Endogenous Risk-Taking. *Management Science*, page Forthcoming, 2018.

Appendix

A Proof of Proposition 1

For $t \in [0, \infty)$, let Υ_t be the $\mathcal{F}_t$–conditional expectation of the manager's life-time utility under the owners's policy p and the manager's action policy a. Namely

$$\begin{aligned} \Upsilon_t &\equiv E^{a,p}\left[\int_0^\infty e^{-\rho s} u\left(c_s + \Delta_s\right) ds | \mathcal{F}_t\right] \\ &= \int_0^t e^{-\rho s} u\left(c_s + \Delta_s\right) ds + e^{-\rho t} q_t. \end{aligned} \tag{A.20}$$

Then $\{\Upsilon_t\}$ is a $\{\mathcal{F}_t\}$–adapted martingale. By a martingale representation theorem[8] we have:

$$d\Upsilon_t = e^{-\rho t}\gamma_t^W dW_t^a - e^{-\rho_t}\gamma_t^N dM_t^\lambda. \tag{A.21}$$

Then (A.20) and (A.21) imply (2).

B Proof of Proposition 2

We need to show that (a, p), with $a \equiv (0, \lambda)$, is incentive compatible if and only if for any $t \in [0, \infty)$

$$(0, \lambda_t) = \arg\max_{\tilde{\Delta}, \tilde{\lambda}} u\left(c_t + \tilde{\Delta}\right) - \gamma_t^W \tilde{\Delta} + \gamma_t^W \phi\left(y_t\right) \tilde{\lambda} - \gamma_t^N \tilde{\lambda}. \tag{A.22}$$

Suppose not, so that there is another $a' = (\Delta', \lambda')$ such that

$$u\left(c_t\right) + \gamma_t^W \phi\left(y_t\right) \lambda_t - \gamma_t^N \lambda_t < u\left(c_t + \tilde{\Delta}\right) - \gamma_t^W \Delta_t' + \gamma_t^W \phi\left(y_t\right) \lambda_t' - \gamma_t^N \lambda_t' \tag{A.23}$$

over a time interval with a positive measure. Let q be the manager's promised utility process under a. Define:

$$\hat{\Upsilon}_t \equiv \int_0^t e^{-\rho s} u\left(c_s + \Delta_s'\right) ds + e^{-\rho t} q_t.$$

So $\hat{\Upsilon}_t$ is the $\mathcal{F}_t$–conditional expectation of the manager's life-time utility if he acts according to a' from date zero and switches to a at t. Then:

$$e^{-\rho t} d\hat{\Upsilon} t = u\left(c_t + \Delta_t'\right) dt - \rho q_t dt + dq_t. \tag{A.24}$$

[8]See Jacod and Shiryaev [2002] Chapter III, Theorem 4.34 for a generalized description of the Martingale Representation Theorem.

Equation (2) implies:

$$dq_t - \rho q_t dt = -u\left(c_t\right)dt + \gamma_t^W \sigma\left(y_t\right) dW_t^a - \gamma_t^N dM_t^\lambda$$

and

$$\begin{aligned} dW_t^a &= dW_t^{a'} + \frac{1}{\sigma\left(y_t\right)}\left[-\phi\left(y_t\right)\left(\lambda_t - \lambda_t'\right) - \Delta_t'\right]dt, \\ dM_t^\lambda &= dM_t^{\lambda'} - \left(\lambda_t - \lambda_t'\right)dt. \end{aligned}$$

Therefore, (A.24) can be rewritten as:

$$\begin{aligned} e^{-\rho t} d\hat{\Upsilon}_t &= \left[u\left(c_t + \Delta_t'\right) - u\left(c_t\right) - \gamma_t^W \phi\left(y_t\right)\left(\lambda_t - \lambda_t'\right) - \gamma_t^W \Delta_t' + \gamma_t^N\left(\lambda_t - \lambda_t'\right)\right]dt \\ &\quad + \gamma_t^W \sigma\left(y_t\right) dW_t^{a'} - \gamma_t^N dM_t^{\lambda'}. \end{aligned} \tag{A.25}$$

Notice that the last two terms on the right hand side of (A.25) are martingale increments. Then (A.23) implies that $\left\{\hat{\Upsilon}_t\right\}$ is a sub-martingale with a negative trend over a positive measure of time. Therefore, there is a $\bar{t} \in [0, \infty)$ such that

$$E^{a',p}\left[\hat{\Upsilon}_{\bar{t}}\right] > \hat{\Upsilon}_0 = q_0.$$

Hence, adopting a' up to $\bar{t}$ and switching to a is a strictly dominant choice of the manager. So (a, p) is not incentive compatible, and we have a contradiction.

