

Why Do Firms Use Incentives That Have No Incentive Effects?

PAUL OYER*

ABSTRACT

This paper illustrates why firms might choose to implement stock option plans or other pay instruments that reward “luck.” I consider a model where adjusting compensation contracts is costly and where employees’ outside opportunities are correlated with their firms’ performance. The model may help to explain the use and recent rise of broad-based stock option plans, as well as other financial instruments, even when these pay plans have no effect on employees’ on-the-job behavior. The model suggests that agency theory’s often-overlooked participation constraint may be an important determinant of some common compensation schemes, particularly for employees below the highest executive ranks.

MANY COMPENSATION PLANS REWARD or punish employees for factors they cannot control. An often-discussed example of this phenomenon is executive compensation (where stock options are not generally indexed to the overall market). However, many firms also offer firm-wide stock options and profit sharing plans that provide even weaker incentives than executive plans—after all, most workers can expect to reap trivial personal gain from their contribution to firm value or profits. I consider an explanation for this phenomenon in which firms contract with employees as a means of indexing wages to market rates rather than to provide incentives.

While much of the agency literature has concentrated on inducing optimal effort, I consider the importance of the generally overlooked participation constraint. I study a model in which employees’ outside opportunities are correlated with firm profits or stock price, and in which both turnover and adjusting the pay scheme parameters are costly. Given these assumptions, the firm may find it most profitable to pay the employee in a way that is correlated with the outside options presented by the outside labor market (despite the requisite risk premium) rather than pay a fixed wage that insures participation in all states.

The model derived below starts from the assumption that wage adjustments are costly, but that workers are willing to make part of their pay contingent

*Oyer is with the Graduate School of Business, Stanford University. I thank George Baker, Kenneth Corts, Adlai Fisher, Robert Gibbons, Richard Green, Phanwadee Khananusapkul, Edward Lazear, Bentley MacLeod, Kevin J. Murphy, Canice Prendergast, James Rebitzer, Scott Schaefer, Andy Skrzypacz, Charles Thomas, Jan Zabochnik, a referee, and seminar participants at Harvard Business School, Northwestern, University of Rochester, University of Michigan, USC, the 2000 Stanford Personnel Economics Conference, the 2000 Econometric Society World Congress, the 2001 AEA Meetings, and the 2001 Western Finance Association meetings for helpful comments.

on firm performance if they are compensated for the corresponding risk. After signing a contract with the firm, the employee receives an outside offer that varies with the state of the economy. Depending on the particular circumstances, the firm will choose one of three contracts: First, the firm may initially pay a fixed wage. In this case, upon being faced with the outside offer, the firm adjusts its wage to the market wage, though doing so may lead to information gathering, negotiation, and turnover costs. Second, the firm can structure the original contract such that enough of the employee's pay is contingent on firm value that the worker's ex ante participation constraint is always met exactly. Because firm value is imperfectly correlated with the economy (and hence with the worker's outside options), pay based on firm performance will not, ex post, be the same as the worker's market wage at the time the contract was arranged. This passes some risk along to the employee. Finally, the firm can make some amount of pay contingent on firm performance, but lower the risk premium by setting the guaranteed portion of compensation such that the worker earns rents in some states of the economy. The firm has to find the optimal trade-off between the risk premium associated with variable compensation and the rents associated with fixed compensation. The firm can opt to use a spot labor market if these two costs become too onerous, though it incurs other transaction costs if it chooses to do so.

The model provides a potential explanation for the puzzle of why firms do not filter common risk out of executive compensation so that pay will be more directly tied to factors that executives can control (see Section V.B). If nonindexed equity compensation varies with executives' outside opportunities, firms may find it optimal to let executive pay vary with general market conditions. However, the model raises just the opposite of the executive compensation indexing puzzle when looking at broad-based stock option plans for rank-and-file employees. That is, in the case of executives, the puzzle is why firms do not index equity pay so as to filter out market-wide risk. But if the model's suggestion that firms use stock options to keep employees' compensation in line with market rates is correct, then the puzzle for nonexecutives is why firms do not base compensation *solely* on an index so as to shelter their employees from uncontrollable idiosyncratic risk? I discuss several possible explanations for this, including some theoretical extensions of the basic model and practical considerations in the specific case of broad-based stock option plans. I also discuss examples of firms that have created financial instruments that more closely tie employees' pay to their labor market opportunities.

In addition to this contribution to the study of appropriate rewards for market and idiosyncratic risk, the paper also addresses the labor and contract theory literatures. The model's emphasis on the "stabilization" effect of pay based on firm performance is similar to the Weitzman (1984) share economy, though I focus on optimal firm choices rather than on social planning. Also, the model in this paper formalizes much of the intuitive discussion (and is consistent with the empirical findings) of sharecropping in Alston and Higgs (1982). I start with a simplified version of Harris and Holmstrom's (1982) model. As in their work, I take employees' on-the-job behavior as fixed and analyze the firm's

optimal contract for insuring the employees' participation. While Harris and Holmstrom focus on learning the employee's actual ability, I study the effects of changes in labor demand.¹ Lazear's (2001) model is similar to the model in this paper, in that it addresses the issue of why firms would want to make pay sensitive to firm performance if that pay has no effect on employees' on-the-job actions. His model focuses on top managers and how the firm can use pay to extract information from an employee.

The rest of this paper proceeds as follows. Sections I and II, respectively, describe the model and three potentially optimal contracts. Section III analyzes comparative statics. I explore how the choice of pay instruments and, if relevant, the amount of firm-performance-based pay is affected by such factors as the magnitude of common market shocks and idiosyncratic firm shocks, the range of workers' reservation wage, and employees' risk aversion. I also discuss the empirical implications of the model. Section IV briefly considers a few extensions of the basic model. I examine how the model may help in explaining the observed patterns in broad-based stock option plans and executive compensation in Section V.² I discuss the model's ability to explain a pay phenomenon that is at odds with previous agency theories—the high level of stock options in risky industries such as the technology sector. I also relate the model to the rise of tracking stocks and to venture capital funds used by some large technology companies. Finally, Section VI summarizes and discusses possible extensions.

I. The Model

I consider the use of a firm-wide performance measure in compensation contracts with individual employees. Many large firms reward workers based on such measures. Though these contracts are typically described as “incentive” contracts, there are reasons to believe that many of these contracts are not adopted to motivate increased effort. For example, when individuals each own a very small share of a firm, the incentives to free ride are likely to dominate the returns to increased effort.³ The value of firm-performance-based contingent pay is further reduced in any case where a firm can monitor employees well enough to assure some lower bound on effort. Assuming convex employee cost of effort, the marginal cost of effort beyond this lower bound is likely to be higher than the marginal benefits available from employee ownership. Also,

¹ See Thomas and Worrall (1988) and Beaudry and DiNardo (1991) for other models in which wages are sensitive to market rates and firms balance compensation and turnover costs.

² While I focus on labor market settings in this paper, the model can be applied to other contracts in which apparent “incentive” clauses may actually represent a convenient means of pricing. See Lafontaine and Masten (2002) for an application of this idea in the trucking business.

³ Holmstrom (1982), Rasmusen (1987), MacLeod (1988), Legros and Matthews (1993), and others develop models where the free rider problem can be overcome in groups. While these models apply in some settings, they are unlikely to completely overcome free riding issues in the contracts discussed in Sections V.A and V.C. Using a detailed data set of stock option grants by over 200 firms, Oyer and Schaefer (2004) quantify the negligible incentives provided to middle managers by broad-based stock option plans.




Figure 1. Timeline.

firms typically have access to individual performance measures (either objective or subjective) that can provide much stronger incentives while causing employees to bear less risk. For these reasons, I ignore the incentive effects of contingent pay throughout the model.

I consider a two-period economy with identical firms and workers (employees). I assume that firms are risk neutral. Firms can offer a contract for the first period and promise the employee the option of renewing the same contract in the second period. Firms pay all profits as dividends at the end of each period.⁴ Workers cannot make binding commitments and there is no market in which they can insure their intertemporal labor market risk.⁵

The firm's gross profits (i.e., before employee compensation) in period i are θ^i if there is a worker during the production phase, so net profits are θ^i less payments to the worker. Short of closing down, there is nothing the firm or worker can do to affect expected gross profits. Maximizing profits is equivalent to minimizing labor costs. The employee has constant absolute risk aversion.

The model unfolds as follows (see timeline in Figure 1):

1. Firms and the employee observe the spot market wage (the employee's reservation utility) for the first period, $s^1 \in \{s_h, s_l\}$ (where $s_h > s_l$), which depends on the expected state of the economy for period 1. The firm offers the employee a contract that specifies a fixed wage (w) and a fractional share (b) of the firm's gross profits. The employee accepts the offer if his expected utility over both periods is at least as great under the contract (given that he has the option of leaving after the first period) as it is if he rejects the offer.
2. Gross profits, θ^1 , are revealed and the worker is paid $b\theta^1$. Gross profits are drawn from a normal distribution with variance σ^2 and expected value θ_h , if $s^1 = s_h$ and expected value θ_l (where $\theta_h > \theta_l$) if $s^1 = s_l$.

