

ACQUISITION OF INFORMATION TO DIVERSIFY CONTRACTUAL RISK*

BY BOĞAÇHAN ÇELEN AND SALTUK ÖZERTÜRK¹

Columbia University, U.S.A.; Southern Methodist University, U.S.A.

Are hedging transactions that diversify a manager's compensation risk detrimental to incentives, or can they improve contracting efficiency? If hedging provides efficiency benefits, should the manager or the firm undertake it? In our model, both the firm and the manager can trade financial portfolios to diversify the manager's compensation risk. Prior to the portfolio selection, the parties need to acquire information on how different financial portfolios fit their diversification purposes. We illustrate that financial portfolios correlated with firm-specific risk improve contracting efficiency. For equal information costs, it is optimal for the firm to undertake the hedging on the manager's behalf.

1. INTRODUCTION

The theory of optimal managerial compensation suggests that firms design their managers' compensation packages in order to provide incentives to maximize the firm's value. Therefore, firms create stock-based compensation schemes to align the managers' incentives with those of the shareholders' by creating a correlation between the managers' wealth and the firm's value. The stock-based compensation schemes introduce risk into the managers' wealth, which in turn creates incentives for poorly diversified and risk-averse managers to exploit opportunities to diversify the risks associated with their compensation.

Exploring the idea that risk-averse managers seek diversification, Stulz (1984) proposes managerial risk aversion as a driver of corporate risk management. By engaging in corporate hedging activities, managers can reduce the variance of their firm's value and obtain some insurance for their compensation schemes. Stulz's (1984) theory of corporate risk management has been criticized by Froot et al. (1993), who argue that, instead of employing corporate hedging practices, the managers can go to the corresponding markets themselves to obtain such insurance. They suggest that Stulz's theory implicitly assumes that managers face significant costs when trading hedging contracts for their own account, as otherwise they would be able to diversify their risks without having to involve their firms directly in any hedging activities.

Indeed, there is ample evidence that suggests that corporate executives do engage in personal hedging transactions in the financial markets to diversify the risk associated with their stock ownership positions. Bettis et al. (2001) document the increasing popularity of sophisticated financial instruments such as equity swaps and zero-cost collars that enable corporate executives to hedge their compensation risks. Similar hedging practices undertaken on managers' personal accounts have also been documented in the legal profession by Bank (1995), Schizer (2000), and Easterbrook (2002). The general view on these managerial hedging transactions has been quite negative. It is typically argued that if the managers have unrestricted access to financial markets to trade on their own accounts, they will hedge away the performance incentives in their

*Manuscript received February 2008; revised March 2009.

¹We are grateful to Onur Özgür and Paolo Siconolfi for numerous discussions. We also thank two anonymous referees, Patrick Bolton, Piero Gottardi, Michael Riordan, Tano Santos, and the participants in the Finance Lunch Seminar at Columbia Business School for valuable comments. Please address correspondence to: Boğaçhan Çelen, Graduate School of Business, Columbia University, 3022 Broadway, 602 Uris Hall, New York, NY 10027. Phone: 212-854 7209. E-mail: bc2132@columbia.edu.

compensation schemes, rendering the incentive justification for managerial stock ownership invalid.²

The above observations suggest that, whereas Froot et al. (1993) criticize Stulz's (1984) theory of corporate risk management on the grounds that the managers can obtain the same insurance by personally trading in the financial markets, the actual managerial hedging trades undertaken in practice are criticized in the business press and the legal profession on the basis that the managers can use these trades to unwind their performance incentives. Given the controversy surrounding managerial hedging, it seems important to understand the implications of these trades for contracting efficiency and the optimal way to implement them. In this article, we address the following questions in an optimal contracting framework. First, we tackle the question of whether hedging transactions that diversify a manager's compensation risk are necessarily detrimental to incentives, or whether they can improve contracting efficiency by reducing the risks that the manager faces without hurting incentives. The second question is: if there are efficiency benefits from hedging the manager's compensation risk, which party—the manager or the firm on the manager's behalf—should undertake the hedging transaction?

To address these questions, we consider an extension of the principal-agent model that allows both the manager and the firm to *trade portfolios correlated with the firm-specific risk*. In particular, as well as tying the manager's compensation to firm value, the firm can choose to trade portfolios correlated with firm-specific risk and offer the manager a compensation contract that is also tied to such portfolios. Upon receiving the contract, the manager can also trade portfolios in the financial market. A key feature in our model is that, although both the manager and the firm have access to financial markets, access to such markets is subject to a costly information acquisition activity. We explore a model where available financial portfolios differ in their effectiveness for diversifying the firm-specific risk, which is measured by the correlation of the portfolio payoff with that risk. Neither the manager nor the firm has complete information about how the available portfolios in the market are correlated with the firm-specific risk. By sampling from the set of available financial portfolios, which we model as a search process, the party that undertakes the hedging can learn the correlation of a given portfolio and decide whether to hold a position in that portfolio or continue searching. Our main results are as follows.

We show that regardless of which party—the manager or the firm—undertakes the hedging transactions, the availability of financial portfolios correlated with firm-specific risk yields a more efficient contracting outcome and elicits higher managerial effort. The ability to diversify more risk from the manager's compensation contract makes the incentive motives more dominant than the insurance consideration in determining the optimal compensation scheme, which results in higher pay-performance sensitivity and higher managerial effort.

On the question of which party should undertake this beneficial hedging activity, we show that there is no equilibrium in which both parties engage in hedging. Furthermore, and perhaps more surprisingly, for equal information acquisition/search costs, it is always optimal for the firm to undertake the search and do the hedging on the manager's behalf. One possible explanation for this result is as follows: The search process introduces another source of randomness into the manager's wealth, namely, the uncertain search cost to achieve a given level of diversification. If the risk-averse manager undertakes the search in financial markets, then he needs to be compensated *ex ante* with an extra risk premium to undertake the desired level of search and diversification. The alternative case, in which the risk-neutral firm undertakes the search, does not cost the firm this extra risk premium, as the contract that can then be offered to the manager involves no uncertainty with respect to the cost of any subsequent hedging activity. As a result, to achieve any desired level of diversification, it is always less costly for the firm to undertake the search itself, instead of letting the risk-averse manager do the search and compensating him for the expected search cost.

² For example, an editorial in *The Economist* ("Executive Relief," April 3, 1999, p. 64) states that "Further justifying scepticism is the current popularity of derivatives that allow managers to hedge their exposure to their own company's shares . . . Such hedging is wholly against the spirit of the massive awards of shares and share options in recent years." (See also Lavalley, 2001.)

1.1. *Literature on Managerial Hedging.* In related earlier work on managerial hedging, Garvey (1993) and Bisin et al. (2008) allow the manager to trade side contracts contingent on his own firm value. They show that such hedging transactions undermine incentive provision.³ Although we describe a hedging transaction that improves, instead of undermines, the efficiency of incentive provision, we should point out that this difference crucially depends on the particular hedging transaction that we consider in our model, namely, trade in *portfolios correlated with the firm-specific risk*.

A portfolio correlated with the firm-specific risk reduces the randomness in the manager's wealth but preserves the link between the manager's wealth distribution and his subsequent choice of effort. In contrast, the hedging transactions contingent on the value of the firm, as analyzed in Garvey (1993) and Bisin et al. (2008), serve to undo the link between the manager's effort choice and his wealth.⁴ With these trades, the manager simply promises his share of firm value to third parties in exchange for a fixed payment and unwinds the effort incentives provided by his stock ownership. Although the availability of financial portfolios correlated with the firm-specific risk reduces the noise in the manager's wealth distribution for a given effort choice, the contracts contingent on the value of the firm serve to sever the link between a given effort choice and managerial wealth, undermining incentive provision.⁵

1.2. *Literature on Corporate Risk Management.* As mentioned, Stulz (1984) views corporate risk management as an outgrowth of managerial risk aversion. According to Stulz (1984), although the company shareholders do have access to trading opportunities to diversify risks related to their share ownership, the company manager, who holds a relatively large portion of his wealth in the firm's stock, lacks such trading opportunities and hence is poorly diversified against the same risk. Therefore, the manager may use corporate risk management activity to diversify his own risk.⁶ However, as Froot et al. (1993) and Holmström and Tirole (2000) convincingly argue, Stulz's (1984) theory seems to implicitly rely on a transaction cost differential between the manager and the firm. Therefore it fails to address why the same diversification outcome cannot be attained (i) by the manager directly using the financial markets to diversify the risk exposure through trades on his own account or (ii) by the firm simply building a risk-reducing filter directly into the manager's contract.

These points indicate that for a better understanding of the relative merits of hedging by the manager or by the firm, one needs to analyze a setting in which the respective hedging costs of the two parties are explicitly modeled. This is the key point that we aim to capture by introducing friction into the financial market access of the manager and the firm in our costly search framework. By explicitly allowing both the manager and the firm to have costly access to financial markets for hedging purposes, we provide a more complete model for analyzing

³ In a similar vein, our earlier paper, Çelen and Özertürk (2007), shows that a manager's ability to trade nonexclusive swap contracts (promising the return from his shares to third parties in exchange for a fixed payment) can undermine incentives completely.

⁴ The nonexclusive contracting literature in which an agent can contract with *multiple principals* for the same moral-hazard activity also typically reports that the agent's unobservable side trades undermine contracting efficiency (see, for example, Kahn and Mookherjee, 1995, 1998; Bisin and Guaitoli, 2004, on the implications for insurance and credit markets, and Parlour and Rajan, 2001, on the implications for loan contracts from multiple creditors).

⁵ Other related contributions with single principal/single agent frameworks include Rogerson (1985) and Park (2004). In a two-period principal-agent model, Rogerson (1985) shows that the optimal contract if the agent has no access to riskless saving and borrowing will leave the agent with a precautionary demand for saving. Riskless saving benefits a risk-averse agent by providing partial insurance against future wage uncertainty, but this insurance weakens incentives and lowers the equilibrium effort that is optimally elicited. Park (2004) allows the agent to privately borrow and save in the *precontracting stage* and shows that the agent's precontracting access to private credit markets results in a severe loss of incentive provision.

⁶ In an empirical study, Tufano (1996) examines corporate risk management activity in the North American gold mining industry and shows that firms whose managers hold more stock-based compensation manage more gold price risk. His study concludes that risk-reduction policies may be set to satisfy the needs of poorly diversified managers.

some issues raised in the corporate risk management and managerial hedging literatures.⁷ Our results illustrate that hedging may enhance contracting efficiency by filtering out the noise from the manager's contract, and even when the search costs of the two parties are equal, the firm, not the manager, should undertake the hedging activity on the manager's behalf.

The plan of the article is as follows: The next section describes the basic model. In Subsection 3.1, we focus on the case where there is only one portfolio correlated with the firm-specific risk, in order to illustrate how the manager's or firm's ability to trade such portfolio improves contracting efficiency. In Subsection 3.2, we consider the general case where there is not only one, but a variety of financial portfolios, and describe the manager's and the firm's optimal search behavior in the financial market. In that subsection, we show that for the same unit search costs, it is always optimal for the firm (but not the manager) to undertake the search and do the hedging. Section 4 provides a discussion of our results and our conclusion. The Appendix contains the proofs that are not presented in the text.