To prove the other direction, suppose that a satisfies (3) and (4) so that:

$$u\left(c_t\right) + \gamma_t^W \phi\left(y_t\right)\lambda_t - \gamma_t^N \lambda_t \geq u\left(c_t + \tilde{\Delta}\right) - \gamma_t^W \Delta_t' + \gamma_t^W \phi\left(y_t\right)\lambda_t' - \gamma_t^N \lambda_t'$$

for any a'. Then $\left\{\hat{\Upsilon}_t\right\}$ is a super-martingale and:

$$E^{a',p}\left[\hat{\Upsilon}_\infty\right] \leq \hat{\Upsilon}_0 = q_0.$$

So a dominates a' and we have the desired result.

C Proof of Lemma 1

We define a function $\tilde{c}(c) \equiv (-q)^{-1}\exp\left(-\theta_A c\right)$ characterizing the relationship between $\bar{c}^M$ and c. Notice that J^M is negative and then the optimal c solves the following problem:

$$\begin{aligned} \min_{\tilde{c}} \Psi_M\left(\tilde{c}\right) &\equiv \mu\tilde{c} - \frac{1}{\theta_A}(-q)^{-1}\left(-\exp\left(-\theta_A\tilde{c}\right)\right) - \sigma^2\theta_P\mu(-q)^{-1}\left(\exp\left(-\theta_A\tilde{c}\right)\right) \\ &\quad + \frac{1}{2}\sigma^2\left(\theta_P + \theta_A\right)\left((-q)^{-1}\exp\left(-\theta_A\tilde{c}\right)\right)^2. \end{aligned}$$

Notice that:

$$\Psi_M'(c) = \mu - \tilde{c}(c) + \sigma^2\theta_A\theta_P\mu\tilde{c}(c) - \sigma^2(\theta_P+\theta_A))\theta_A\tilde{c}(c) \tag{A.26}$$

Since $\tilde{c}'(c) < 0$, the first-order condition implies:

$$\mu - \tilde{c}(c) + \sigma^2\theta_A\theta_P\mu\tilde{c}(c) - \sigma^2(\theta_P+\theta_A))\theta_A\tilde{c}(c) = 0$$

and $\tilde{c}(c)$ must be a root of the quadratic function:

$$f(x) = \mu + \left(\sigma^2\theta_A\theta_P\mu - 1\right)x - \sigma^2(\theta_P+\theta_A))\theta_A x^2, \tag{A.27}$$

which has a positive root and a negative one because $f(0) = \mu > 0$ and the coefficient of x^2 is negative. According to the definition of $\tilde{c}(c)$, we must choose the positive root. So we check the second order condition. Notice that:

$$\Psi_M''(c) = f'(\tilde{c}(c))\underbrace{\tilde{c}'(c)}_{<0}$$

and $f'(\tilde{c}(c)) < 0$ as $\tilde{c}(c)$ is the larger root of $f(x)$, and we have the desired result.

To prove the result on the sensitivities γ^W, note that from the expression for $\bar{c}^M$ in the text we have:

$$\begin{aligned}
(\theta_A+\theta_P)\bar{c}^M &= -\frac{1}{2\sigma^2\theta_A} + \frac{\theta_P\mu}{2} + \frac{\sqrt{\left(\sigma^2\theta_A\theta_P\mu+1\right)^2 + 4\mu\sigma^2\theta_A^2}}{2\sigma^2\theta_A} \\
&\geq -\frac{1}{2\sigma^2\theta_A} + \frac{\theta_P\mu}{2} + \frac{\sqrt{\left(\sigma^2\theta_A\theta_P\mu+1\right)^2}}{2\sigma^2\theta_A} \\
&= \theta_P\mu.
\end{aligned}$$

Therefore $\bar{c}^M \geq \frac{\theta_P}{\theta_A+\theta_P}\mu$ and so:

$$\gamma^{W,M} = -\theta_A\bar{c}^M q \geq -\theta_A\frac{\theta_P}{\theta_A+\theta_P}\mu q = \gamma^{W,F}.$$