⁴ I use the term "employee ownership" to refer to employment contracts based on firm-wide performance. In the model, "profit sharing" is the same as employee ownership because the firm distributes all profits each period.

⁵ I ignore discounting. Previous versions of the model included discounting, which added some complexity but no insight.

3. The employee receives an outside offer that reflects the second period spot market wage. The employee's expected utility if he takes the offer is $s^2 \in \{s_h, s_l\}$.⁶ Let $q = \Pr(s^2 = s_h | s^1)$. By incurring costs of k , the firm can either adjust the terms of the contract, or if the worker leaves, replace the worker at a wage of s_h or s_l .⁷ These costs, which create a quasi-rent for the firm and/or the employee, may reflect the cost of search, loss of specific human capital acquired during the first period, or other costs.
4. Second period profits, with expected value based on the second period spot market wage, are revealed and the employee is paid his share, as in step 2.

The firm can pay the worker, either in part or in full, through a fixed per-period wage (w). The variable w can depend on s^1 , but contracts written in period 1 cannot be made contingent on s^2 . Also, it is costly to adjust w once s^2 is revealed. The firm can pay a fixed wage of s_h in order to guarantee the employee's participation regardless of outside offers, or the parties can write a contract based on the value of θ . Therefore, an employment contract consists of a fixed wage (w) paid in all states and a share of profits (b). I rule out contracts in which the firm prespecifies w_1 and w_2 , where $w_1 \neq w_2$. I make this restriction because while the model is limited to two periods, I hope it captures some longer term arrangements where all parameters cannot be specified in advance. Partially relaxing this assumption leads to some interesting additional results (see Section IV.A).

I refer to $\theta_h - \theta_l$, or $\Delta\theta$, the range of the *expected* outcomes, as the "common shock range." I refer to σ^2 , the variance of firm profits conditional on the realization of s , as "idiosyncratic volatility." As σ^2 decreases, the idiosyncratic shock becomes less volatile and the market wage is a better predictor of the firm's prospects. If θ^i is perfectly predictable given s^i , then $\sigma^2 = 0$ and none of the volatility is idiosyncratic.

By assuming that pay can be made contingent only on θ , I am assuming that, among the measures on which the employee is willing to base a contract, θ is the measure that is most correlated with spot market wages. This could be because workers will not accept contracts with some portion of pay based on a stock market or industry index; the firm's success is more indicative of the worker's reservation wage than a broader measure (perhaps because other firms infer the firm's workers' average ability from θ); or the bargaining costs to define the

⁶ In the interests of focusing on the most relevant issues, I have simplified the outside offer process and I assume that turnover is always inefficient. A potential enhancement of the model would have workers receive a random draw from an outside offer distribution. Then the firm would have to choose pay parameters so as to minimize costs of compensation, turnover, and renegotiation. Such an enhancement would be a labor market application similar to Rochet and Stole (2002). Also, while the outside offers in this model are determined solely by the state of the economy, others have explored the private information used in or revealed by outside offers. See, for example, Lazear (1986).

⁷ For some firms, replacing workers may be less expensive than adjusting the terms of contracts with current workers, while contract adjustments may be the preferable alternative at other firms. The variable k should therefore be interpreted as the minimum of replacement and renegotiation costs.

proper index are so great that the firm and employee agree to contract on θ . Note that, for simplicity, I assume that the contract is linear in θ . This will not necessarily be optimal, but it allows me to compare fixed wages with a form of (if not *the best* form of) contingent pay.

Because the employee has constant absolute risk aversion (with a coefficient of r) and his income is drawn from a normal distribution, the period i certainty equivalent of the employee's utility (CE), conditional on s^i , can be expressed $CE^i = w + bE[\theta^i | s^i] - rb^2\sigma^2$. To avoid paying the adjustment cost while insuring that the employee will continue to work for the firm in the second period regardless of s^2 , the pair w and b must meet two participation constraints. If $s^2 = s_h$, then the worker will stay with the firm if

$$s_h \leq w + b\theta^h - rb^2\sigma^2. \quad (1)$$

Similarly, if $s^2 = s_l$, then the worker will stay if

$$s_l \leq w + b\theta^l - rb^2\sigma^2. \quad (2)$$

Hereafter, I refer to (1) and (2) as the "bull constraint" and the "bear constraint," respectively. Under the maintained assumption that the firm cannot stipulate different wages in a contract arranged at the beginning of the first period, satisfying (1) and (2) insures that the employee will participate in the first period as well. Any contract that satisfies the bull constraint (if $s^1 = s_h$) or the bear constraint (if $s^1 = s_l$) will satisfy the employee, because his first period utility is at least as high as his market alternative and he can resort to the market in the second period.

The firm can use two different strategies to make both the bull and bear constraints bind. First, it can adjust the wage to exactly meet the outside offer. Alternatively, it can select a positive b such that the employee's expected second period wages exactly match his reservation wage when $s^2 = s_h$ and when $s^2 = s_l$. Because both of these choices impose costs on the firm, it may prove profitable to pay the employee in such a way that his participation constraint does not always bind.

Initial contract offers differ according to the state of the first period economy. In discussing the possible forms the employment contract may take in Section II and comparative statics in Sections III.A-III.C, I assume that $s^1 = s_h$. I focus on the good state for two reasons. First, workers are more likely to be hired when firms' prospects, and the state of the economy in general, are relatively promising. Second, one reason that is often cited for the difficulty of adjusting wages to the spot market is that wages are downwardly rigid. That is, one interpretation of k is that it is costly to the firm to decrease wages, either because lowering wages requires the firm to replace the current employees or because wage reductions hurt productivity (see Bewley (2000)). The sign of all relationships and comparative statics that follow are unaffected when $s^1 = s_l$, with one exception that I discuss below. I return to the effects of s^1 and discuss the effect of allowing k to depend on the *direction* in which w is changed in Section III.D.

II. Three Possible Outcomes

There are three possible classes of contracts the firm may choose to offer the employee. Depending on the parameters of the model, any of these three may be optimal from the firm's perspective. In this section, I describe each of these three regimes and then consider the conditions under which any particular outcome will be profit-maximizing.

A. Spot Labor Markets

Given $s^1 = s_h$, the firm may choose to offer a wage of s_h and then spend k to adjust wages down to s_l if $s^2 = s_l$. In such circumstances, the firm is transacting in a spot labor market. The firm pays no risk premium or rents to the worker. However, sometimes the worker imposes an externality on the firm by leaving or demanding to renegotiate. The profits of a spot labor market strategy are gross profits minus wages minus transaction costs. Expected gross profits, which are not affected by the employment contract, are θ_h in period 1 and $q\theta_h + (1 - q)\theta_l$ in period 2. Wages are s_h in period 1, but they change to s_l with probability $1 - q$ in period 2. Transaction costs are k times the probability that wages have to be adjusted. Total profits over the two periods in a spot market regime can therefore be expressed

$$\Pi^{\text{sm}} = (1 + q)(\theta_h - s_h) + (1 - q)(\theta_l - s_l) - (1 - q)k. \tag{3}$$

B. Dual Binding Participation Constraints (DPC)

Spot markets lead to transaction costs. The firm can avoid these costs if it chooses w and b such that, for both realizations of s^2 , the employee is exactly indifferent between the choice of staying at the firm and leaving. That is, the firm can choose the wage and a share of profits such that the bull and bear participation constraints (i.e., (1) and (2)) hold with equality.⁸ Combining the two constraints when they bind yields

$$s_h - s_l = b(\theta_h - \theta_l) \tag{4}$$

and (let Δ denote the difference between variables in the h and l states) the optimal employee share is therefore

$$b = \frac{\Delta s}{\Delta \theta}. \tag{5}$$

⁸Note that I assume that σ^2 is not related to s . This makes the model much more tractable. However, it is possible that better economies (that is, those with spot wages of s^h) are more variable. I relaxed this assumption and then solved for the optimal contract using many parameter sets. I then determined how the contract changed when the main parameters of interest varied. The results were entirely consistent with the results in the rest of the paper, with a minor exception that is noted in Section III.C.

The employee's share increases with the difference between the two possible reservation utilities because the firm wants to lower the base wage it has to pay when spot wages are low. This result assumes that the firm's optimal type of contract is unaffected by changes to the exogenous parameters. Changes could lead to a new optimal contract in which the employee's share no longer solves equation (5).

In a case where two participation constraints bind (DPC hereafter), profits and the guaranteed wage are

$$\Pi^{\text{dpc}} = (1 + q)(\theta_h - s_h) + (1 - q)(\theta_l - s_l) - \frac{2r\sigma^2(\Delta s)^2}{(\Delta\theta)^2} \quad (6)$$

and

$$w = s_h - \frac{\theta_h \Delta s}{\Delta\theta} + \frac{r\sigma^2(\Delta s)^2}{(\Delta\theta)^2}. \quad (7)$$

The term that includes $r\sigma^2$ in wages and profits represents the risk premium necessary to induce the worker to accept a given share. If the risk cost is prohibitive (e.g., when r is large), then the firm may choose to use the spot labor market (as discussed in the previous subsection) or it may choose not to keep the worker bound by his participation constraint in all states (as discussed in the following subsection).