2. THE MODEL

The basic model is based on the commonly used CARA-normal principal-agent model with linear contracts. It is detailed below.

2.1. Technology and Preferences. A risk-neutral principal (whom we refer to as *she*) hires a risk-averse agent (whom we refer to as *he*). The agent chooses unobservable effort, $e \in [\underline{e}, \bar{e}] \subset \mathbb{R}$, which, together with the realization of a random shock ϵ , determines the final output (firm value) x :

$$x := f(e) + \epsilon.$$

The function $f : [\underline{e}, \bar{e}] \mapsto \mathbb{R}_+$ is nondecreasing, concave, and differentiable and represents the productivity of effort. We assume that the idiosyncratic output shock is a random variable \mathbf{E} , normally distributed with mean 0 and variance σ_ϵ^2 . Therefore, given e , output is a random variable \mathbf{X}_e induced by \mathbf{E} , normally distributed with mean $f(e)$ and variance σ_ϵ^2 . Supplying effort is costly for the agent/manager. The cost of effort is determined by a nondecreasing, convex, differentiable function

$$c : [\underline{e}, \bar{e}] \mapsto \mathbb{R}_+.$$

The risk-averse agent has constant absolute risk aversion (CARA) preferences, represented by the utility function $u : \mathbb{R}_+ \mapsto \mathbb{R}$ specified as

$$u(w) := -\exp\{-\alpha w\},$$

where the parameter $\alpha > 0$ is the agent's coefficient of absolute risk aversion and w is his final wealth.

⁷ In our framework, the sole reason that the firm is interested in engaging in the risk management activity is to improve contracting efficiency by reducing the noise in the risk-averse manager's contract. The literature has also provided different explanations for corporate risk management in addition to the managerial risk aversion rationale of Stulz (1984, 1996). In Breeden and Viswanathan (1998), it is argued that some managers may undertake hedging to influence the labor market's perception of their ability. Smith and Stulz (1985) argue that if taxes are a convex function of earnings, it will generally be optimal for firms to hedge. Also in Smith and Stulz (1985), it is argued that hedging may be used as a way to increase debt capacity and avoid costly bankruptcy. Stulz (1990) proposes that hedging can add value by reducing the distortions associated with debt finance. In a model where external financing is more costly than internally generated funds, Froot et al. (1993) show that corporate hedging adds value to the extent that it helps to ensure that the company has sufficient internal funds to take advantage of attractive investment opportunities.

2.2. Linear Compensation. The principal's objective is to maximize the expected output net of the agent's compensation. Since the agent's effort choice is unobservable, the principal cannot compensate the agent contingent on his effort. However, the principal can provide effort incentives by tying the agent's compensation to the realized output. Specifically, we assume that the agent's compensation scheme takes the linear form $sx + t$, where t is a fixed payment and s is the agent's share of the final output x . In what follows, we denote this contract by (s, t) and we refer to s as the *pay-performance sensitivity* of the agent's compensation scheme.⁸

We should note that, by construction, the contract (s, t) , which is only contingent on output x , cannot make use of all possible available information about the agent's effort decision. It is well known in the literature that, since the agent is risk averse, the optimal second-best contract reflects a trade-off between insuring the agent and providing incentives for efficient effort choice. Since the best possible contract that is only contingent on output x induces the risk-averse agent to supply less than the first-best level of effort, the principal may be able to improve the contract by making pay contingent on a second variable, say y , in addition to output x . Instead of relying only on x , the principal can make the contract contingent on other available information and better assess the agent's effort choice by filtering out some of the noise in the compensation contract. Indeed, as Holmström (1979) has shown, the optimal contract should be made contingent not only on x , but also on y , whenever y contains incremental information in assessing the agent's unobservable effort.⁹ More precisely, the optimal contract will depend on y as well if x is not a sufficient statistic for the pair (x, y) with respect to the agent's effort. Since we focus on an initial contract contingent only on x , effectively our contract (s, t) does not satisfy Holmström's (1979) sufficient statistic condition.¹⁰ In practice, however, there are few such contingencies; therefore, considering an initial contract only contingent on \mathbf{X} seems to be a reasonable, yet critical, assumption. This critical assumption implies that the contract (s, t) offered to the agent in our setting does not filter out the noise associated with \mathbf{E} and hence leaves the risk-averse agent with a diversification incentive.¹¹ Accordingly, hedging transactions undertaken by the principal or the agent to reduce the randomness in the agent's initial contract may involve efficiency benefits. Later, we describe in detail the particular way we construct these hedging transactions.

2.3. Trading in the Financial Market. Let us assume that both the agent and the principal have access to a financial market where they can trade portfolios that are correlated with the

⁸ Linear contracts have been highly attractive in the literature due to their tractability and the intuitive solution they deliver. Holmström and Milgrom (1987) show that linear contracts are optimal in a dynamic principal-agent setting with *CARA* preferences and with a binomial output process. This result has provided the primary justification for restricting attention to linear contracts in a variety of applications. Although linear contracts are not necessarily optimal in the present model, we restrict attention to them due to their tractability, as many other papers in the literature do. For general treatments of the principal-agent problem, we refer the reader to Ross (1973), Harris and Raviv (1979), Holmström (1979), and Grossman and Hart (1983).

⁹ See Gibbons and Murphy (1990) for an excellent discussion on the optimality of relative performance evaluation contracts.

¹⁰ In a seminal paper on relative performance evaluation, Holmström (1982) has shown that with n workers who all face a common shock as well as an idiosyncratic shock, for the purpose of efficient risk sharing the optimal contract for the i th worker must be contingent not only on his own output x_i , but also on the average output \bar{x} of all workers. Typically, the wage paid to the i th worker will be positively related to his own output x_i but negatively related to the average output \bar{x} of all workers subject to the same common shock. Holmström (1982) also showed that as the number of workers becomes very large, the optimal contract is contingent on $x_i - \bar{x}$; that is, the performance of the i th worker is measured relative to the performance of others in the reference group.

¹¹ Similarly, when the agent's output is also subject to a common and observable market-wide shock, it is optimal to filter out the noise associated with this common shock by measuring the agent's performance relative to a market index. Despite what agency theory would predict, Garvey and Milbourn (2003) argue that there is very little evidence that firms index executive compensation to aggregate market variables (relative performance evaluation) to remove market-wide risks from compensation schemes. They point out that executives can trade market indices themselves, and hence the lack of relative performance evaluation might be explained by an executive's own ability to remove market-wide risks by trading in the financial markets.

firm-specific risk. Therefore, in principle, both parties can engage in hedging transactions that can help to diversify the risk in the agent's compensation contract (s, t) .

There is a variety of financial portfolios that the agent and the principal can potentially trade. The payoff from a generic portfolio is represented by a random variable \mathbf{Y} , which is assumed to be distributed normally with mean μ and variance σ_y^2 . Each \mathbf{Y} is correlated with the output shock \mathbf{E} according to $\text{cov}(\mathbf{Y}, \mathbf{E})$; therefore, each \mathbf{Y} serves as a potential instrument for diversifying \mathbf{E} from the agent's compensation scheme. The degree of diversification that a portfolio provides, however, will depend on the correlation of its payoff with \mathbf{E} . A portfolio that is highly correlated with \mathbf{E} —either positively or negatively—and with a low standard deviation σ_y is clearly more desirable for hedging purposes. To capture this, we define

$$q := \frac{\text{cov}(\mathbf{Y}, \mathbf{E})}{\sigma_y},$$

and refer to $q^2 =: z$ as the quality of a portfolio: A portfolio with a higher z is better for diversification purposes.

Should the party undertaking the hedging transaction have perfect knowledge of all available portfolios, then (s)he would choose the portfolio with the highest z in order to diversify as much risk as possible. In our framework, however, access to the financial market is subject to imperfections that we model by using a costly information acquisition (search) framework. A key feature in our model is that, when they attempt to diversify the contractual risk that is specific to the firm, the agent or the principal is uncertain about how different financial portfolios are correlated with this particular risk. Accordingly, if they trade in the financial market they need to engage in costly information acquisition and learn how different portfolio returns are correlated with the firm-specific risk.

Each available portfolio has a different quality, and the two parties know only the distribution of qualities, but not the quality of a particular portfolio. The set of financial portfolios can be described as the collection of all available qs . More precisely, there is a random variable \mathbf{Q} that describes the quality of all available portfolios. Let us denote the principal by p and the agent by a . Any party can search for a portfolio from the distribution of all qs . By sampling a portfolio from \mathbf{Q} , a party can find out its quality—i.e., its correlation with \mathbf{E} . If the party $i = a, p$ undertaking the search is “satisfied” with the quality of a portfolio, then (s)he stops and creates a position $d_i \in \mathbb{R}$ in that portfolio; otherwise, (s)he pays a cost $\kappa_i > 0$ and samples another portfolio, and so on, until deciding to stop the search process.

If the principal does a search and creates a position $d_p \in \mathbb{R}$ in the portfolio at which she stopped, then subsequently she offers the agent a contract $s\mathbf{X} + d_p\mathbf{Y} + t$. If the agent decides to search further, then he creates a position $d_a \in \mathbb{R}$ in the portfolio at which he stopped, and he is entitled to a claim $d_a\mathbf{Y}$. We should note that even if the principal undertakes the search and trades a portfolio, this does not preclude the agent from undertaking another search; that is, in principle both the principal and the agent can undertake separate hedging transactions. Finally, we assume that the participants of the financial market are risk neutral and competitive, so that each share of the portfolio \mathbf{Y} is priced at $E[\mathbf{Y}] = \mu$.¹²

For convenience, Figure 1 summarizes the sequential form of the game and the notation: First, the principal decides whether to undertake any search and trading in the financial market. If the principal does the search at a unit search cost $\kappa_p > 0$ and creates a position $d_p \in \mathbb{R}$ in a portfolio \mathbf{Y} , then she offers the agent a compensation scheme of the form $s\mathbf{X} + d_p\mathbf{Y} + t$. If the principal chooses not to do any search and trading, the agent receives a contract of the form $s\mathbf{X} + t$. If the agent accepts the contract, then he decides whether or not to undertake another search in the financial market. If the agent decides to trade, he can sample portfolios at a unit search cost $\kappa_a > 0$ and can create a position $d_a \in \mathbb{R}$ in the portfolio at which he stopped. Alternatively, the agent can choose not to do another search. Subsequently, the agent chooses

¹² Therefore, the party i who creates a position d_i pays μd_i .

- ① The principal decides whether to search a portfolio \mathbf{Y} at a unit cost $\kappa_p > 0$. If she searches, she takes a position $d_p \in \mathbb{R}$ at a price μ and offers the agent a contract $s\mathbf{X} + d_p\mathbf{Y} + t$. If she does not search, she offers a contract $s\mathbf{X} + t$.
- ② The agent decides whether to accept the contract or not. If the agent does not accept the contract, both parties opt out for their outside options. If the agent accepts, the relationship moves to the next stage.
- ③ The agent can further search a portfolio \mathbf{Y} at a unit cost $\kappa_a > 0$. If he searches, he takes a position $d_a \in \mathbb{R}$ at a price μ
- ④ The agent chooses costly effort $e \in [\underline{e}, \bar{e}]$.
- ⑤ The output, x , and the portfolio payoffs, y , are realized.