D Proof of Proposition 4

We only need to consider the case in which (12) is satisfied with strict inequality. As above, we define the function $\tilde{c}(c) \equiv (-q)^{-1}\exp(-\theta_A c)$ characterizing the relationship between $\bar{c}^M$ and c. Notice that J^R is negative and then the optimal c solves the following problem:

$$\begin{aligned}
\min_{\tilde{c}} \Psi_R(\tilde{c}) \equiv\ & \mu\tilde{c} - \frac{1}{\theta_A}(-q)^{-1}(-\exp(-\theta_A\tilde{c})) - \sigma^2\theta_P\mu(-q)^{-1}(\exp(-\theta_A\tilde{c})) \\
& + \frac{1}{2}\sigma^2(\theta_P+\theta_A)\left((-q)^{-1}\exp(-\theta_A\tilde{c})\right)^2 \\
& + \underline{\lambda}\left[\phi(-q)^{-1}\exp(-\theta_A c) + \frac{1}{\theta_P}\exp(\theta_P\mu\phi)\left(1+\phi\theta_A i(-q)^{-1}\exp(-\theta_A c)\right)^{-\frac{\theta_P}{\theta_A}}\right].
\end{aligned}$$

Let $f(\cdot)$ be the quadratic function defined in (A.27) and define the function $g(\cdot)$ such that:

$$g(x) = \underline{\lambda}\theta_A \phi x \left[1 - \exp\left(\theta_P \mu \phi\right)\left(1 + \phi\theta_A x\right)^{-\frac{\theta_A + \theta_A}{\theta_A}}\right].$$

The first order condition for the optimization above is then:

$$\Psi_R'\left(c\right) = f(\tilde{c}\left(c\right)) - g(\tilde{c}\left(c\right)) = 0$$

or:

$$f\left(\bar{c}^R\right) = g\left(\bar{c}^R\right). \tag{A.28}$$

According to its definition, $\bar{c}^R > 0$ and we focus on the positive real line. Notice that:

$$\lim_{x \downarrow 0} f(x) > 0,\ \lim_{x \to \infty} f(x) = -\infty,\ \lim_{x \downarrow 0} g(x) = 0, \text{ and } \lim_{x \to \infty} g(x) = \infty.$$

Thus there must be a positive value of $\bar{c}^R$ satisfying (A.28). We assume that there is a unique such solution, as the extension of the argument to the case with multiple solutions is straightforward. Since $f(x) < g(x)$ if $x \in \left(0, \bar{c}^R\right)$ and $f(x) > g(x)$ if $x \in \left(\bar{c}^R, \infty\right)$, then $f'\left(\bar{c}^R\right) - g'\left(\bar{c}^R\right) < 0$ and:

$$\Psi_R''(c) = \left[f'(\tilde{c}\left(c\right)) - g'(\tilde{c}\left(c\right))\right]\tilde{c}'(c) > 0$$

so that the second order condition is satisfied. According to its definition:

$$\begin{aligned} g'(x) &= \underbrace{\underline{\lambda}\theta_A\left[1 - \exp\left(\theta_P\mu\phi\right)\left(1 + \phi\theta_A x\right)^{-\frac{\theta_P + \theta_A}{\theta_A}}\right]}_{\geq 0 \text{ if and only if } g(x) \geq 0} \\ &\quad \underbrace{+ \underline{\lambda}\theta_A x \frac{\theta_P + \theta_A}{\theta_A} \exp\left(\theta_P\mu\phi\right)\left(1 + \phi\theta_A x\right)^{-\frac{\theta_P + \theta_A}{\theta_A} - 1}}_{> 0}. \end{aligned}$$

So $\lim_{x \downarrow 0} g'(x) = \underline{\lambda}\theta_A\left[1 - \exp\left(\theta_P\mu\phi\right)\right] < 0$ and then continuity implies that $g(x) < 0$ for x close to zero. Furthermore, $g'(x) > 0$ if $x \geq 0$. Therefore there is a unique solution for $g(x) = 0$ which is $\bar{c}_0^M$ defined in (12).

The fact that $f(0) > 0$ and the coefficient of the quadratic term of $f(\cdot)$ is negative implies $f(x) > 0$ for all $x \in (0, \bar{c}^M)$. Therefore $\bar{c}_0^M < \bar{c}^M$ implies that $f\left(\bar{c}_0^M\right) > 0 = g\left(\bar{c}_0^M\right)$ and then:

$$0 < \bar{c}_0^M < \bar{c}^R < \bar{c}^M.$$

Since:

$$\bar{\gamma}^{N,R} = -\phi\theta_A\bar{c}^R \text{ and } \bar{\gamma}^{N,M} = -\phi\theta_A\bar{c}_0^M$$

the above equation implies $\bar{\gamma}^{N,R} < \bar{\gamma}^{N,M}$ and we have the desired result.