C. Single Binding Participation Constraint (SPC)

The previous subsections considered two ways in which the firm can extract all the surplus from its relationship with the employee. However, these strategies involve either transaction costs (spot markets) or compensating the worker for risk (DPC). There may be situations in which the firm is better off with a less extreme strategy. If the firm sets w and b so as to maximize profits while insuring the worker's participation when spot wages are high, the firm may choose to pay rents to the worker when spot wages are low. This is because, were the firm to drive the worker down to his participation constraint when spot wages are low, it would have to pay him more in risk premium than the expected rents if the worker is overpaid when spot wages are low. I refer to an optimal contract in which a single participation constraint binds as an SPC contract.⁹

Note that the firm can avoid all transaction and risk costs by setting $w = s_h$ and $b = 0$. Such a contract would satisfy the bull constraint exactly and enable the employee to earn rents of Δs in period 2 if $s^2 = s_l$. However, it will generally

⁹ It is trivial to show that the firm would never offer a contract where the participation constraint only binds when spot wages are low. Note that, assuming that there are multiple identical firms, the "rents" I discuss would disappear as each firm would offer the same contract. Strictly speaking, these rents refer to increases in the agent's utility relative to an economy where only the spot labor market is available.

be optimal for the firm to save some of these rents by passing along some of the risk to the employee.

The firm will choose b and w to maximize $E[\Pi] = (1 + q)(1 - b)\theta_h + (1 - q)(1 - b)\theta_l - 2w$, subject to the bull constraint. After generating the first-order condition and rearranging terms, the optimal employee share can be written

$$b = \frac{(1 - q)\Delta\theta}{4r\sigma^2}. \tag{8}$$

The employee's share is decreasing in σ^2 because a greater idiosyncratic shock transfers more risk to the employee. The employee's share is also decreasing in r , because the firm finds it cost effective to shoulder more of the risk when the employee is more risk averse.

After some manipulation, profits and the wage under an SPC contract can be expressed as

$$\Pi^{\text{SPC}} = 2(\theta_h - s_h) - (1 - q)\Delta\theta + \frac{(1 - q)^2(\Delta\theta)^2}{8r\sigma^2} \tag{9}$$

and

$$w = s_h - \frac{\Delta\theta(1 - q)[(3 + q)\theta_h + (1 - q)\theta_l]}{16r\sigma^2}. \tag{10}$$

The profit and wage terms now include both a risk premium and rents that the worker receives if spot wages are low. The relative importance of transaction costs, rents, and the risk premium determines which of the three possible contract types is optimal (and feasible), as discussed in the next subsection.

D. Which of the Three Possibilities Will Hold?

When does the optimal contract make only the bull participation constraint bind and when does it make both the bull and bear constraints bind? If the bear constraint does not bind when b is the optimal SPC employee share, an SPC regime will hold. This is because in order to get the bear participation constraint to bind, w would have to be lowered and b would have to increase. But any change in the employee share from equation (8) lowers profits, and the firm will prefer to pay the worker the necessary rents rather than impose additional risk on the employee. However, if the bear constraint does not hold when b is the optimal SPC employee share, then the best option available to the firm is to pay the full risk premium necessary to get both participation constraints to bind.

By combining the bull and bear constraints, substituting the optimal SPC employee share, and rearranging terms, I find that the optimal SPC contract is feasible (and therefore dominates the optimal DPC contract) when

$$\Delta s > \frac{(1 - q)(\Delta\theta)^2}{4r\sigma^2}. \tag{11}$$

The trade-off between risk costs and rents, which drives (11), applies only when choosing between the regimes with no adjustment costs (i.e., DPC and SPC). If the profits under the optimal zero adjustment plan are lower than those in a spot market, then the firm will opt for the spot market. The firm's preferred contract emerges from combining and comparing the previously derived profit functions for each of the three types of contracts, Π^{sm} , Π^{dpc} , and Π^{spc} . The spot labor market is preferred to the DPC contract when

$$\Delta s > \frac{\Delta\theta}{\sigma} \sqrt{\frac{k(1-q)}{2r}}. \quad (12)$$

Combining (9) and (3), while insuring that the bear constraint is satisfied, I find that an SPC contract is optimal when

$$\frac{(1-q)(\Delta\theta)^2}{4r\sigma^2} < \Delta s < k + \frac{(1-q)(\Delta\theta)^2}{8r\sigma^2}. \quad (13)$$

III. Comparative Statics

A. Common and Idiosyncratic Shocks

This section considers how the shocks that are common to all firms in this labor market ($\Delta\theta$) and idiosyncratic to the individual firm (σ^2) affect the firm's optimal contract. I establish three main results. First, employee ownership becomes more valuable as the range of common shocks increases and as the idiosyncratic shock becomes less variable. Second, if employee ownership is optimal, a greater common shock or a less variable idiosyncratic shock induces the firm to lower (and eventually eliminate) the rents it shares. The third finding is that as the shock common to all employers in this labor market rises from zero to very large or as the idiosyncratic shock diminishes from very large to zero, the employee's share is weakly increasing in common shock and weakly decreasing in idiosyncratic shock.

PROPOSITION 1:

1. If spot wages are unrelated to firm value (i.e., $\theta_h = \theta_l$), the firm sets $b = 0$.
2. A DPC contract is optimal below some strictly positive threshold idiosyncratic shock volatility. That is, there exists $\sigma_D > 0$ such that for any $\sigma < \sigma_D$, the firm will choose a DPC contract with $b > 0$.
3. Spot markets are optimal above some threshold idiosyncratic shock volatility. That is, there exists σ_{SM} such that for any $\sigma > \sigma_{SM}$, the firm will set $b = 0$ and, if necessary, incur cost of k to set $w_2 = s^2$.
4. If $\Delta s < 2k$, then $\sigma_D < \sigma_{SM}$ and an SPC contract is optimal for all σ such that $\sigma_D < \sigma < \sigma_{SM}$.
5. If $\Delta s \geq 2k$, then $\sigma_D = \sigma_{SM}$ and an SPC contract is never optimal.
6. A DPC contract is optimal above some threshold common shock range; spot markets are optimal below some threshold common shock range; and if $\Delta s < 2k$, an SPC contract is optimal between the DPC and spot market

thresholds. That is, there exists $\Delta\theta_D$ such that the firm will choose a DPC for all $\Delta\theta > \Delta\theta_D$; there exists $\Delta\theta_{SM}$ such that the firm will use the spot market for all $\Delta\theta < \Delta\theta_{SM}$; and, if $\Delta s < 2k$, $\Delta\theta_D > \Delta\theta_{SM}$ and an SPC contract is optimal for all $\Delta\theta$ such that $\Delta\theta_D > \Delta\theta > \Delta\theta_{SM}$.

Proof: If $\theta_h = \theta_l$, then because $s_h > s_l$, (1) can only hold if (2) does not bind. Therefore, (8) implies $b = 0$, which proves part 1. Define $M = (1 - q)(\Delta\theta)^2 / 2r$, $\sigma_1^2 = M/2(\Delta s)$, $\sigma_2^2 = M/4(\Delta s - k)$, and $\sigma_3^2 = Mk/(\Delta s)^2$. Equations (12) and (11) imply that a DPC contract will be optimal for any $\sigma^2 < \min(\sigma_1^2, \sigma_3^2)$. Letting $\sigma_D = \min(\sigma_1, \sigma_3)$ proves part 2. Equations (12) and (13) imply that spot markets are optimal for any $\sigma^2 > \max(\sigma_2^2, \sigma_3^2)$. Letting $\sigma_{SM} = \max(\sigma_2, \sigma_3)$ proves part 3. Manipulation of (13) shows that an SPC contract is optimal and feasible if and only if $\Delta s < 2k$. Note that if $\Delta s < 2k$, $\sigma_1^2 < \sigma_3^2 < \sigma_2^2$. Combining (11), (12), and (13), a DPC contract is optimal for $\sigma < \sigma_1$, an SPC contract is optimal for any σ where $\sigma_1 < \sigma < \sigma_2$, and spot markets are optimal for $\sigma > \sigma_2$, concluding the proof of part 4. Now note that if $\Delta s > 2k$, $\sigma_2^2 > \sigma_3^2 > \sigma_1^2$. In this case, (13) cannot hold. Part 5 follows from the fact that a DPC contract will therefore be optimal for any $\sigma < \sigma_3$ and that spot markets will be optimal for $\sigma \geq \sigma_3$. Part 6 can be proven by the same logic as parts 2–5. Q.E.D.

The proposition starts by establishing the somewhat obvious, but intuitively appealing, result that the firm would not want to tie the employee’s pay to a measure that is not correlated with either output or outside opportunities. At the other extreme, consider the case in which spot wages resolve all uncertainty—i.e., $\sigma^2 = 0$. Then the risk premium disappears, and the solution is to pay a fixed wage of s_l and let the employee earn Δs through his ownership interest when profits are high. Because profits are high only when the spot wages are high, workers are paid exactly their reservation wage with no risk premium and no turnover. As part 2 of the proposition shows, for any sufficiently low idiosyncratic volatility, the firm will not pay the worker any rents.