FIGURE 1

TIMING OF THE GAME

his effort level, e . Finally, the firm value, x , and the portfolio payoffs, y , are realized. Figure 1 depicts the timing and the notation of the game.

In the following section, we will analyze some features of the subgame-perfect equilibrium of the game.

3. ANALYSIS

3.1. *Single Portfolio.* To illustrate how the availability of financial portfolios correlated with firm-specific risk improves contracting efficiency, we begin our analysis by focusing on the problem when *only one portfolio is available* and *only the agent can trade this portfolio*. This exercise provides us with a benchmark that simplifies the analysis of the case where the principal can also trade the portfolio. We analyze the general case with a financial market later in Subsection 3.2.

Suppose there is only one portfolio \mathbf{Y} with a given quality q . Furthermore, *the quality of the portfolio is common knowledge*. Also assume that only the agent can trade portfolio \mathbf{Y} . Since, by assumption, the principal cannot trade, the agent receives a compensation contract (s, t) . If the agent chooses a position $d_a \in \mathbb{R}$ in the portfolio and expends an effort level e , his wealth can be written as a random variable

$$\mathbf{W}_{e,d_a} = s\mathbf{X}_e + d_a\mathbf{Y} + t - c(e) - \mu d_a.$$

The CARA preferences and the normality assumptions on \mathbf{E} and \mathbf{Y} imply that the agent's problem of choosing e to maximize $E(u(\mathbf{W}_{e,d_a}))$ is equivalent to maximizing the *certainty equivalent wealth*¹³

$$\begin{aligned} w(e, d_a) &:= E(\mathbf{W}_{e,d_a}) - \frac{\alpha}{2}\text{Var}(\mathbf{W}_{e,d_a}) \\ &= sf(e) + t - c(e) - \frac{\alpha}{2}(s^2\sigma_\epsilon^2 + d_a^2\sigma_y^2 + 2sd_aq\sigma_y). \end{aligned}$$

Notice that $w(e, d_a)$ is concave, differentiable, and nondecreasing in $e \in [\underline{e}, \bar{e}]$. Furthermore, the choice of effort only affects the expected wealth, whereas the choice of a position d_a only affects the variance of the agent's wealth distribution. The optimal effort e^* satisfies one of the

¹³ To see this, it is sufficient to check the joint moment-generating function of \mathbf{X}_e and \mathbf{Y} .

following two conditions:

$$(1) \quad \begin{aligned} e^* &\in (e, \bar{e}) \text{ and } sf(e^*) - c'(e^*) = 0, \\ e^* &= \bar{e} \text{ and } sf'(e^*) - c'(e^*) \geq 0. \end{aligned}$$

The above expression for the optimal effort choice is standard: As the agent is given a higher pay–performance sensitivity s , he expends more effort.

Turning to the optimal choice of a position in portfolio \mathbf{Y} , notice that $w(e^*, d_a)$ is also concave, differentiable, and nonmonotonic in d_a . Therefore, straightforward maximization yields the optimal position in \mathbf{Y} as

$$(2) \quad d_a^*(s) = -\frac{q}{\sigma_y} s.$$

The agent's optimal position in \mathbf{Y} increases in the pay–performance sensitivity s of his compensation scheme and the quality q of the portfolio. The sign of q determines whether the optimal position involves buying or short-selling the portfolio.

By using (1) and (2), we can illustrate how the optimal position in \mathbf{Y} provides diversification. At the optimal e^* and d_a^* , the agent's certainty equivalent wealth becomes

$$w(e^*, d_a^*) = sf(e^*) + t - c(e^*) - \frac{\alpha}{2} s^2 (\sigma_\epsilon^2 - z).$$

A position d_a^* in a portfolio with quality q reduces the variance of the agent's wealth distribution to $s^2(\sigma_\epsilon^2 - z)$ and hence lowers the associated premium that the agent demands for bearing risk. From the principal's perspective, this diversification decreases the insurance cost of eliciting a given effort level. Accordingly, the agent's equilibrium pay–performance sensitivity s , and hence the equilibrium effort that is optimally elicited should be increasing in the quality q of the portfolio available. We verify this prediction next.

3.1.1. Optimal pay–performance sensitivity. Taking into account the agent's optimal portfolio and effort choices, the principal chooses (s, t) to maximize the expected firm value net of the agent's compensation, which is given by

$$(1 - s)E(\mathbf{X}_{e^*}) - t,$$

subject to agent's participation constraint $u(w(e^*, d_a^*)) \geq u(\bar{w})$, where \bar{w} is the agent's reservation certainty equivalent wealth. We can write this participation constraint in terms of certainty equivalent wealth as

$$sf(e^*) + t - c(e^*) - \frac{\alpha}{2} s^2 (\sigma_\epsilon^2 - z) \geq \bar{w}.$$

At the optimal contract, this participation constraint holds at equality. Solving for t and substituting it into the principal's objective function, one can describe the problem as choosing s to maximize the net expected surplus

$$f(e^*) - c(e^*) - \frac{\alpha}{2} s^2 (\sigma_\epsilon^2 - z) - \bar{w}.$$

Since the agent's trade in a portfolio \mathbf{Y} reduces the premium that must be paid to the agent for bearing firm-specific risk, the expected net surplus in the contractual relationship increases

in the quality z of the portfolio. The optimal pay–performance sensitivity s^* that maximizes the above net expected surplus satisfies one of the following two conditions:

$$s^* \in (0, \bar{s}) \text{ and } (1 - s)f'(e^*) \frac{de^*}{ds} = \alpha s(\sigma_\epsilon^2 - z),$$

$$s^* = \bar{s} \text{ and } (1 - s)f'(e^*) \frac{de^*}{ds} \geq \alpha s(\sigma_\epsilon^2 - z),$$

where $e^*(\bar{s}) = \bar{e}$ and $\bar{s} = \min\{\bar{s}, 1\}$. This condition equates the marginal benefit of increasing s (which works through eliciting higher effort) to the marginal cost of this incentive provision given by $\alpha s(\sigma_\epsilon^2 - z)$. Using the above condition, one can show that the optimal s^* , and hence $e^*(s^*)$, increases in the quality of the portfolio. We summarize this result below.

PROPOSITION 1. *Suppose there is only one portfolio whose quality is known and only the agent can trade this portfolio. The optimal pay–performance sensitivity that the principal sets and the equilibrium effort that is optimally elicited increases in the quality of the portfolio.*

Now we assume that in addition to the agent, the principal can also trade in the financial market. The agent’s trade in a correlated portfolio \mathbf{Y} is desirable for the principal, since it creates efficiency benefits by reducing the insurance cost of incentive provision. Suppose now that the principal can also trade the correlated portfolio as described in Figure 1 and base the initial contract on the portfolio payoff \mathbf{Y} as well as the firm value \mathbf{X} .

We now show that, when the principal can also trade, the same efficiency level is achieved in terms of the optimal effort elicited. In particular, suppose the principal could offer a contract (s, d_p, t) that provides a wealth distribution $s\mathbf{X}_e + d_p\mathbf{Y} + t$ to the agent. It is easy to see that, for a given contract (s, d_p, t) , the agent will choose the optimal position $d_a^*(s) - d_p$, where $d_a^*(s)$ is as defined in (2). In other words, the agent will adjust his optimal position in the portfolio \mathbf{Y} in such a way that the optimal level $-qs/\sigma_y$ is achieved. In this case, any pair (d_a, d_p) is a part of the equilibrium as long as

$$d_p + d_a = -\frac{q}{\sigma_y}s.$$

As a result, if the principal could also trade the correlated portfolio, the optimal pay–performance sensitivity and the equilibrium effort would be exactly the same as described in Proposition 1.

In order to understand why the present hedging framework delivers a positive result on contracting efficiency note that the particular transaction we consider does not change the agent’s exposure to the firm’s value \mathbf{X} in his overall wealth distribution. The agent trades a portfolio that is correlated with \mathbf{X} , but the portfolio does not include any asset that can be described as a function of \mathbf{X} . Hence, for a given effort level, the firm’s value \mathbf{X} is not affected by the hedging instrument \mathbf{Y} . Put differently, the agent’s optimal effort choice does not change because \mathbf{Y} becomes part of his wealth. Consequently, the hedging activity reduces the risk that the agent is exposed to—without changing his optimal effort behavior—which in turn lowers the associated premium that he demands for bearing risk. As a result, the hedging activity serves to expand the contracting space for the principal. This enables the principal to make the agent more sensitive to output and to elicit more effort from the agent. This is how the access to financial markets leads to a welfare improvement in the contractual relationship.

Previous papers, such as Garvey (1993) and Bisin et al. (2008), have considered hedging transactions in which the agent can trade side contracts based on his own firm value \mathbf{X} . Thus, in contrast with the hedging transaction we consider, the agent’s effort choice is affected by this type of hedging activity, which therefore serves to undo the link between the agent’s wealth and the performance of his firm (which depends on the effort choice). As a result, the agent’s effort

choice is primarily determined through the hedging market that is characterized by the presence of moral hazard. Similarly, in the context of our model, if the agent could trade side contracts based on his own firm value \mathbf{X} , he would effectively do so to reduce the pay–performance sensitivity of his compensation contract.¹⁴ Such trades would reduce the agent’s disutility from exposure to firm-specific risk as well, but would do so by lowering s instead of diversifying \mathbf{E} . The negative results in the literature thus rely on hedging transactions that undermine the agent’s subsequent effort incentives.¹⁵ In contrast, we describe a type of hedging transaction that diversifies the firm-specific risk \mathbf{E} while preserving the link between the agent’s effort choice and his wealth distribution provided by s .

Clearly, *CARA* preferences and the normality assumptions on \mathbf{E} and \mathbf{Y} provide tractability and help us to express the results conveniently in closed form. Clearly, if \mathbf{Y} is a function of \mathbf{X} since the agent can trade \mathbf{Y} to eliminate the randomness introduced by \mathbf{X} this positive result cannot survive. Hence, it is essential that \mathbf{Y} is not a function of \mathbf{X} . However, the robustness of our results remain a question. We discuss this issue in the Concluding Remarks section.¹⁶

The analysis up to this point shows that the availability of a portfolio correlated with firm-specific risk increases contracting efficiency and that the extent of this efficiency benefit increases in the correlation of the portfolio available. Furthermore, whether it is the agent or the principal who trades this portfolio is irrelevant in the case of a single portfolio: The same efficiency benefit is achieved regardless of the party undertaking the trade. Although a single correlated portfolio setting is convenient for illustrative purposes, it does not allow one to be more specific about the type of frictions and costs that determines the amount of diversification that can be achieved nor, more importantly, about which party—the principal or the agent—should undertake the diversification trade. In what follows, we focus on the general case with a variety of financial portfolios available in the market. Furthermore, we introduce a friction for both the agent’s and the principal’s access to the financial market, by assuming that neither of the two parties knows perfectly how different portfolios in the market are correlated with the firm-specific risk.