There can be two transitions in regime as the idiosyncratic shock improves from highly volatile to insignificant (or as shocks affecting firms in this labor market move from uncorrelated to large and homogeneous)—from spot markets to SPC to DPC. For firms with volatile idiosyncratic shocks, employee ownership is a relatively expensive means of retaining workers. In fact, firms may choose to incur wage adjustment costs periodically, rather than pay the necessary rents during bad times or the risk premium that workers would demand to take a share of the firm. Inequality (13) suggests that, for a volatile enough firm (i.e., high enough σ^2), the firm will opt to use spot markets if the difference between the two reservation wages gets sufficiently large or if the costs of adjusting the wage are small. That is, spot markets become a more attractive option when the firm faces a large idiosyncratic shock, when s_l decreases, and when k decreases. But as the firm’s shock gets smaller (meaning that spot market wages become more informative), the risk premium decreases. The firm may decide to pay a risk premium and some rents to the worker in an SPC regime rather than pay the adjustment costs associated with using spot markets.

The second transition can occur as the shock becomes less volatile if the firm is in an SPC state. A less volatile shock leads the firm to give the employee a greater stake at a lower premium, driving the worker to be bound by the lower participation constraint.

PROPOSITION 2: *The employee's share decreases with idiosyncratic shock volatility (σ^2) and increases in common shock range ($\Delta\theta$) in an SPC state. The employee's share is constant in σ^2 and $\Delta\theta$ in a DPC state. In either state, profits (total payments to the worker) are decreasing (increasing) in idiosyncratic shock volatility and increasing (decreasing) in common shock range. If $b > 0$, w is increasing in idiosyncratic shock volatility and decreasing in common shock range.*

The proof follows from simple comparative statics and is omitted. The firm engages in spot markets with no employee ownership when common shocks are small or when idiosyncratic shocks are large because the turnover costs are outweighed by the risk-sharing and rents necessary to retain employees. Then the employee's share (b) takes a discrete jump up when the firm switches to an SPC regime. The employee's share continues to increase with common shock or with decreases in idiosyncratic shocks because the employee is willing to bear more of the risk as the labor market becomes more closely related to the firm's outcome. But then the employee's share hits a point at which the lower participation constraint binds. Further reductions in the idiosyncratic shock do not affect the employee's share, but allow the firm to lower the risk premium by lowering the guaranteed wage. This orthogonal relationship between contingent pay and idiosyncratic risk, as well as the positive relationship between common shocks and contingent pay in an SPC regime, distinguishes the model from standard agency models that predict a universally negative correlation between risk and incentive pay. As discussed in Section V.A, this may help explain why firms in relatively risky industries are relatively likely to offer stock options to all employees. An example of the effect of idiosyncratic risk on the employee's share is shown graphically in Figure 2.¹⁰

Because idiosyncratic volatility induces rents and risk compensation, profits decrease with idiosyncratic volatility. Figure 3 shows which factors contribute to the monotonic decline in profits as idiosyncratic risk increases. In the DPC regime, the risk premium paid to the employee increases with idiosyncratic volatility. When idiosyncratic volatility becomes high enough that the firm switches to an SPC contract, the firm buys down some of the risk premium with rents through salary increases. Above some threshold idiosyncratic volatility, the combined rents and risk premium exceed expected turnover costs and the firm chooses to use the spot market.

The findings in this section lead to potentially testable empirical predictions. In particular, within firms, industry-specific human capital is more valuable in some occupations than in others. When markets for certain types of labor are industry-specific, the firms in that labor market are liable to be subject to

¹⁰ All figures assume that $k = 0.35$, $q = 0.5$, $\theta_h = 2$, and $\theta_l = 0$. When not being used as the varying parameter, $r = 1$, $s_h = 1.25$, $s_l = 0.75$, and $\sigma^2 = 1$.




Figure 2. Idiosyncratic risk and employee's share. Shows the optimal level of employee ownership as shocks that are idiosyncratic to the firm (and therefore not related to market wages) change. The type of contract offered by the firm also changes with idiosyncratic volatility. At low enough levels, the firm shares ownership, but pays no rents (a DPC contract). At high levels, the firm abandons employee ownership and engages in the spot labor market.

more common shocks than are firms in broader labor markets. Consider a software company. The engineers in the company are likely to have technology and software-specific skills. Therefore, the firm's profit level or stock price will be subject to significant shocks common among the other firms that might employ these engineers. Though profit sharing or stock options require the engineers to bear substantial risk, the incremental risk relative to what they would face in the spot market may not be large. But the skills of the firm's accountants are not tied as closely to the software industry, and as a result, employee ownership will not be as effective, because it would force a lot more risk on these workers than would the spot market. This argument suggests that within-firm variations in employee ownership should be related to how concentrated occupations are within the firm's industry. Also, broad employee ownership should be more prevalent in industries that dominate the labor market for the occupations of a large share of their employees.

Similar to these cross-occupation implications, geographic variation may be used to analyze the model. Certain geographic areas have more concentrations of one or a few industries than others. As a result, all workers at the firms in the concentrated regions are likely to be in a labor market affected by the same industry forces that affect their employers' profitability. So while a software firm's accountants may have more idiosyncratic risk, on average, than its




Figure 3. Idiosyncratic risk effect on costs. Shows how various components of labor cost vary with idiosyncratic volatility. Turnover costs occur only when the firm engages in spot markets. “Rents” only occur under SPC contracts when the risk premium necessary to make the participation constraint bind in all cases becomes prohibitively expensive.

software engineers, one might expect accountants in Silicon Valley or Detroit to have more common shocks and lower idiosyncratic risk (and therefore to be more likely to earn stock options or profit sharing) than accountants in more diversified economies like those in Chicago or Los Angeles.

B. Range of Reservation Utility

While shocks are an important consideration, they are irrelevant if the worker’s reservation utility is invariant. This section studies how the firm responds to changes in the range of the worker’s outside options (“reservation range”), which is captured in the model by Δs (i.e., $s_h - s_l$).¹¹ I establish three main results. First, as reservation utility can vary across a wider range, the firm is likely to transition from a DPC regime to an SPC regime, and ultimately, to spot markets. Second, as the range of reservation utility becomes greater, the employee’s share increases to a plateau, then stays constant before returning to zero for a high enough reservation range. Third, profits decrease in reservation range.

¹¹ Increases in reservation range refer to increases in Δs that preserve the mean reservation utility. That is, any change in s_h is coupled with an opposite change of $\frac{q}{1-q}$ in s_l . Note that the variance in reservation utility is $q(1-q)(\Delta s)^2$, so reservation range is monotonically related to variance in reservation utility for a fixed q .

PROPOSITION 3:

1. A DPC contract is optimal below some strictly positive reservation range. That is, there exists $\Delta s_D > 0$ such that for any $\Delta s < \Delta s_D$, the firm will choose a DPC contract with $b > 0$.
2. Spot markets are optimal above some threshold reservation range. That is, there exists Δs_{SM} such that for any $\Delta s > \Delta s_{SM}$, the firm will set $b = 0$ and, if necessary, incur the cost of k to set $w_2 = s^2$.
3. If $\frac{(1-q)(\Delta\theta)^2}{8r\sigma^2} < k$, then $\Delta s_D < \Delta s_{SM}$ and an SPC contract is optimal for all Δs such that $\Delta s_D < \Delta s < \Delta s_{SM}$.
4. If $\frac{(1-q)(\Delta\theta)^2}{8r\sigma^2} \geq k$, then $\Delta s_D = \Delta s_{SM}$ and an SPC contract is never optimal.

Proof: Define $M = \frac{\Delta\theta}{\sigma} \sqrt{\frac{(1-q)}{2r}}$, $s_a = \frac{M^2}{2}$, $s_b = k + \frac{M^2}{4}$, and $s_c = M\sqrt{k}$. Equations (12) and (11) imply that a DPC contract will be optimal for any $\Delta s < \min(s_a, s_c)$. Letting $s_D = \min(s_a, s_c)$ proves part 1. Similarly, (12) and (13) imply that spot markets are optimal for any $\Delta s > \max(s_b, s_c)$. Letting $s_D = \max(s_a, s_c)$ proves part 2. Manipulation of (13) shows that an SPC contract is optimal and feasible if and only if $\frac{M^2}{4} < k$. Note that if $\frac{M^2}{4} < k$, $s_a < s_c < s_b$. Combining (11), (12), and (13), a DPC contract is optimal for $\Delta s < s_a$; an SPC contract is optimal for any Δs where $s_a < \Delta s < s_c$; and spot markets are optimal for $\Delta s > s_b$, concluding the proof of part 3. Now note that if $\frac{M^2}{4} > k$, $s_c > s_b > s_a$. In this case, (13) cannot hold. A DPC contract will therefore be optimal for any $\Delta s < s_c$ and spot markets will be optimal for $\Delta s \geq s_c$. Q.E.D.

The proposition establishes that employee ownership is a suboptimal strategy if the spread in reservation utilities is extreme. This is because a high range of possible outside offers leads, ultimately, to higher rents paid to the employee in slow times. It is rents, rather than a risk premium, that lead to spot markets, because the firm always has the option of just paying the high reservation utility in all states and thus can eliminate the risk premium.