The literature on corporate risk management provides some background to motivate our introduction of the friction that the parties face in their access to financial markets for diversification purposes. As we discussed in the Introduction, Stulz (1984) has argued that risk-averse managers would use their company’s risk management policies to diversify their own risks. Criticizing this point of view, Froot et al. (1993, p. 1632) point out that “[o]ne weakness of Stulz theory is that it implicitly relies on the assumption that managers face significant costs when trading in hedging contracts for their own account—otherwise, they would be able to adjust the risks they face without having to involve the firm directly in their hedging activities.” It seems natural to argue that in practice it is costly for both the manager and the firm to trade in the financial markets. Furthermore, since the risk that the parties want to diversify is firm specific, it is unlikely that they will have *ex ante* full knowledge of how all the different portfolios are correlated with this particular risk. To capture this type of friction, and to address the question of which party should undertake the hedging activity, we consider a search-theoretic framework in which the agent or the principal has to pay an information acquisition cost to learn about the correlation of a portfolio before creating a position in that portfolio.

We should point out that one could also exogenously specify a “cost of diversification” function for the agent and the principal and address the question of which party should undertake the hedging activity. Yet this reduced-form approach does not seem to be satisfactory in terms of explaining why it is costly for the parties to find and trade the best portfolio in the first place.

¹⁴ For example, one such hedging transaction based on own firm value is an equity swap contract in which the manager promises the return from his company shares to a third party in exchange for a fixed payment (cf. Bettis et al., 2001). A swap transaction reduces the manager’s effective share ownership and dilutes the link between the manager’s wealth and his effort decision.

¹⁵ Of course, the third parties would anticipate the diminished effort incentives and make the price of the hedging transaction conditional on the manager expending less effort subsequent to hedging (see, for example, Bisin et al., 2008; Çelen and Özertürk, 2007).

¹⁶ We are thankful to a referee for encouraging us to clarify the level of generality of our result.

Similarly, this reduced-form approach would not be precise in capturing how certain financial market characteristics affect contracting efficiency. Finally, and perhaps most importantly, without the search-theoretic costly information acquisition framework, with the same exogenous diversification costs for the two parties, which of the two should undertake the hedging activity would be a matter of indifference. As we illustrate in our analysis, the search framework introduces an uncertainty in the costs that parties face *ex ante* when they begin uncovering trading opportunities. Because of this *ex ante* uncertain diversification cost, to achieve any desired level of diversification *with equal unit search costs*, it is always less costly for the risk-neutral principal to undertake the hedging activity herself, instead of letting the risk-averse agent do the hedging and compensating him for the *ex ante* uncertain search cost. This interesting observation would be missing in a setting with an exogenously specified cost of diversification. The search model we analyze later captures this point and also allows us to introduce two financial market characteristics that seem to be relevant for the diversification problem of the two parties: the extent of the asset variety in the financial market (financial market sophistication) and the cost of acquiring information on the statistical characteristics of different available portfolios.¹⁷

3.2. Search for a Portfolio. In this subsection, we consider the general case where there is not just one, but a variety of financial portfolios that both the agent and the principal can trade. For the reader’s convenience, let us briefly review the notation. There is a random variable \mathbf{Q} that describes the quality of all available portfolios. Both the principal and the agent can search for a portfolio from the distribution of all qs . By sampling a portfolio from \mathbf{Q} , party $i \in \{a, p\}$ can find out its quality; that is, (s)he can learn its correlation with \mathbf{E} . If party i is “satisfied” with the quality of the portfolio, then (s)he stops and creates a position $d_i \in \mathbb{R}$ in that portfolio; otherwise, (s)he pays a cost $\kappa_i > 0$ and samples another portfolio and so on until (s)he decides to stop the search process. The main question we now address is: Which party should undertake the search and trading activity to diversify the agent’s compensation risk?

3.2.1. The agent’s optimal search behavior. We first characterize the optimal search strategy of an agent in the financial market who holds a contract (s, t) . Suppose that the principal does not do any search and hedging *ex ante*, offers the agent a contract (s, t) , and delegates the hedging decision to the agent. This analysis characterizes the portfolio quality z that the agent holds and hence the amount of compensation risk that he diversifies in the financial market.

Formally, the agent’s search process is modeled as follows: Let $\mathbf{Q}_1, \mathbf{Q}_2, \dots$ be an i.i.d. sequence of random variables with an absolutely continuous distribution function F ; we denote the density function by f . The agent can observe the sequence $\mathbf{Q}_1, \mathbf{Q}_2, \dots$ at a unit cost $\kappa_a > 0$ as long as he wishes.¹⁸ For each $n = 1, 2, \dots$, after observing $\mathbf{Q}_1 = q_1, \mathbf{Q}_2 = q_2, \dots, \mathbf{Q}_n = q_n$ the agent can stop and enjoy the utility

$$-\exp \left\{ -\alpha \left(sf(e^*) + t - c(e^*) - \frac{\alpha}{2} s^2 (\sigma_\epsilon^2 - q_n^2) - \kappa_a n \right) \right\}$$

or he may continue and observe \mathbf{Q}_{n+1} .

¹⁷ This informational obstacle for diversification has also been emphasized by Allen and Gale (2001, p. 449), who argue that “it is easy to imagine situations in which different agents are interested in quite different kinds of information. Suppose that agents want to hedge some idiosyncratic risks, for example, their future incomes from employment. Then, each agent will want to know the correlation between the return to a particular asset and his own labor income.”

¹⁸ Note that we study the problem as if the agent has no memory, in the sense that he does not have access to the realizations observed in the past. However, this assumption is only for expositional ease: Analysis of the case where the agent perfectly recalls the past is exactly the same. To see this, simply observe that if the agent’s optimal strategy does not require him to stop after a realization q_n , he will not stop at any future realization smaller than q_n , because of the i.i.d. nature of the process. Therefore, the agent will keep searching only because he aims to improve what he currently holds.

To formally state the agent's search problem, let us write $\beta := \exp\{\alpha\kappa_a\}$ and define the random variable

$$\mathbf{U}_n := -\exp\left\{-\alpha\left(sf(e^*) + t - c(e^*) - \frac{\alpha}{2}s^2(\sigma_\epsilon^2 - \mathbf{Q}_n^2)\right)\right\}.$$

DEFINITION 1. A stopping rule is a function $\sigma : \bigcup_{\tau=1}^{\infty} \mathbb{R}^\tau \mapsto \{0, 1\}$ such that after the observation $\mathbf{q}_n = (q_1, q_2, \dots, q_n)$, the agent stops if $\sigma(\mathbf{q}_n) = 0$ and continues sampling if $\sigma(\mathbf{q}_n) = 1$.

The agent's problem is to find a stopping rule $\sigma \in \Sigma$ that maximizes the expected utility

$$(3) \quad E(\beta^n \mathbf{U}_n),$$

where Σ is the set of all stopping rules.

For any stopping rule $\sigma \in \Sigma$, let $v(\sigma)$ denote the value of the expected utility given in (3) and define $v^* = \sup_{\sigma \in \Sigma} v(\sigma)$. We first establish the existence of the stopping rule.

PROPOSITION 2. *For the search problem in Definition 1, there exists an optimal stopping rule $\sigma^* \in \Sigma$ such that $v(\sigma^*) = v^*$.*

PROOF. If $E(|\mathbf{U}_n|) < \infty$, and $\lim_{n \rightarrow \infty} \beta^n \mathbf{U}_n = -\infty$, then the existence of a stopping rule is standard (see DeGroot, 1970). Let us write $\mathbf{U}_n := \psi \exp\{\varphi \mathbf{Q}_n^2\}$ where $\psi = -\exp\left\{-\alpha\left(sf(e^*) + t - c(e^*) - \frac{\alpha}{2}s^2\sigma_\epsilon^2\right)\right\}$ and $\varphi = -\frac{\alpha^2}{2}s^2$. Observe that $E(|\mathbf{U}_n|) = |\psi| \int \exp\{\varphi t^2\} dF(t)$. Because $\varphi < 0$ we immediately get $E(|\mathbf{U}_n|) < \infty$. $\lim_{n \rightarrow \infty} \beta^n \mathbf{U}_n = -\infty$ is obvious. ■

At any given stage of the search, the agent knows that if he continues sampling, the optimal stopping rule σ^* yields the value v^* . Therefore, at the continuation, the agent pays the sampling cost and gets the maximum of βv^* and the realization of $\beta \mathbf{U}$. This means that the continuation value should be equal to the expectation of the maximum of βv^* and $\beta \mathbf{U}$. Hence, the optimality equation of the agent's search problem is

$$(4) \quad v^* = \beta E(\max\{\mathbf{U}, v^*\}).$$

The agent stops searching when the realization q yields a utility (excluding the sampling cost) larger than v^* :

$$v^* \leq -\exp\left\{-\alpha\left(sf(e^*) + t - c(e^*) - \frac{\alpha}{2}s^2(\sigma_\epsilon^2 - q^2)\right)\right\}.$$

Rearranging and solving for q , one observes that the agent's optimal search strategy is a cutoff rule with the threshold $z^* := (q^*)^2$. The agent stops sampling if and only if

$$(5) \quad q^2 \geq z^* := \sigma_\epsilon^2 - \frac{2}{\alpha s^2} \left(sf(e^*) + t - c(e^*) + \frac{1}{\alpha} \ln(-v^*) \right).$$

The agent continues sampling until he finds a portfolio with quality q^* so that he can reduce the variance of his wealth distribution to $s^2(\sigma_\epsilon^2 - z^*)$. In particular, the agent's optimal search procedure is the threshold strategy: If the realization of \mathbf{Q} is between $-|q^*|$ and $|q^*|$, then the agent keeps searching, and he stops when $-|q^*| \geq q$ or $q \geq |q^*|$. The following proposition

characterizes the threshold portfolio quality q^* :

PROPOSITION 3. *For a given compensation contract (s, t) , the agent's optimal stopping rule σ^* in the financial market is characterized by a pair (v, q^*) such that*

$$(6) \quad \beta^{-1} - [F(q^*) - F(-q^*)] = \int^{-q^*} \exp \left\{ -\frac{\alpha^2 s^2}{2} (\tau^2 - z^*) \right\} dF(\tau) + \int_{q^*} \exp \left\{ -\frac{\alpha^2 s^2}{2} (\tau^2 - z^*) \right\} dF(\tau)$$

and the agent stops at

$$v = \min \{n \in \{1, 2, \dots\} : q_n^2 \geq z^*\},$$

where $z^* = (q^*)^2$.

PROOF. See the Appendix.

For future reference, we now describe the effect of the unit search cost κ_a on the agent's optimal search behavior. Not surprisingly, as the search becomes more costly, the agent searches less aggressively.¹⁹

PROPOSITION 4. *Suppose f is differentiable. As the search cost κ_a increases, the agent searches less aggressively in the asset market. That is, $dz^*/d\kappa_a \leq 0$: As κ_a increases, the agent's optimal stopping threshold z^* decreases.*

PROOF. See the Appendix.

3.2.2. *Optimal linear contract when only the agent trades.* We now analyze the principal's problem of setting the optimal linear compensation contract (s, t) when the hedging activity is completely delegated to the agent. This exercise provides a benchmark for comparison in determining whether the principal should undertake another search before offering the agent a compensation contract.