The proposition also shows that a DPC will be in place for a low enough reservation range because the employee can be held to his indifference point in all states without much of a risk premium. But beyond some threshold reservation range, it is worth paying the rents in an SPC regime rather than impose the risk costs needed to make the employee's participation constraint bind in all states (though this threshold may be *after* spot markets become optimal.) Taken together, these clauses of the proposition suggest that as the reservation range increases from zero to extreme, the firm is likely to transition from a DPC regime to an SPC regime to spot markets. I now explore the profit and employee ownership implications of this evolution.

PROPOSITION 4:

1. The optimal employee share (b) weakly increases in the reservation range for reservation ranges below the level where spot markets become optimal.

2. *Firm profits monotonically decrease in the reservation range. Profits are convex in the reservation range in a DPC regime and are linear in the reservation range in an SPC regime.*

Again, the proof follows from simple differentiation and is omitted. The profit result is intuitively simple—increasing the range of the employee’s outside options leads to either an increase in the risk premium (in a DPC regime) or in the employee’s rent (in an SPC regime). Either one of these effects increases the payments to the worker and lowers the firm’s profits. The quadratic risk premium drives the effect on profits in a DPC regime, leading labor costs to increase at an increasing rate. But once the regime shifts to SPC, increases in reservation range cause the firm to pay a constant q units of rents in bad states for every unit increase in reservation range. This increase in rents as reservation range increases is similar to Harris and Holmstrom’s (1982) finding that “insurance” increases with variance in worker ability. In both cases, the firm finds it optimal to protect risk averse employees from bad realizations of their market wage.¹²

The employee’s share increases with reservation range in a DPC regime because more variance in pay is required to keep the worker on both participation constraints as the difference in reservation utility widens. Once the firm lets the lower participation constraint go slack, the optimal trade-off between risk and rents is unrelated to the spread between the reservation utilities. As a result, the employee’s share increases with initial increases in reservation range, reaches a plateau, and then employee ownership is abandoned altogether when the rents overwhelm the adjustment costs associated with spot markets. Figure 4 graphs the effect of reservation range on the employee’s share.

These results generate the empirical implication that firms that employ people in occupations or industries in which the market wages fluctuate widely are likely to have greater employee ownership than more stable industries. However, if the variation becomes too great, employee ownership is less likely to be adopted.

C. Risk Aversion

The level of employee risk aversion is critical to the model. If the employee is risk neutral (i.e., $r = 0$), then the firm can ensure participation in both states at minimal cost. Note that risk aversion has exactly the same mathematical effects as idiosyncratic shock volatility throughout the model. I therefore state the key intuition regarding risk aversion, because the calculations and graphical relationships are the same as those in Section III.A.

At low enough levels of risk aversion, the risk premium necessary to keep a worker bound by both participation constraints is not great. As the employee

¹² Note that the profit/reservation range relationship is the one result that differs when $s^0 = s_1$. If the first-period economy is weak, then a mean preserving spread in the reservation range can (under certain parameter values) lower the reservation utility enough to compensate for the added risk.




Figure 4. Reservation range and the employee's share. Shows how the optimal level of employee ownership and the form of the employment contract change as the amount of variation in market wages changes. When wage variation is low enough, the firm shares ownership, but pays no rents (a DPC contract). At high levels, the firm abandons employee ownership and engages in the spot labor market.

becomes more risk averse, the risk premium increases relative to the rents that the firm would have to pay the worker if it let the lower participation constraint go slack. Further increases in risk aversion make the risk premium so high that employee ownership eventually gives way to spot markets.

Equation (5) implies that the employee's share is not affected by risk aversion in a DPC.¹³ However, as r continues to increase, the risk premium in a DPC regime eventually becomes sufficiently high that the firm prefers the rents of an SPC to the risk premium needed to keep the employee bound by both participation constraints. Once the SPC regime is in place, further increases in r require further reductions in the employee's share as the firm trades even more rents to keep the risk premium in line. Eventually, if the rents are too large, then the firm may move to a spot market.

D. The Initial State of the Economy

I now consider how the compensation plans would differ from those in Section II if employment contracts are written when the economy is relatively

¹³ This is the only result that I have found to rely on the assumption that σ^2 is independent of s . When this assumption is relaxed, optimal b increases (decreases) with r if σ_h^2 is greater (less) than σ_l^2 .

weak (i.e, if $s^1 = s_l$).¹⁴ If k does not depend on the direction in which wages are changed, then when the firm uses spot markets, changing starting conditions only affects the firm by decreasing gross profits. The wages paid in any given state of the economy are the same as above. Similarly, the compensation parameters (b and w) of a DPC contract are unchanged because these values are driven by second-period reservation utility. The most important change in compensation is in an SPC. When the economy is weak in the first period, the expected rents in an SPC are higher and the firm chooses a higher b . Specifically, the firm will set

$$b = \frac{3\Delta\theta}{8r\sigma^2}. \quad (14)$$

Expected profits for all three contracts are lower when $s^1 = s_l$ than when $s^1 = s_h$. For spot and DPC contracts, the profit changes are exactly the same, but SPC profits fall further due to the extra rents. Therefore, while the condition under which the firm prefers spot markets to a DPC contract (12) is unchanged, the conditions that determine when an SPC is best ((11) and (13)) change somewhat. An SPC contract dominates a DPC contract when

$$\Delta s > \frac{3(\Delta\theta)^2}{8r\sigma^2}. \quad (15)$$

The firm will use an SPC contract when

$$\frac{3(\Delta\theta)^2}{8r\sigma^2} < \Delta s < \frac{k}{3} + \frac{3(\Delta\theta)^2}{16r\sigma^2}. \quad (16)$$

Because the SPC becomes less valuable if the contract is written during a weak economy, fewer firms will use employee ownership because spot markets are more profitable for a larger set of parameter values. Also, more firms will use DPC contracts rather than SPC contracts. Together, these changes suggest that when the economy is relatively weak, firms are less likely to issue shares to employees, but conditional on using employee ownership, shares will be as high or higher a share of pay than for similar idiosyncratic volatility and risk aversion levels in a strong economy.

One interpretation of k is that the firm cannot lower wages and must replace its workforce in order to do so, or at least it is costly to lower wages due to morale, influence activity, and other factors. As a result, k may be a function of the direction of the change in wages, with k being larger for wage decreases than for wage increases. That is, k may be greater when $s^1 = s_h$ (in which case any

¹⁴ Throughout this section, I assume that $q = 0.5$. By assuming that the two states of the economy are equally likely (which is equivalent to the less restrictive assumption that any serial correlation in the state of the economy is unrelated to s^1), I can abstract from the effects of the extreme likelihood of switching states. It is possible to imagine markets in which this assumption is unreasonable, however.




Figure 5. Effects of initial state and wage rigidity. Like Figure 2, this shows the optimal level of employee ownership as idiosyncratic shocks change. The solid line is the same as the line in Figure 2. The dashed line shows the same relationship, but assumes that the employment agreement is formed during unfavorable economic conditions. The line with dots and dashes shows this relationship under the assumptions that market wages are initially low and that adjusting wages upward is less costly than reducing wages. This reduction in adjustment costs leads to a smaller region of idiosyncratic shocks for which the firm chooses to share ownership with the employee.

change to w would be downward) than when $s^1 = s_l$.¹⁵ An economy-dependent k would therefore make spot markets even more attractive relative to employee ownership when $s^1 = s_l$. Figure 5 contrasts the optimal contract at various idiosyncratic volatilities under three scenarios: (1) $s^1 = s_h$ (which is a re-creation of Figure 2); (2) $s^1 = s_l$ and nothing else changes; and (3) $s^1 = s_l$ and k is lower than in (1) and (2). Given the set of parameters used in the graph, an SPC is never optimal. The graph and the discussion above lead to the empirical prediction that firms are less likely to have employee owners in a weak economy

¹⁵ Allowing k to differ with the direction of the wage change may introduce an opportunity for the firm. For some parameter values, the optimal contract would set first-period wages below market if the contract is signed when the bull constraint holds ($s^1 = s_h$). Second-period wages would rise to s_h if the bull constraint continued to hold, but the firm would not lower wages in an unfavorable second period if k were larger than the wage savings. This could reduce the range of parameter values over which employee ownership is optimal. But the only effect on the previous comparative statics discussions would be to make employee ownership nonoptimal below some threshold reservation range. Note that this upwardly flexible contract would become less relevant if more periods were added to the model.

than in a strong one, and that this effect will be exacerbated when wages are downwardly rigid.

IV. Extensions

A. Time-Varying Wages

To this point, I have assumed that the firm cannot specify wages for the two periods where $w_1 \neq w_2$ without paying k in adjustment costs. Allowing the firm to prespecify different wages for the two periods will under some parameter values allow the firm and worker to avoid risk and renegotiation costs by setting $w_2 = s_h$ and w_1 , such that $w_2 + w_1 = s^1 + qs_h + (1 - q)s_l$. That is, the firm will promise to pay the worker his expected reservation utility over the two periods and remove his labor market risk. This arrangement, by paying the worker at or above his spot wage in the second period and below his spot wage in the first period, is similar to the contract in Lazear (1979). In effect, the employee posts a bond of $(1 - q)\Delta s$ in the first period and is repaid this amount in the second period.