In setting the optimal contract (s, t) , the principal takes into account the agent's optimal search for a portfolio, his subsequent optimal position in the portfolio, and his optimal effort choice, e^* . Formally, the principal chooses the linear contract (s, t) to maximize

$$(1 - s)E(\mathbf{X}_{e^*}) - t$$

subject to the participation constraint

$$E(\beta^{\mathbf{N}_{v^*}})v^* \geq \bar{v},$$

where v^* is the optimal utility attained in the search process defined in Equation (4), $\bar{v} = -\exp\{-\alpha\bar{w}\}$ is the agent's reservation utility, and \mathbf{N}_{v^*} is the random variable determining the number of times the agent has to search in order to reach v^* . Hence, the term $E(\beta^{\mathbf{N}_{v^*}}) =: \phi(z^*)$

¹⁹ Our exercise is akin to random search models in monetary economics, industrial organization, finance, and labor economics. Similar results are obtained in the theory of job search. In particular, the notion of reservation wage is analogous to the threshold quality of a portfolio, and they react similarly to changes in search cost. For a general discussion of job search models, see Cahuc and Zylberberg (2004), Devine and Kiefer (1991), and Rogerson et al. (2005). Also, although there are modeling differences, Burdett and Ondrich (1985), Loberg and Wright (1987), and Boldrin et al. (1993) obtain similar comparative static results.

is the *expected cost of optimal search*. We can write this participation constraint in terms of certainty equivalent wealth as

$$sf(e^*(s)) + t - c(e^*(s)) - \frac{\alpha}{2}s^2(\sigma_\epsilon^2 - z^*(s)) - \bar{w} - \frac{1}{\alpha} \ln(\phi(z^*)) \geq 0.$$

In the above expression, the term $1/\alpha \ln(\phi(z^*))$ is the *ex ante search cost* of having the agent (measured in terms of wealth) perform an optimal search with threshold z^* . In equilibrium, the participation constraint holds as an equality. Solving for t and substituting it into the principal's objective function, the problem becomes choosing the agent's pay-performance sensitivity s to maximize the net surplus

$$(7) \quad f(e^*(s)) - c(e^*(s)) - \frac{\alpha}{2}s^2(\sigma_\epsilon^2 - z^*(s)) - \bar{w} - \frac{1}{\alpha} \ln(\phi(z^*)).$$

Differentiating (7) with respect to s , and using the optimality condition for e^* from (1) we obtain the following result:

PROPOSITION 5. *Suppose the principal delegates the hedging decision to the agent and offers a contract (s, t) . Then, the optimal pay-performance sensitivity s^* satisfies one of the following three conditions:*

$$\begin{aligned} s^* \in (0, \bar{s}] \text{ and } (1-s)f'(e^*) \frac{de^*}{ds} &= \alpha s \left(\sigma_\epsilon^2 - z^* - \frac{s}{2} \frac{dz^*}{ds} \right) + \frac{1}{\alpha} \frac{\phi'(z^*(s))}{\phi(z^*(s))} \frac{dz^*(s)}{ds}, \\ s^* \in (\bar{s}, 1) \text{ and } \alpha s \left(\sigma_\epsilon^2 - z^*(s) - \frac{s}{2} \frac{dz^*}{ds} \right) &+ \frac{1}{\alpha} \frac{\phi'(z^*(s))}{\phi(z^*(s))} \frac{dz^*(s)}{ds} = 0, \\ s^* = 1 \text{ and } \alpha \left(\sigma_\epsilon^2 - z^*(1) - \frac{1}{2} \frac{dz^*}{ds} \Big|_{s=1} \right) &+ \frac{1}{\alpha} \frac{\phi'(z^*(1))}{\phi(z^*(1))} \frac{dz^*(s)}{ds} \Big|_{s=1} \leq 0, \end{aligned}$$

where $\bar{s} = \min\{\bar{s}, 1\}$ such that $e^*(\bar{s}) = \bar{e}$. Moreover, s^* is nonincreasing in the agent's search cost κ_a in the asset market and nondecreasing in the financial market's sophistication.

The first condition characterizes the interior solution when $e^*(s^*) \in (\underline{e}, \bar{e})$. In the second condition, the agent's optimal effort for s^* exceeds the maximum effort \bar{e} , and s^* is still in the interior of $(0, 1)$. In that case, although the incentive provided for the agent does not lead to a marginal increase in the effort level, it makes the agent search more aggressively. Hence, at $e^* = \bar{e}$ the optimal choice of s equates the marginal decrease in the risk with the marginal increase in the expected cost of the search. Finally, the last condition refers to the boundary case in which the objective function is increasing at $s = 1$. The optimal incentive scheme s^* is determined by the trade-off between providing effort incentives by making pay more sensitive to output and compensating the risk-averse agent for bearing the associated risk exposure (insurance). The more risk the agent can diversify, the lower the premium required for a given risk exposure is, and hence the higher the pay-performance sensitivity of the agent's compensation scheme is.

3.2.3. The principal's search behavior. We now analyze the principal's search behavior to determine which party should undertake the search. To distinguish the principal's search problem, let us denote the realization of the quality of a portfolio for her by \tilde{q} and $\tilde{q}^2 =: \zeta$. The problem of the principal is as follows: For each $n = 1, 2, \dots$, after observing $\mathbf{Q}_1 = \tilde{q}_1, \mathbf{Q}_2 = \tilde{q}_2, \dots, \mathbf{Q}_n = \tilde{q}_n$, the principal can stop and offer the agent a contract (s, d_p, t) of the form $s\mathbf{X}_e + d_p \mathbf{Y} + t$, where \mathbf{Y} is the portfolio with quality \tilde{q}_n^2 , or pay the unit search cost κ_p to continue sampling and observe \mathbf{Q}_{n+1} .

We should emphasize that when the principal stops and offers a contract, in the following subgame the agent decides whether to accept the contract or not. If he accepts, the agent can

also undertake another search and make the optimal effort choice. It is important to clarify that if the agent undertakes another search, that does not mean he will *append* anything to the portfolio he is already offered in his contract, but rather that he will *replace* it. Indeed, this is a plausible interpretation, since we assume that each sampling in the search aims to improve the portfolio the agent—or the principal—builds during the search.

In particular, if the initial contract specifies s and an optimal position in a portfolio with quality ζ , the optimal quality the agent can obtain in the search should be $\max\{\zeta, z^*(s)\}$, for z^* defined in Proposition 3. Consequently, if the agent accepts the contract he may engage in another search only when $\zeta < z^*$. An immediate implication of this observation is that, in any contract that incentivizes the agent to undertake another search, the principal never finds it optimal to do any search and provide a portfolio as part of the initial contract. In other words, *there is no equilibrium in which both the principal and the agent search.*

In order to see this point formally, suppose that the principal offers the agent a contract that includes a portfolio with quality $\zeta < z^*(s)$. Let $\rho(\zeta)$ denote the principal's expected search cost that corresponds to the stopping threshold ζ . Then the principal's problem is to maximize $(1-s)E(\mathbf{X}_{e^*}) - t - \rho(\zeta)$ subject to the agent's participation constraint $\phi(z^*)v^* \geq \bar{v}$, where $\phi(z^*)$ is the agent's expected cost of optimal search with the threshold z^* . When we write the participation constraint in terms of the certainty equivalent wealth and rearrange the objective function, the principal's problem can be rewritten as choosing s and ζ to maximize

$$f(e^*(s)) - c(e^*(s)) - \frac{\alpha}{2}s^2(\sigma_\varepsilon^2 - z^*(s)) - \bar{w} - \frac{1}{\alpha} \ln(\phi(z^*(s))) - \rho(\zeta).$$

It is clearly suboptimal to incur $\rho(\zeta)$ and also compensate the agent for his expected optimal search cost by $\frac{1}{\alpha} \ln(\phi(z^*(s)))$. Therefore, in any optimal contract the principal either offers $\zeta \geq z^*$ or lets the agent undertake the search completely.

The above argument rules out the possibility of an equilibrium contract in which both the principal and the agent undertake a search. In what follows, instead of fully characterizing the principal's optimal search behavior, we will compare the principal's objective functions when only the agent undertakes the search and when only the principal undertakes the search. This comparison allows us to address the question of which party undertakes the search in an optimal linear contract.

If only the agent does the search, the principal's objective function is given by

$$f(e^*(s)) - c(e^*(s)) - \frac{\alpha}{2}s^2(\sigma_\varepsilon^2 - z^*(s)) - \bar{w} - \frac{1}{\alpha} \ln(\phi(z^*(s))),$$

whereas if only the principal does the search, her objective function takes the form

$$f(e^*(s)) - c(e^*(s)) - \frac{\alpha}{2}s^2(\sigma_\varepsilon^2 - \zeta) - \bar{w} - \rho(\zeta).$$

Earlier in our analysis, we showed that as his search cost κ_a decreases, the agent's search intensity $z^*(s)$ increases. Accordingly, if the agent's search cost is different than the principal's, which party should undertake the search depends on the relative search cost. In particular, if κ_a is sufficiently small compared to the principal's cost, κ_p , then the principal may find it optimal to let the agent do the search. The exact comparison of the two problems with different search costs requires a full characterization of the principal's search problem. However, one can readily analyze the most interesting case—when the search costs are identical—which we do next.

Suppose that the search costs of the two parties are identical, $\kappa_a = \kappa_p = \kappa$, and that the *optimal* contract specifies a pay–performance sensitivity $\hat{\delta}$. If the contract makes the agent do the search, the threshold portfolio quality that the agent would seek is given by $z^*(\hat{\delta})$. Alternatively, if the contract makes the principal do the search, the principal would specify a portfolio with quality,

say, $\hat{\zeta}$. We already know that $\hat{\zeta} \geq z^*(\hat{s})$. We now show that when the search costs are equal, the principal always prefers to undertake the search herself. Consider the case where $\hat{\zeta} = z^*(\hat{s})$. Clearly, if the principal prefers to do the search when $\hat{\zeta} = z^*(\hat{s})$, in an optimal contract she would continue to do so even if $\hat{\zeta} > z^*(\hat{s})$. Note that the only difference between the two problems is that when the principal does the search she incurs an expected search cost $\rho(\hat{\zeta})$, whereas if she lets the agent do the search, she needs to compensate the agent for his expected search cost $1/\alpha \ln(\phi(z^*(\hat{s})))$. The proposition below shows that when the unit search costs of the two parties are the same, for the same threshold portfolio quality the principal always prefers to undertake the search herself, instead of letting the agent do the search and compensating the agent for his search cost.

PROPOSITION 6. *Let z be a stopping threshold for the stopping problem stated in Subsection 3.2. Denote $\Delta(z) = F(z) - F(-z)$ and assume that $\Delta(z)e^{\alpha\kappa_a} < 1$. The (expected) utility costs of a search with a stopping threshold z for the agent and for the principal are*

$$\phi(z) := \frac{e^{\alpha\kappa_a}(1-\Delta(z))}{1-e^{\alpha\kappa_a}\Delta(z)}, \quad \rho(z) := \frac{\kappa_p}{1-\Delta(z)},$$

respectively.

Suppose that $\kappa_a = \kappa_p$. Compensating the agent in order to give incentives to him to undertake a portfolio search is more costly for the principal than her own expected search cost if she undertakes the same search herself. That is, $\frac{1}{\alpha} \ln(\phi(z)) > \rho(z)$ for any z .

PROOF. See the Appendix.