Several factors can keep this contract from being feasible or optimal, however. First, the firm may not be able to credibly commit to pay a second-period wage above the market wage. This would be particularly likely to be true in any case where $s_h - s_l > k$. In such cases, the firm would be better off firing the worker and replacing him if $s^2 = s_l$. This provides another reason to think that the incidence of employee ownership will increase in Δs . Second, if the utility function were altered such that the worker had preferences for smooth consumption, a sharp increase in wages between the two periods might be utility-decreasing. Third, if the worker accumulates additional general human capital in the process of performing his job, then the firm may reach some first-period liability limitation if it tries to extract all rents through the first-period wage. Finally, if in addition to being correlated with spot market wages, firm profits are correlated with the employee's marginal product at that firm, prespecifying wages might inhibit efficient separations.

I now briefly consider a variant on the base model in which the firm can specify w_1 and w_2 at the beginning of period 1 where $w_1 + d \geq w_2$. In other words, the firm can credibly promise to repay a bond up to some limit. This variation turns out to have fairly minor effects on most of the previous results, so I do not derive all the details.

In any case where $d \geq s_h - s_l$, the firm can write a two-period contract that eliminates all costs of risk, turnover, and rents. Therefore, an additional comparative static of the model is that, as the ability to commit to future cash wages increases, the value of employee ownership diminishes. This suggests that relatively young firms that have not had a chance to establish a reputation for honoring implicit wage commitments and firms whose future existence is questionable may choose employee ownership because they do not have the option of writing a long-term contract.

If $d < s_h - s_l$, the firm cannot avoid paying k (with some probability) without employee ownership or rents. Therefore, the firm can implement the same three

possible contracts (spot markets, SPC, and DPC) as in the basic model when d was assumed to be 0. The optimal spot market contract and DPC contract are the same when the firm can prespecify different w_1 and w_2 as they are when $d = 0$. To see why the spot market contract is unaffected, note that if the firm may have to adjust wages anyway, then the optimal spot market contract sets $w = s$ for each period. The DPC contract is unaffected because any contract that keeps the employee exactly at his participation constraint in all states requires setting b independently of the base wage.

The main benefit (to the firm) of a positive d is in the case of an SPC contract. The firm will set $w_2 = w_1 + d$. This increases profits by d relative to the base model. However, because nothing changed *on the margin*, it does not change the optimal employee share (b). That is, the optimal trade-off between risk costs and rents is not affected by the opportunity to commit to a wage increase.

The main effects of allowing wages for the two periods to vary in a prespecified manner are summarized in Figure 6. The graph redraws the employee share/idiosyncratic volatility relationship in Figure 2, adding lines for various levels of d . For high enough d , the firm can commit to a second-period wage




Figure 6. Effects of commitment to wage increase. Like Figure 2, this shows the optimal level of employee ownership as idiosyncratic shocks change. The solid line is the same as the line in Figure 2. The other lines show the same relationship, but while allowing the firm to prespecify a second-period wage of up to d greater than the first-period wage. For a high enough d , the firm can commit to a second-period wage that insures the employee's labor market risk and there is no employee ownership. For lower levels of d , the firm pays lower rents under a DPC contract relative to when $d = 0$. As a result, the firm prefers employee ownership to spot markets for a larger set of idiosyncratic volatilities.

such that it can insure the employee's labor market risk. In this case, there is no employee ownership. For lower values of d , the value to the firm of employee ownership actually increases because it can lower the level of rents it has to pay the employee in an SPC. As a result, an SPC becomes preferable to a spot market for some higher levels of idiosyncratic volatility.

The results in this section suggest that allowing the firm to commit to a wage structure (i.e., allowing the worker to post a bond) can have two effects. If the firm's ability to commit is limited, commitment enhances the value of employee ownership. This is especially true at risky firms where commitment to a wage increase can enhance the value of employee ownership by lowering rents paid to the employee. If the firm has more flexibility in the amount it can commit to a wage increase (or if spot market wages do not vary much), then employee ownership becomes unnecessary and the firm can eliminate turnover costs, risk costs, and rents. Taken together, these results provide further potential justification for the use of employee ownership in risky industries (such as stock options at high-tech firms), because employee ownership is more likely to be adopted when high variance in future prospects limits commitment capability, when idiosyncratic risk is high, and when spot labor market variation is high.

B. Correlated Firm Performance and Employee Value

The main model assumes that the employee's labor market prospects are independent of the firm's performance *relative to the rest of the market*. I now consider a case in which working for a successful firm makes an employee more valuable. This could happen, for example, if firms make bets on various technologies and employees who are trained in those technologies that ultimately prove successful are more valuable. Alternatively, if part of what makes a firm successful is its ability to identify talented employees, the market's estimate of a person's ability could be related to the success of his current employer. Formally, suppose that the market's expectation of the employee's second-period productivity is an increasing function of θ^1 . Then the employee's outside offer at the beginning of the second period would be increasing in his firm's first-period profit (θ^1). Now assume that the firm waits until the end of the second period before paying out profits.

In this case, the employee's second-period outside offer (s^2) would be a function of his firm's first-period performance. But his expected compensation if he stayed at the initial employer would also be correlated with the firm's first-period performance and his second-period market opportunities. In cases where this correlation between firm performance and employee market value exists, it would mitigate some of the costs of idiosyncratic shocks discussed in Section III.A. Also, as discussed in Section V.A below, firms that issue stock options broadly impose significant firm-specific risks on their employees. But if firm performance is related to those employees' outside opportunities, that firm-specific risk may not be overly costly.

C. Incentives Created by Vesting

Throughout the paper to this point, I have assumed that pay based on firm performance has no effect on the employee's behavior. However, suppose the firm pays out profits at the end of two periods and that the firm can fire the employee at the end of the first period if the worker is observed to be "shirking" in period 1. Suppose further that "working" (as opposed to shirking) is costly to the employee, but is efficient in the sense that working creates more benefits for the firm than costs to the employee. Though not critical, also suppose that searching for a new job is more costly to the employee in periods when spot wages are low.

If these conditions hold, then employee ownership can have the additional benefits of providing marginal incentives when they are most valuable. When labor market conditions are favorable, the threat of being fired may provide relatively weaker incentives. However, the value of unvested payments to the employee will be high when the other costs of being fired are low. On the other hand, slack labor market conditions provide incentive when the expected value of unvested employee shares is low. In other words, the firm may be able to use employee ownership to tie the value of unvested payments to labor market conditions in a way that provides consistent incentives at a relatively low cost.

V. Applications

A. Employee Stock Options

Many high-tech companies give stock options to all employees. At a small company, options can have powerful incentive effects. However, it is hard to imagine that most new workers at Microsoft and Cisco Systems, both of which have over 35,000 workers and offer stock options to all of them, believe that their actions affect the value of their stock options. So why do these and other firms offer options as part of their standard compensation plan?¹⁶

In this section, I argue that the model in this paper can help explain this phenomenon. I first discuss how some relevant empirical facts are consistent with the comparative static predictions of Section III and with the importance of turnover costs (k). I then consider the relevance of vesting patterns in employee stock option grants. Finally, I address why a firm might use options *in its own stock* rather than tying pay to a market index that might better reflect its workers' outside opportunities and shelter employees from the firm's idiosyncratic volatility.

As noted in Sections III.A and III.B, a key distinguishing feature of the model is that higher variance in the overall fortunes of firms in a given labor market

¹⁶ Recently, Microsoft announced major changes to its broad-based equity plan. See Section III.D for a possible explanation of why broad-based plans would become less prevalent when labor market competition becomes less intense.

(either in the form of large common industry shocks or high reservation range as firms' demand for labor comes and goes with their markets) can lead to more contingent pay. While most agency models predict relatively little contingent pay in industries such as high-tech, where there are large industry shocks to firm performance and market wage, the model in this paper is consistent with observed patterns. Oyer and Schaefer (2004) provide empirical evidence consistent with this relationship between industry shocks and industry volatility. They show that broad-based option plans are more common in industries with higher average volatility. They also show that option plans are more common in industries in which a typical firm's return is more closely correlated with the industry average return.

A few examples highlight this relationship. The average R^2 in a regression of a firm's monthly stock return on the average return of other firms in its 2-digit industry is 0.1202 in Oyer and Schaefer's sample of 1,000 randomly selected public firms. The industries with the highest proportion of broad-based option plans are SIC codes 35 (industrial machines, including computers), 36 (electronic equipment), and 73 (business services, including software). The average industry regression R^2 in those industries are 0.1367, 0.2019, and 0.1986, respectively. The R^2 in the industries with the lowest proportion of broad plans (textiles, primary metals, and fabricated metals) are 0.1052, 0.1800, and 0.1129, respectively. Similarly, the average industry volatilities in the industries with more broad plans are 0.0732, 0.1008, and 0.1131, respectively. The volatilities in the rare-broad-plan industries are 0.0530, 0.0616, and 0.0459 and the average for all industries is 0.0605. Though these examples are not conclusive, they illustrate that volatility and common shocks are generally higher in industries where broad option grants are found more frequently.

There was an economy-wide increase in stock-based pay in the late 1990s (see Biddle (2000)). The analysis in Section III.D may help explain this trend. During the robust economy of the late 1990s, spot market wages were high. In the model, this makes stock options more attractive. Also, if the tight job market of the late 1990s was indicative of high k , then this could help explain why stock options became a more widely used pay instrument in that period.