Proposition 6 proves useful for understanding the basic elements that determine which party should undertake the search. Clearly, the most obvious of those elements is the difference between the search costs κ_a and κ_p of the two parties. The comparison in the proposition assumes that the search cost is identical for the agent and the principal. However, a quick look at the terms $\phi(z)$ and $\rho(z)$ that determine the utility cost of the search shows—not surprisingly—that they both monotonically increase in the unit search cost. Thus, a large enough difference between unit costs κ_a and κ_p can offset the other factors (which we discuss later) and determine who will undertake the search.

The second, and perhaps more interesting, factor that determines who makes the search is the differences between the risk preferences of the agent and the principal. Note that the search process per se involves an element of risk for the agent if he makes the search himself. A given pay–performance sensitivity s determines the optimal threshold $z^*(s)$ at which the agent will stop in his search. From an ex ante point of view, this means that the agent is uncertain about how much searching it will take to achieve this threshold. There is an expected search cost for the agent to achieve $z^*(s)$, but the actual search cost that will be incurred is ex ante uncertain. As such, from the principal’s perspective, if she prefers to delegate the search to the agent, he needs to be compensated with a premium to undertake the desired amount of search. The agent prefers to receive the hedging instrument in his contract rather than to receive a compensation in his contract to perform a search for the same amount of hedging, simply because the former completely eliminates the risk of an uncertain search cost for him. On the flip side of the relationship, for the same exact reason, the principal will find it more costly to delegate the same amount of search (assuming $\kappa_a = \kappa_p$) than undertaking it herself. Therefore, the fact that the agent is risk averse, and hence needs to be compensated ex ante for an uncertain search cost, proves important in determining who should undertake the search when the unit search costs are equal.

The last and perhaps most subtle factor is the flexibility that the principal gains when she undertakes the search herself. Let us focus on the benchmark case where the unit search costs

are equal. Also recall that for a given s , the agent's search behavior is characterized by the threshold $z^*(s)$. If the principal undertakes the search, then her problem involves choosing a stopping rule ζ and a contract (s, d_p, t) . However, if she lets the agent do the search, her only choice variable is (s, t) for a given $z^*(s)$. Clearly, in the former scenario both the agent and the principal are better off since the optimization problem is relaxed. In particular, the principal can set the threshold $\zeta = z^*(s)$ and the position $d_p = d_a$ and optimally choose (s, t) and do at least as well as in the latter case. Therefore, if $\kappa_a = \kappa_p$ and the agent is more risk averse than the principal, the principal prefers to do the search herself.

In the light of Proposition 6, it is also interesting to point out that modeling the search frictions explicitly proves important in our framework. Recall that with only one correlated portfolio available, our analysis established that it is irrelevant which party undertakes the hedging transaction. The particular point captured by our costly search framework with a variety of portfolios is that now uncovering the hedging opportunity is subject to an ex ante uncertain search cost. This point would be missing if we had instead modeled the financial market frictions with an exogenously specified diversification cost. Such an exogenous cost specification would yield a similar irrelevance result as long as there is no ex ante uncertainty regarding the cost of diversification. With the costly search framework we describe, if the principal delegates the hedging decision to the risk-averse agent, it is not ex ante certain how much search cost the agent will incur to achieve a given level of diversification. As we argued, this ex ante "risk" embedded in the search process implies that, when the search costs are equal, the risk-neutral principal should undertake the search, as otherwise the risk-averse agent should be compensated ex ante with an additional risk premium to achieve a given level of diversification. In that respect, our analysis captures a novel feature that is not previously explored in the corporate risk management literature. As we argued in the Introduction, this literature has correctly emphasized that the transaction cost differential between the two parties is likely to determine which party should optimally do the hedging (Froot et al., 1993). However, the literature has not recognized that it may be better that the principal undertake the hedging on the agent's behalf due to uncertainties involved in uncovering the optimal hedge portfolio for the agent, which we capture through explicitly modeling the frictions parties face in the financial market.

3.2.4. Hedging and financial market sophistication. Clearly, the extent to which a wide variety of financial portfolios are available in the market affects the optimal search behavior in our model. To be able to build a link between the intensity of the search behavior and the availability of portfolios correlated with firm-specific risk, we now introduce a simple definition of financial market sophistication.

An important characteristic of the financial market is captured by the distribution of the random variable \mathbf{Q} . Notice that F describes the availability of different portfolios in the market and measures the market's "sophistication." Let us illustrate with a simple example: Let u_x denote the uniform density function with support $[-x, x]$. Consider u_x and u_y for $x > y$, and let us refer to them as u_x - and u_y -market, respectively. Because $[-y, y] \subset [-x, x]$, there are portfolios available in u_x -market that are not available in u_y -market. In other words, u_x -market is larger than u_y -market for diversification purposes. Obviously, comparing the support of different markets alone is not sufficient for defining "market sophistication." It may simply be that although one market is larger than the other, the former has too little weight on the tails, and it is therefore costly to attain them. To capture these ideas, we restrict our attention to a class of markets (pdfs) ordered according to a well-known relation: *mean-preserving spread*. We refer to a market as more sophisticated than another if the former's density function is a mean-preserving spread of the latter's. We write $g \succeq f$ when g is a mean-preserving spread of f .

Our goal is to show that the party who undertakes the hedging transaction searches more aggressively in the g -market than in the f -market if $g \succeq f$. In fact, our simple example immediately indicates why this is true. We immediately observe that $u_x \succeq u_y$. A party searching for

portfolios in the u_x -market can easily restrict his search in $[-y, y]$ and reach a utility at least as good as the one he can get in the u_y -market. It is easy to verify our claim and to see what this simple example demonstrates. In what follows, we focus on a large class of markets and therein prove our claim.

DEFINITION 2. Let $\mathfrak{F} := \{f_\lambda\}$ be a family of distribution functions indexed by λ such that (i) $\int f_\lambda(t)dt = 0$ for any λ and $f(t) = f(-t)$ for all t , (ii) $f_{\lambda'} \geq f_\lambda$ for any $\lambda' > \lambda$, and (iii) for any $\lambda, \lambda', f_\lambda$, and $f_{\lambda'}$ satisfy the following single-crossing property in \mathbb{R}_+ : If $f_\lambda(x) - f_{\lambda'}(x) = 0$ and $f_\lambda(x') - f_{\lambda'}(x') = 0$ for some $x' > x > 0$, then $f_\lambda(y) - f_{\lambda'}(y) = 0$ for any $y \in [x, x']$.

Let us briefly discuss the assumptions. We first restrict our attention to a class of density functions that are symmetric around zero. Although this seems like a technical restriction, it is merely an assumption in terms of the qualitative aspect of the problem at hand: From the viewpoint of our agent, what really matters is $|q|$. The second assumption restricts our analysis to a class of functions such that we can compare any two distribution functions with respect to the relation \geq that is central to the idea of market sophistication. Finally, for technical reasons we assume the single-crossing condition. Although it narrows the class of functions we work with, we still have many well-known distributions that satisfy these assumptions. The family of normal distributions with zero mean and variance λ is just one example. For a proper choice of parameters, uniform, logistic, cauchy, Student's t , and beta distributions also satisfy the assumptions.

For the purpose of simple illustration, we describe the comparative static result for the case in which the agent undertakes the search. Consider the agent's optimal stopping rule as described in Equation (6). Note that for $f \in \mathfrak{F}$ we can rewrite Equation (6) as

$$1 - \beta^{-1} = 2 \int_{q^*} 1 - \exp\left\{-\frac{\alpha^2 s^2}{2}(\tau^2 - z^*)\right\} dF(\tau).$$

Now for any $f \in \mathfrak{F}$, define

$$\varphi(\tau; f) := \int_\tau h(t)f(t)dt,$$

where $h(t) := 1 - \exp\left\{-\frac{\alpha^2 s^2}{2}(t^2 - \tau^2)\right\}$.

LEMMA 1. For any $f \in \mathfrak{F}$, $\varphi(\tau; f)$ is nonincreasing in τ . Furthermore, for any $f, g \in \mathfrak{F}$, if $g \geq f$ then $\varphi(\tau; f) \leq \varphi(\tau; g)$.

PROOF. Let $f \in \mathfrak{F}$. It is clear that

$$\frac{d\varphi(\tau; f)}{d\tau} = -\alpha^2 s^2 \int_\tau (1 - h(t))dF(t) < 0.$$

Let $f, g \in \mathfrak{F}$, and assume that g is a mean-preserving spread of f . Also define $\iota := g - f$. We want to show that

$$\int_\tau h(t)\iota(t)dt \geq 0.$$

Note that $h(t) > 0$ for all $t \in \mathbb{R}$. Also, recall that g and f satisfy the single-crossing condition. Because g is a mean-preserving spread of f , there exists at most one interval $[x_*, x^*]$ such that for any $x < x_*$ $\iota(x) < 0$ and for any $x > x^*$ $\iota(x) > 0$. The problem is trivial if $\iota(x) \geq 0$ for all $x \in \mathbb{R}_+$.

Let $x_*, x^* \in \mathbb{R}_+$ such that $u(x_*) = u(x^*) = 0$. Denote $h_* := h(x_*)$ and $h^* := h(x^*)$. Now observe that $h(x)u(x) \geq h_*u(x)$ for all $x < x_*$ and $h(x)u(x) \geq h^*u(x)$ for all $x > x^*$ because h increases monotonically. Therefore, we have

$$\int_{\tau} h(t)u(t)dt \geq h^* \int_{\tau} u(t)dt \geq 0.$$

Hence, we get $\varphi(\tau; f) \leq \varphi(\tau; g)$. ■

COROLLARY 1. *Let $f, g \in \mathfrak{F}$. There exists $\tau, \varsigma \in \mathbb{R}$ such that $\varphi(\tau; f) = 1/2(1 - \beta^{-1}) = \varphi(\varsigma; g)$. If $g \succeq f$, then $\varsigma \geq \tau$.*

PROOF. Let $f, g \in \mathfrak{F}$ and $\tau \in \mathbb{R}_{++}$ such that $\varphi(\tau; f) = 1/2(1 - \beta^{-1})$. Suppose f second order stochastically dominates g . Then by Lemma 1, $\varphi(\tau; g) \geq 1/2(1 - \beta^{-1})$. Since $\varphi(\cdot; g)$ is nonincreasing, $\varphi(\varsigma; g) = 1/2(1 - \beta^{-1})$ implies $\varsigma \geq \tau$. ■

In light of these observations, we are ready to state the result:

PROPOSITION 7. *Suppose that there are two financial markets that are characterized by $f, g \in F$, respectively, and that the agent undertakes the search. If g is more sophisticated than f , the agent searches more aggressively in the g -market than in the f -market. That is, the optimal stopping threshold in the g -market is larger than the threshold in the f -market.*

PROOF. The result follows from Corollary 1. ■

4. CONCLUDING REMARKS

4.1. Robustness of the Results. We derive our results from a simple model in which the agent has CARA preferences, the principal offers a linear contract, and the output is composed of a deterministic production technology and an additive normal noise structure. These assumptions provide an analytically tractable structure, although they restrict the generality of the results. In this subsection, we examine the robustness of our results with regard to these specific assumptions.

Linear contracts are prevalent in common practice. As Milgrom and Roberts (1992) argue “[I]inear contracts are commonly observed in the form the commissions paid to sales agents, contingency fees paid to attorneys, piece rates paid to tree planters or knitters, crop shares paid to sharecropping farmers, and so on.” Moreover, some authors (viz., Lazear, 1986; Brown, 1990) discuss conditions under which linear contracts can perform well in practice. Bose et al. (2011) use computational techniques to compare the performance of optimal contracts vis-à-vis linear contracts. They show that in a large class of models, if the principal uses a linear contract, she can secure 95% of the profits that she could generate under an optimal contract.

Another aspect of linear contracts that makes them appealing in agency-theory literature is their tractability. In that literature, attention has been restricted to linear contracts in a variety of applications; analyses of those contracts has provided insights that seem generalizable to other contracts as well. We intuit that the same is true of the present model. As we discussed earlier, the main reason that our hedging framework has a positive effect on contracting efficiency is that it does not change the agent’s exposure to the firm’s value \mathbf{X} in his overall wealth distribution. In other words, the existence of hedging possibility does not directly affect the agent’s optimal choice of effort for a given pay–performance sensitivity. Hence, it seems that this result can be extended to contracts that preserve the positive relation between the agent’s optimal choice of effort and his exposure to the firm’s value.

The two other critical assumptions are the exponential (CARA) utility function and the fact that the noise term is normally distributed. These assumptions, like that of the linear contract, are of great help in maintaining the tractability of the analysis. However, recall that availability of hedging is critical for our results simply because the agent's wealth increases for a given hedging portfolio; that relaxes the participation constraint, which in turn enables the principal to expose the agent to more risk. This seems to suggest that as long as an optimal position in a hedging instrument increases the agent's expected utility, the results will still hold.

Finally, we want to discuss the particular additive functional form that determines the final output. Note that the assumption that the final output is additive isolates any possible effect of the effort on the variability of the output; this assumption was also crucial in the above discussions. Clearly, in a more general model the agent's effort choice might have a direct effect on output. To carefully examine this issue, suppose that \mathbf{X}_e is the (output) random variable distributed with $F(x|e)$, and $F(x, y|e)$ is the joint distribution of \mathbf{X}_e and a given portfolio \mathbf{Y} at some effort level e . In that case, the agent's problem will be to choose e to maximize

$$U(x, y, e) := \int u(s(x), y, e) dF(x, y|e),$$

where $s(x)$ is the contract offered by the principal. Suppose that the agent makes the search and takes an optimal position in a portfolio \mathbf{Y} . In order to preserve the results that we obtain, taking an optimal position in a portfolio must not lead to a decrease in the agent's optimal effort choice. This is certainly not true if \mathbf{Y} can be written as a function of \mathbf{X}_e , since the agent can take a position to reduce or even eliminate the incentives that are provided by \mathbf{X}_e .

Note that what we need is to maintain at least the same level of effort for a given contract when hedging (i.e., availability of \mathbf{Y}) is introduced. The literature on comparative monotone statics (viz., Athey, 2002) discusses conditions that are sufficient for providing the results we need. In particular, we want $U(x, y, e)$ to behave well, and Athey (2002) provides conditions on primitives u and F (such as log-supermodularity) to obtain such results. The only remaining issue is whether these conditions hold when the contract that the principal offers, $s(x)$, is incorporated. Although the linearity of contracts preserves log-supermodularity of u , we cannot yet make this claim for more general contracts. As such we leave this exercise for future research.

4.2. Conclusion. The main insights of our analysis are as follows: In a principal-agent framework, the availability of financial portfolios correlated with firm-specific risk can improve contracting efficiency by reducing the noise in the risk-averse agent's compensation scheme. In a search-theoretic setting where both the principal and the agent can uncover and undertake such beneficial trades, for the same unit search costs it is optimal for the principal to do the hedging on the manager's behalf. One explanation for this result is that the search process per se introduces an element of risk into the agent's wealth distribution. From an ex ante point of view, the agent is uncertain about how much search it will take to achieve a given level of noise reduction, and therefore needs to be compensated with a premium to undertake the desired amount of search. Due to this extra premium, the principal finds it more costly to delegate the same amount of search to the agent instead of undertaking it herself.

Previous papers in the managerial hedging literature have considered settings in which the agent can trade financial instruments based on his own firm value: Such trades are shown to undermine incentives, since they reduce the sensitivity of the agent's wealth to firm performance. In this article, we consider a different type of hedging transaction: trading portfolios correlated with firm-specific risk serve to diversify compensation risk while preserving the link between the agent's effort decision and his wealth. As a result, the trades in these portfolios reduce the insurance cost of incentive provision and enable the principal to elicit a higher level of effort.

Our article is also related to the arguments proposed in the corporate risk management literature with regard to the issue of reducing the risk-averse manager's compensation risk. It has been proposed that the manager can trade in the markets on his own account to achieve diversification (Froot et al., 1993) or that the firm can build a filter into the manager's compensation

scheme to diversify away some noise (Holmström and Tirole, 2000). Our model addresses the question of which one of these two options achieves a more efficient contracting outcome in a setting where the respective diversification costs of the manager and the firm are explicitly modeled. Finally, this article also contributes to the recent debate on the desirability of managerial hedging transactions. Our analysis suggests that the manager’s trades in the financial markets can improve efficiency by reducing the noise in his compensation contract; however, it might even prove better if such beneficial hedging transactions are undertaken by the firm on the manager’s behalf, instead of letting the manager hedge himself.

Before concluding, it seems important to discuss more specifically how an efficiency-enhancing portfolio might be constructed to hedge idiosyncratic risk in the compensation contract. As an example, consider the CEO of a gold mining corporation who holds his own company stock as part of his compensation scheme. At the risk of oversimplification, but for the sake of a simple argument, let us assume that the company’s revenues, and hence its stock price, are subject to gold price risk (which is common to all gold mining companies) and also production/output risk specific to this company. More concretely, suppose that this gold company is heavily engaged in exploring and developing gold mines in a Latin American country that recently elected a left-oriented populist government with a nationalization agenda that may target foreign companies. If this company’s main existing and future gold production ability is heavily dependent on the mines they currently operate in that particular country, one would imagine that a high proportion of this company’s firm-specific production risk stems from this nationalization threat.

The CEO, or the shareholders who design the compensation scheme, can hedge the CEO’s gold price risk with a long position in U.S. dollar index, which is typically negatively correlated with gold prices. This type of transaction would hedge an industry-specific risk factor in the compensation scheme—namely, the gold price—and would be akin to the type of hedging considered in Jin (2002), Acharya and Bisin (2005), and Özertürk (2006). However, in terms of hedging against a possible plunge in production capacity due to nationalization, which constitutes the main firm-specific risk in our example, a long position in the U.S. dollar index would not be helpful at all. Consider then an investment portfolio that tracks foreign investment flows to the country in question. The value of this investment portfolio would be negatively correlated with the nationalization efforts of the government: to the extent that the nationalization agenda is carried out, foreign investment flow would decline. In that respect, a short position in a portfolio that tracks foreign investment flows to the country in our example would work to diversify some of the company’s firm-specific production risk.²⁰

APPENDIX

PROOF OF PROPOSITION 3. Note that Equation (4) can be written as

$$\begin{aligned}
 v^* &= \beta \int_{-q^*}^{q^*} v^* dF(t) + \beta \int^{-q^*} - \exp \{ -\alpha(\varphi - \psi\sigma_\epsilon^2) - \alpha\psi t^2 \} dF(t) \\
 &\quad + \beta \int_{q^*} - \exp \{ -\alpha(\varphi - \psi\sigma_\epsilon^2) - \alpha\psi t^2 \} dF(t) \\
 &= \frac{\beta}{1 - \beta\Delta F^*} \left(\int^{-q^*} - \exp \{ -\alpha(\varphi - \psi\sigma_\epsilon^2) - \alpha\psi t^2 \} dF(t) \right. \\
 &\quad \left. + \int_{q^*} - \exp \{ -\alpha(\varphi - \psi\sigma_\epsilon^2) - \alpha\psi t^2 \} dF(t) \right),
 \end{aligned}$$

²⁰ One might argue that constructing such an investment vehicle might be cumbersome, and furthermore that the extent to which such a hedge works depends on the degree of correlation between the company’s production risk in that country and how the prospect of nationalization affects foreign investment flows. We should point out that this is precisely the point we aim to make by establishing a framework of a costly search of a portfolio. Unlike instruments that are useful in hedging industry-wide aggregate risks (such as gold price risk in our example), financial instruments that can help to hedge idiosyncratic risk are much less readily available. Successful hedging of firm-specific risk requires identifying the singular adverse scenarios a company may face (i.e., what constitutes firm-specific risk) and seeking out investment vehicles that are correlated with the possibility of such adverse events.

where $\Delta F^* = F(q^*) - F(-q^*)$. Expressing v^* in terms of z^* by use of (5) and rearranging the expression above yields the result

$$\begin{aligned}
 \text{(A.1)} \quad -\exp\left\{-\alpha(\varphi - \psi(\sigma_\epsilon^2 - z^*))\right\} &= \frac{\beta}{1 - \beta\Delta F^*} \left(\int^{-q^*} -\exp\left\{-\alpha(\varphi - \psi\sigma_\epsilon^2) - \alpha\psi t^2\right\} dF(t) \right. \\
 &\quad \left. + \int_{q^*} -\exp\left\{-\alpha(\varphi - \psi\sigma_\epsilon^2) - \alpha\psi t^2\right\} dF(t) \right) \\
 \frac{1 - \beta\Delta F^*}{\beta} &= \int^{-q^*} \exp\left\{-\frac{\alpha^2 s^2}{2}(t^2 - z^*)\right\} dF(t) \\
 &\quad + \int_{q^*} \exp\left\{-\frac{\alpha^2 s^2}{2}(t^2 - z^*)\right\} dF(t). \quad \blacksquare
 \end{aligned}$$

PROOF OF PROPOSITION 4. Taking the derivatives of both sides of (A.1) with respect to κ , we get

$$\begin{aligned}
 -\alpha \exp\{-\alpha\kappa_a\} - \frac{dz^*}{d\kappa_a}(f(q^*) + f(-q^*)) &= \frac{dz^*}{d\kappa_a} \int^{-q^*} \left(\frac{\alpha^2 s^2}{2}\right) \exp\left\{-\frac{\alpha^2 s^2}{2}(t^2 - z^*)\right\} dF(t) \\
 &\quad + \frac{dz^*}{d\kappa_a} \int_{q^*} \left(\frac{\alpha^2 s^2}{2}\right) \exp\left\{-\frac{\alpha^2 s^2}{2}(t^2 - z^*)\right\} dF(t) \\
 &\quad - \frac{dz^*}{d\kappa_a}(f(q^*) + f(-q^*)).
 \end{aligned}$$

Rearranging this expression yields the result

$$\frac{dz^*}{d\kappa_a} = -\frac{2 \exp\{-\alpha\kappa_a\}}{\alpha s^2 \left(\int^{-q^*} \exp\left\{-\frac{\alpha^2 s^2}{2}(t^2 - z^*)\right\} dF(t) + \int_{q^*} \exp\left\{-\frac{\alpha^2 s^2}{2}(t^2 - z^*)\right\} dF(t) \right)} < 0. \quad \blacksquare$$