The fact that stock options plans are so common among technology firms in the small geographic area of Silicon Valley is also consistent with the model's implications regarding turnover costs. Until very recently, brisk demand for technology workers made workers mobile while making it difficult for firms to fill openings (that is, high k). The conditions under which employee ownership yields higher profit than spot markets ((12) and (13)) suggest that high k makes stock options profitable relative to spot markets.

The model relies critically on the notion that the employee gives up something by leaving the firm. In practical terms, this suggests that the model requires options to vest over a period of time and for the value of those options that have not yet vested to vary with the state of the economy. It is a well-established empirical fact that employee stock options typically vest over a period of a few years. It is also well documented that employees typically exercise their options shortly

after they vest if the options are in-the-money.¹⁷ Oyer and Schaefer (2004) show that the value of unvested grants to middle managers in their sample is quite substantial and varies significantly with common shocks. Specifically, a middle manager at a typical firm in their sample has an annual cash salary of \$90,000. After one year on the job, the value of that manager's unvested options would be approximately \$70,000 if the firm's industry did well in that year, but only \$15,000 if the industry had a bad year. At a firm in the sample that uses options more liberally, the manager's unvested option wealth varies from \$183,000 in a good economy to \$27,000 in a slow period. These numbers suggest that the value of unvested options is likely to have a significant effect on employees' propensity to take an outside offer.

To the extent that the model in this paper applies to broad-based stock option plans, it also raises a question related to a puzzle in the use of stock options in executive compensation. As discussed in the next section, there has been significant recent attention among financial economists to the question of why firms do not insulate their executives from market-wide risk that they cannot influence. The puzzle for lower-level employees, however, is just the opposite—why would firms expose their employees to substantial idiosyncratic firm risk when those employees' market wages might be better reflected by some broader index? I offer several possible answers to this question.

First, firms have the option to lower the strike price of employee stock options in the event of a negative firm-specific shock. Though the empirical work on the repricing of stock options has focused on executives, the results are supportive of the model. Specifically, Carter and Lynch (2001) show that firms reprice executive stock options after *firm-specific* shocks (that is, shocks that do not affect market demand for their employees' services). Firms do not, however, reprice in response to poor *industry* performance. They conclude that patterns in option repricing are consistent with firms attempting to "prevent management in competitive labor markets from going to work for other firms" (p. 222).

Section IV.B provides an alternative explanation of why a firm might issue options in its own stock rather than pay based on a broader index. If other potential employers infer that employees of successful firms are likely to be of higher ability, then the market wages of those employees will be a function both of the market wages for similar employees and the firm's idiosyncratic performance. In this case, stock options may be a more cost-effective contracting tool than a market or industry index.

There may also be practical barriers that make stock option contracts preferable to contracts based on an index. For example, if an appropriate index would

¹⁷ Oyer and Schaefer (2004), for example, find that, in a sample of firms that grant options to a majority of employees, most use options that vest over a four year period and most options are exercised within a year of vesting. Huddart and Lang (1996) study eight firms with broad option plans, all of which grant options that vest over three to five years. They also show that most options are exercised soon after vesting. Using a somewhat older and larger sample (478 firms), Aboody (1996) finds similar vesting and exercise patterns.

be a small set of similar firms, an employer may be reluctant to give its employees stock in its competitors. Also, the appropriate index could be different for various sets of employees in the same firm. Rather than try to tailor the appropriate instrument to the appropriate group, it may be simpler (and therefore less costly in terms of administration and negotiations) to use a universal metric for the whole firm. Finally, to the extent that managers value higher current earnings, the favorable accounting treatment of option grants may sway managers to use them over other tools.

Having provided some reasons why firms may choose options, it is also worth noting that at least some firms have used creative financial instruments that appear to be designed to make employment contracts more efficient. For example, according to the *Economist* (2000a), "DLJ direct, an online brokerage, boasts that selling its tracking stock has reduced turnover among technical staff from 30 to 5%, even though its share price is now below its issue price" (p. 76). While the firm could have offered its technical workers options in the parent company (Donaldson, Lufkin, and Jenrette) stock, this would have exposed them to the ups and downs of the brokerage and financial services markets. But the technical staff was more likely to be courted by other internet businesses, so the parent company's stock experienced significant shocks that are not common to other firms that might employ the online workers. These shocks make stock options expensive. However, by issuing a tracking stock that more accurately reflects the internet economy, and by extension, the technical workers' market opportunities, the firm was able to better incorporate relevant common labor market shocks into its stock price while lowering idiosyncratic shocks. This allowed them to retain employees more cost effectively.

Another example (and one that cannot possibly square with any incentive story) of the creation of a financial instrument as a means of contracting with employees involves venture capital investing by large firms. Several Wall Street firms, in conjunction with large technology firms such as Cisco Systems and Intel, have started venture capital pools that invest in the types of startups that often lure away key personnel from these large technology companies (see *Economist* (2000b) and Rabinovitz (2000)). The firms then offer shares in these funds to the people they are most afraid of losing, conditional on continued employment. In the model's terms, firms are attempting to decrease σ^2 rather than pay the costs to renegotiate or replace workers. One impediment to these funds has been allotting shares in the funds, which is another type of negotiation or transaction cost. In the framework of the model, these financial instruments involve the firm paying some amount of negotiation/transaction fees (that is, some k , but less than losing the person) to decrease σ^2 . Because profits are decreasing in σ^2 , if the cost in terms of k is not too extreme, then the firms stand to gain from these venture funds.

B. Executive Compensation

Are executive compensation contracts consistent with agency theory? That is an often-studied and widely argued question. Models of efficient risk sharing

typically imply that wherever possible, executives should be measured on relative, rather than absolute, performance. That is, to the extent that some underlying shock affects performance across multiple firms in a way that individual executives cannot affect, firms will find it profitable to filter this common shock out of the measures of executive performance. Most executive stock options and other executive incentives do not adjust for market-wide or industry-wide shocks, and empirical studies have found little evidence of relative performance evaluation in executive pay (see Murphy (1999)). A series of recent papers, each of which explores a specific version of the general issue of distortions in relative performance evaluation noted by Gibbons and Murphy (1990), have presented potential explanations for this lack of indexing. Notable examples include Garvey and Milbourn (2003) (who argue that executives can hedge the market risk if they choose to), Aggarwal and Samwick (1999) (who argue that, under some market structures, indexed options would drive managers to compete in a way that would lower firm value), and Meulbroek (2002) (who considers the practical limitations of indexing). However, perhaps one reason relative performance evaluation is not more common is that the models that show it to be optimal typically assume a constant participation constraint.¹⁸ The variable participation constraint presented here may lead firms to reward executives for absolute performance so that compensation varies with the executives' outside opportunities.

Bertrand and Mullainathan (2001) analyze the efficiency of CEO incentive contracts, starting with the hypothesis that an efficient contract will reward (punish) high (low) *relative* performance. They study the oil industry, in which common shocks (especially commodity prices) can be easily identified and removed from measures of executive performance. They show that the pay of oil company CEOs is affected by oil prices, a factor that is surely beyond the executives' control.¹⁹ They go on to establish evidence that CEOs are rewarded for factors beyond their control in a range of industries.

The model presented here, with relabeling of some variables, suggests that it can be optimal to pay CEOs for industry-level performance if industry performance is correlated with the executives' outside opportunities. Suppose that a high oil price (or a high oil futures price) leads to increased demand for oil executive talent, either because of entry or because incumbent firms launch new projects. This should raise the demand for and price of managerial talent in the oil industry. In the terms of the model, high oil prices cause the expected value of new ventures to be θ_h at the same time as they raise a manager's reservation wage to s_h . Paying the manager for the firm's absolute performance (θ), rather than for the firm's relative performance ($\theta - \theta_{\text{market}}$) could be an efficient way to

¹⁸ Himmelberg and Hubbard (2000) develop a model that is similar in spirit to the model above. Rather than considering idiosyncratic shocks and reservation variance, they focus on how an executive's outside opportunities are related to the size of the firm he manages. Using chief executive officer compensation data, they then go on to find empirical evidence that supports their model, as well as the model in this paper.

¹⁹ Some oil companies do, however, make at least some effort to filter out the effect of commodity prices and other factors beyond managers' control (see, e.g., Podolny and Roberts (1998)).

provide incentives for the manager to work hard while minimizing the chances that he will leave for another oil producer.²⁰

Barro and Barro (1990) present further evidence consistent with the model. They find that bank CEO pay is responsive to firm performance but is not adjusted for relative performance, which implies that bank CEO compensation is partially based on conditions in the overall banking market. They also find that bank CEO turnover does respond to relative performance. These facts are consistent with banks letting their compensation schemes fluctuate so as keep pace with market rates, while dismissing executives when firm performance lags behind peer performance.

C. Traditional Profit Sharing

Profit sharing plans are common in a variety of industries, even among large firms where the incentive effects are trivial (see Kruse (1993)). Why would a firm impose the risk of profit sharing on its employees? The model in this paper, as well as that of Weitzman (1984), justify profit sharing based on what Kruse (1993) calls the “stability theory.” When wages are rigid, profit sharing is a means of keeping workers at or near their participation constraints. Kruse suggests that there is some empirical support for this idea, but that it is inconclusive in most studies. He points out that identifying stability effects is difficult given the endogeneity of profit sharing.

Kruse (1993) also found that the best predictor of the adoption of profit sharing is an increase in the profitability of the firm. This is consistent with rent sharing between the firm and workers, but it may also be consistent with the model in this paper. If a firm’s prospects improve, then the value of retaining workers increases, which increases k in the model. As discussed above, increasing k increases the value of profit sharing relative to using the spot labor market. Also note that this contrasts with the empirical implications of Weitzman (1984). While the stability effect of profit sharing in the above model comes from minimizing voluntary turnover, stability in Weitzman derives from the firm not laying off workers in a downturn. The latter idea would make profit sharing relatively more valuable when the firm’s prospects dim.

In their study of Australian firms, Drago and Heywood (1995) find that profit sharing is relatively common at firms whose workers are highly skilled and who have invested in firm-specific human capital (as measured by proportion of managers, proportion of “casual” workers, use of productivity groups, and a propensity to promote from within). If these workers are relatively expensive to replace, then these findings are consistent with profit sharing being more common at firms with relatively high k . Drago and Heywood also find that

²⁰ Bertrand and Mullainathan (2001) briefly discuss the issue of oil CEO reservation utility. They argue that such a model could explain part of their findings but is unlikely to completely explain their results. This argument is particularly compelling when they show that CEOs with weaker corporate governance gain more from “luck” than do CEOs with tight oversight. However, this result is also consistent with more talented CEOs endogenously weakening corporate governance (see Hermalin and Weisbach (1998)).

turnover is negatively associated with profit sharing, which is consistent with profit sharing having a stabilizing effect.²¹

VI. Conclusions and Further Research

I derive a model that provides a possible reconciliation of agency theory with the fact that employees are often rewarded or punished for things they cannot control. I show that variability in an employee's reservation utility may lead the firm to want to transfer risk to the worker as a means of insuring participation during various states of the economy. The model relies critically on the assumption that the terms of an employee's compensation scheme are costly to adjust. I also determine that the firm may find it optimal to pay rents to workers during relatively slow periods and, under certain conditions, may prefer to forsake long-term retention of employees.

The model yields some specific testable hypotheses that are generally consistent with anecdotal evidence and with the empirical analysis in Oyer and Schaefer (2004). The empirical implications of the model include:

1. The adoption of employee ownership increases as the firm's costs of replacing workers increases.
2. As shocks common to the firms employing a set of workers increase and volatility of shocks idiosyncratic to an individual firm decrease, the *adoption* of employee ownership becomes more attractive and the *amount* of employee ownership increases (weakly).
3. Greater variation in local market wages leads to an increase in the *amount* of employee ownership, though extreme variation discourages employee ownership.
4. Employee ownership is relatively attractive in strong economies and tight labor markets.

I briefly consider a few extensions to the model, but other extensions could add further insights. For example, the model in this paper assumes that it is always efficient for the firm to retain the worker in all states of the world. In reality, an employee's marginal product at his current employer may change relative to alternative employers. It would be interesting to investigate the implications of a model in which firm profits or value are correlated with the employee's marginal products *at that firm* as well as (or instead of) with his alternative employment opportunities. Firms and workers may choose to write contracts based on firm performance so as to promote efficient turnover. Such a model would allow outside offers to be drawn from a distribution, rather than having only two possible values of s . This would lead the firm to have to consider the optimal level of turnover, given that it would probably not want to match

²¹ Another common form of contractible variable pay is index-linked wage rates. As Card (1986) shows, these contracts can allow wages to adjust to *the firm's* reservation wage. His discussion of the link between industry prices and the CPI can be interpreted as a labor demand version of idiosyncratic shocks.

all outside offers. A model along these lines might provide further justification for using pay based on the employer's performance, rather than tying pay to some market index.

REFERENCES

- Aboody, David, 1996, Market valuation of employee stock options, *Journal of Accounting and Economics* 22, 357–391.
- Aggarwal, Rajesh, and Andrew Samwick, 1999, Executive compensation, strategic competition, and relative performance evaluation: Theory and evidence, *Journal of Finance* 54, 1999–2043.
- Alston, Lee J., and Robert Higgs, 1982, Contractual mix in Southern agriculture since the Civil War: Facts, hypotheses, and tests, *Journal of Economic History* 42, 327–353.
- Barro, Jason R., and Robert J. Barro, 1990, Pay, performance, and turnover of bank CEOs, *Journal of Labor Economics* 8, 448–481.
- Beaudry, Paul, and John DiNardo, 1991, The effect of implicit contracts on the movement of wages over the business cycle: Evidence from micro data, *Journal of Political Economy* 99, 665–688.
- Bertrand, Marianne, and Sendhil Mullainathan, 2001, Do CEOs set their own pay? The ones without principals do, *Quarterly Journal of Economics* 116, 901–932.
- Bewley, Truman F., 2000, *Why Wages Don't Fall During a Recession* (Harvard University Press, Cambridge, MA).
- Biddle, Frederic M., 2000, A special news report about life on the job—and trends taking shape there, *The Wall Street Journal*, January 18, A1.
- Card, David, 1986, An empirical model of wage indexation provisions in union contracts, *Journal of Political Economy* 94, S144–S175.
- Carter, Mary Ellen, and Luann J. Lynch, 2001, An examination of executive stock option repricing, *Journal of Financial Economics* 61, 207–225.
- Drago, Robert, and John S. Heywood, 1995, The choice of payment schemes: Australian establishment data, *Industrial Relations* 34, 507–531.
- Economist*, 2000a, The real meaning of empowerment, March 25, 75–78.
- Economist*, 2000b, Money to burn, May 27, 71–73.
- Garvey, Gerald T., and Todd T. Milbourn, 2003, Incentive compensation when executives can hedge the market: Evidence of RPE in the cross-section, *Journal of Finance* 58, 1557–1581.
- Gibbons, Robert, and Kevin J. Murphy, 1990, Relative performance evaluation for chief executive officers, *Industrial and Labor Relations Review* 43, 30s–51s.
- Harris, Milton, and Bengt Holmstrom, 1982, A theory of wage dynamics, *Review of Economic Studies* 49, 315–333.
- Hermalin, Benjamin E., and Michael S. Weisbach, 1998, Endogenously chosen boards of directors and their monitoring of the CEO, *American Economic Review* 88, 96–118.
- Himmelberg, Charles P., and R. Glenn Hubbard, 2000, Incentive pay and the market for CEOs: An analysis of pay-for-performance sensitivity, Columbia University.
- Holmstrom, Bengt, 1982, Moral hazard in teams, *Bell Journal of Economics* 13, 324–340.
- Huddart, Steven, and Mark Lang, 1996, Employee stock option exercises: An empirical analysis, *Journal of Accounting and Economics* 21, 5–43.
- Kruse, Douglas L., 1993, *Profit Sharing: Does it Make a Difference?* (W.E. Upjohn Institute for Employment Research, Kalamazoo, MI).
- Lafontaine, Francine, and Scott E. Masten, 2002, Contracting in the absence of specific investments and moral hazard: Understanding carrier-driver relations in U.S. trucking, Working paper, University of Michigan.
- Lazear, Edward P., 1979, Why is there mandatory retirement? *Journal of Political Economy* 87, 1261–1284.
- Lazear, Edward P., 1986, Raids and offer matching, *Research in Labor Economics* 8, 141–165.
- Lazear, Edward P., 2001, Output-based pay: Incentives, retention or sorting? Working paper, Stanford University.

- Legros, Patrick, and Steven A. Matthews, 1993, Efficient and nearly-efficient partnerships, *Review of Economic Studies* 60, 599–611.
- MacLeod, W. Bentley, 1988, Equity, efficiency, and incentives in cooperative teams, *Advances in the Economic Analysis of Participatory and Labor Managed Firms* 3, 5–23.
- Meulbroek, Lisa, 2002, Designing an option plan that rewards relative performance: Indexed options revisited, Working paper, Harvard Business School.
- Murphy, Kevin J., 1999, Executive compensation, in Orley Ashenfelter and David Card, eds.: *Handbook of Labor Economics*, vol. 3 (North-Holland, New York).
- Oyer, Paul, and Scott Schaefer, 2004, Why do some firms give stock options to all employees?: An empirical examination of alternative theories, *Journal of Financial Economics*, forthcoming.
- Podolny, Joel, and John Roberts, 1998, British Petroleum (A2): Organizing for performance at BPX, Graduate School of Business, Stanford University, Case Study IB16A2.
- Rabinovitz, Jonathan, 2000, Free bagels and a venture fund, *The Industry Standard*, May 15, 77.
- Rasmusen, Eric, 1987, Moral hazard in risk averse teams, *RAND Journal of Economics* 18, 428–435.
- Rochet, Jean-Charles, and Lars A. Stole, 2002, Nonlinear pricing with random participation constraints, *Review of Economic Studies* 69, 277–311.
- Thomas, Jonathan, and Tim Worrall, 1988, Self-enforcing wage contracts, *Review of Economic Studies* 55, 541–554.
- Weitzman, Martin L., 1984, *The Share Economy* (Harvard University Press, Cambridge, MA).