PROOF OF PROPOSITION 6. Let z be the stopping threshold. For a given F , let $\Delta(z) = F(z) - F(-z)$. Assume that $e^{\alpha\kappa_a} \Delta(z) < 1$. The expected utility cost of search for the agent is

$$\begin{aligned}
 \phi(z) &:= e^{\alpha\kappa_a}(1 - \Delta(z)) + e^{2\alpha\kappa_a}(1 - \Delta(z))\Delta(z) + e^{3\alpha\kappa_a}(1 - \Delta(z))\Delta(z)^2 + \dots \\
 &= e^{\alpha\kappa_a}(1 - \Delta(z)) [1 + e^{\alpha\kappa_a} \Delta(z) + e^{2\alpha\kappa_a} \Delta(z)^2 + \dots] \\
 &= \frac{e^{\alpha\kappa_a}(1 - \Delta(z))}{1 - e^{\alpha\kappa_a} \Delta(z)}.
 \end{aligned}$$

Obviously, $\phi(z)$ is infinite if $e^{\alpha\kappa_a} \Delta(z) \geq 1$. On the other hand, the expected utility cost of search for the principal is

$$\begin{aligned}
 \rho(z) &:= \kappa_p(1 - \Delta(z)) + 2\kappa_p(1 - \Delta(z))\Delta(z) + 3\kappa_p(1 - \Delta(z))\Delta(z)^2 + \dots \\
 &= \kappa_p(1 - \Delta(z)) [1 + 2\Delta(z) + 3\Delta(z)^2 + \dots] \\
 &= \kappa_p(1 - \Delta(z)) \left[\frac{1}{1 - \Delta(z)} + \frac{\Delta(z)}{1 - \Delta(z)} + \frac{\Delta(z)^2}{1 - \Delta(z)} + \dots \right] \\
 &= \frac{\kappa_p}{1 - \Delta(z)}.
 \end{aligned}$$

Observe that the principal should pay the agent $\frac{1}{\alpha} \ln(\phi(z))$ in order to compensate his expected utility cost. Therefore, for a given z , in order to decide which utility cost is higher for the principal, we have to compare $\frac{1}{\alpha} \ln(\phi(z))$ and $\varphi(z)$.

Let $\kappa = \kappa_a = \kappa_p$. To show that it is more costly for the principal to compensate the agent for his search than to do it herself, it will be enough to show that

$$\begin{aligned} \frac{1}{\alpha} \ln \left(\frac{e^{\alpha\kappa}(1 - \Delta(z))}{1 - \Delta(z)e^{\alpha\kappa}} \right) &> \frac{\kappa}{1 - \Delta(z)} \\ \frac{e^{\alpha\kappa}(1 - \Delta(z))}{1 - \Delta(z)e^{\alpha\kappa}} &> e^{\frac{\alpha\kappa}{1 - \Delta(z)}} \\ (1 - \Delta(z)) + \Delta(z)e^{\frac{\alpha\kappa}{1 - \Delta(z)}} - e^{\alpha\kappa \frac{\Delta(z)}{1 - \Delta(z)}} &> 0. \end{aligned}$$

Let us write $x = \Delta(z)$ and define $f(x) := 1 - x + xe^{\frac{\alpha\kappa}{1-x}} - e^{\frac{\alpha\kappa x}{1-x}}$. Observe that $f(0) = 0$ and

$$\begin{aligned} f'(x) &= -1 + e^{\frac{\alpha\kappa}{1-x}} + \frac{\alpha\kappa x}{(1-x)^2} e^{\frac{\alpha\kappa}{1-x}} - \frac{\alpha\kappa}{(1-x)^2} e^{\frac{\alpha\kappa x}{1-x}} \\ &= -1 + e^{\frac{\alpha\kappa}{1-x}} \left(1 + \frac{\alpha\kappa x}{(1-x)^2} - \frac{\alpha\kappa}{(1-x)^2} e^{-\alpha\kappa} \right) \\ &\geq 0 \end{aligned}$$

since $xe^{\alpha\kappa} < 1$. Therefore, $f(x) > 0$ for all $x > 0$. Hence, for the same threshold, the principal always prefers to make the search herself, instead of letting the agent do the search and compensating him (through the participation constraint.) ■

REFERENCES

ACHARYA, V. V., AND A. BISIN, “Managerial Hedging, Equity Ownership and Firm Value,” Mimeo, New York University, 2005.

ALLEN, F., AND D. GALE, *Comparing Financial Systems* (Cambridge, MA: MIT Press, 2001).

ATHEY, S., “Monotone Comparative Statics under Uncertainty,” *Quarterly Journal of Economics* 117 (2002), 187–223.

BANK, S. A., “Devaluing Reform: The Derivatives Market and the Executive Compensation,” *DePaul Business Law Journal* 7 (1995), 301–32.

BETTIS, J. C., J. M. BIZJAK, AND M. L. LEMMON, “Managerial Ownership, Incentive Contracting and the Use of Zero Cost Collars and Equity Swaps by Corporate Insiders,” *Journal of Financial and Quantitative Analysis* 36 (2001), 345–70.

BISIN, A., AND D. GUAITOLI, “Moral Hazard and Nonexclusive Contracts,” *Rand Journal of Economics* 35 (2004), 306–28.

—, P. GOTTARDI, AND A. A. RAMPINI, “Managerial Hedging and Portfolio Monitoring,” *Journal of the European Economic Association* 6 (2008), 158–209.

BOLDRIN, M., N. KIYOTAKI AND R. WRIGHT, “A Dynamic Equilibrium Model of Search, Production and Exchange,” *Journal of Economic Dynamics and Control* 17 (1993), 723–57.

BOSE, A., D. PAL, AND D. E. M. SAPPINGTON, “On the Performance of Linear Contracts,” *Journal of Economics & Management Strategy* 20 (2011), 159–93.

BREEDEN, D. T., AND S. VISWANATHAN, “Why Do Firms Hedge? An Asymmetric Information Model,” Mimeo, Duke University, 1998.

BROWN, C., “Firms’ Choice of Method of Pay,” *Industrial and Labor Relations Review* 43 (1990), 165–82.

BURDETT, K., AND J. I. ONDRICH, “How Changes in Labor Demand Affect Unemployed Workers,” *Journal of Labor Economics* 3 (1985), 1–10.

CAHUC, P., AND A. ZYLBERBERG, *Labor Economics* (Cambridge, MA: MIT Press, 2004).

ÇELEN, B., AND S. ÖZERTÜRK, “Implications of Executive Hedge Markets for Firm Value Maximization,” *Journal of Economics & Management Strategy* 16 (2007), 319–49.

DEGROOT, M. H., *Optimal Statistical Decisions* (Hoboken, NJ: Wiley Interscience, 1970).

- DEVINE, T. J., AND N. M. KIEFER, *Empirical Labor Economics: The Search Approach* (New York: Oxford University Press, 1991).
- EASTERBROOK, F. H., "Derivative Securities and Corporate Governance," *The University of Chicago Law Review* 69 (2002), 733–47.
- FROOT, K. A., D. S. SCHARFSTEIN, AND J. C. STEIN, "Risk Management: Coordinating Corporate Investment and Financing Policies," *Journal of Finance* 48 (1993), 1629–58.
- GARVEY, G., AND T. MILBOURN, "Incentive Compensation When Executives Can Hedge the Market: Evidence of Relative Performance Evaluation in the Cross-Section," *Journal of Finance* 58 (2003), 1557–82.
- GARVEY, G. T., "The Principal-Agent Problem When the Agent Has Access to Outside Markets," *Economics Letters* 43 (1993), 183–6.
- GIBBONS, R., AND K. J. MURPHY, "Relative Performance Evaluation for Chief Executive Officers," *Industrial and Labor Relations Review* 43 (1990), 30S–51S.
- GROSSMAN, S., AND O. HART, "An Analysis of the Principal-Agent Problem," *Econometrica* 51 (1983), 7–45.
- HARRIS, M., AND A. RAVIV, "Optimal Incentive Contracts with Imperfect Information," *Journal of Economic Theory* 20 (1979), 231–59.
- HOLMSTRÖM, B., "Moral Hazard and Observability," *Bell Journal of Economics* 10 (1979), 74–91.
- , "Moral Hazard in Teams," *Bell Journal of Economics* 13 (1982), 324–40.
- , AND P. MILGROM, "Aggregation and Linearity in the Provision of Intertemporal Incentive," *Econometrica* 55 (1987), 303–28.
- , AND J. TIROLE, "Liquidity and Risk Management," *Journal of Money, Credit and Banking* 32 (2000), 295–319.
- JIN, L., "CEO Compensation, Diversification and Incentives," *Journal of Financial Economics* 66 (2002), 29–63.
- KAHN, C., AND D. MOOKHERJEE, "Competition and Incentives with Nonexclusive Contracts," *Rand Journal of Economics* 29 (1998), 443–65.
- , AND —, "Market Failure and Moral Hazard with Side Trading," *Journal of Public Economics* 58 (1995), 159–84.
- LAVALLE, L., "Commentary: Undermining Pay Performance," *Business Week* 15 (2001), 70–71.
- LAZEAR, E. P., "Salaries and Piece Rate," *Journal of Business* 59 (1986), 405–31.
- LOBERG, J., AND R. WRIGHT, "Unemployment Insurance, Taxes, and Unemployment," *Canadian Journal of Economics* 20 (1987), 36–54.
- MILGROM, P., AND J. ROBERTS, *Economics, Organization & Management* (Englewood Cliffs, NJ: Prentice Hall, 1992).
- ÖZERTÜRK, S., "Managerial Risk Reduction, Incentives and Firm Value," *Economic Theory* 27 (2006), 523–35.
- PARK, I.-U., "Moral Hazard Contracting and Private Credit Markets," *Econometrica* 72 (2004), 701–46.
- PARLOUR, C. A., AND U. RAJAN, "Competition in Loan Contracts," *American Economic Review* 91 (2001), 1311–28.
- ROGERSON, W., R. SHIMER, AND R. WRIGHT, "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature* 43 (2005), 959–88.
- ROGERSON, W. P., "Repeated Moral Hazard," *Econometrica* 53 (1985), 69–76.
- ROSS, S. A., "The Economic Theory of Agency: The Principal's Problem," *American Economic Review* 63 (1973), 134–9.
- SCHIZER, D. M., "Executives and Hedging: The Fragile Foundations of Incentive Compatibility," *Columbia Law Review* 100 (2000), 440–504.
- SMITH, C. W., AND R. M. STULZ, "The Determinants of Firms Hedging Policies," *Journal of Financial and Quantitative Analysis* 20 (1985), 391–405.
- STULZ, R. M., "Optimal Hedging Policies," *Journal of Financial and Quantitative Analysis* 19 (1984), 127–40.
- , "Managerial Discretion and Optimal Financing Policies," *Journal of Financial Economics* 26 (1990), 3–27.
- , "Rethinking Risk Management," *Journal of Applied Corporate Finance* 9 (1996), 8–24.
- TUFANO, P., "Who Manages Risk? An Empirical Examination of Risk Management Practices in the Gold Mining Industry," *Journal of Finance* 51 (1996), 1097–137.