

Incentives and Innovation: A Multi-tasking Approach*

Thomas Hellmann[†] and Veikko Thiele[‡]

August 2009

Abstract

This paper examines how employees trade off planned activities versus unplanned innovation, and how firms can choose incentives to affect these choices. It develops a multi-task model where employees make choices between their assigned standard tasks, for which the firm has a performance measure and provides incentives, and privately observed innovation opportunities that fall outside of the performance metrics, and require ex-post bargaining. The model shows how firms adapt incentive compensations in the presence of such unplanned innovation. If innovation are highly firm-specific, firms provide lower-powered incentives for standard tasks to encourage more innovation, yet in equilibrium employees undertake too few innovation. The opposite occurs if innovation are less firm-specific. We also investigate the effectiveness of several possibilities to encourage innovation, such as tolerance for failure, investing in employee innovation, stock-based compensation, and the allocation of intellectual property rights.

Keywords: Innovation, multi-tasking, incentives.

JEL classification: D82, D86, M52.

*We like to thank Jean-Etienne de Bettignies, Momo Deretic, Mati Dubrovinsky, Ulrich Hege, Christoph Lulfesmann, Gustavo Manso, Joeri Sol, Mikhel Tombak, Masako Ueda, and seminar participants at the University of British Columbia, Simon Fraser University, and participants of the Economics of Entrepreneurship and Innovation Conference in Kingston, the European Summer Symposium in Financial Markets in Gerzensee, the Regional Comparative Advantage and Knowledge-Based Entrepreneurship Conference in Amsterdam, and the Second Annual Economics and Law of the Entrepreneur Conference in Chicago for valuable comments and suggestions.

[†]University of British Columbia, Sauder School of Business, 2053 Main Mall, Vancouver, BC, Canada V6T 1Z2, phone: +1 604 822-8476, fax: +1 604 822-8477, e-mail: hellmann@sauder.ubc.ca.

[‡]Queen's University, Queen's School of Business, Goodes Hall, 143 Union Street, Kingston, Ontario, Canada K7L 3N6, phone: +1 613-533-2783, fax: +1 613-533-6589, e-mail: vthiele@business.queensu.ca.

1 Introduction

In a dynamic economy, firms' ability to innovate is essential for their competitive advantage, even survival. Yet the process of generating internal innovation frequently clashes with firms' established production and management processes. In this paper we develop a multi-tasking model of employees choosing between the execution of their assigned activities, which can be compensated according to a well-defined performance measure, versus the pursuit of previously unplanned innovation. Our central research question is under what circumstances employees deviate from their assigned activities to pursue an innovation, and how firm managers can optimally influence employees' innovation choices.

The tension between exploitation of existing assets versus exploration of new opportunities has long been recognized (March and Olsen, 1991). Established firms typically aim to maximize profits and efficiency by planning activities and creating measurable performance metrics. A large literature on optimal contracts examines optimal employee incentives and performance measurement (Gibbons 1998). However, almost by definition, innovation does not fit into such a structure. While there are some innovative activities that can be anticipated to a limited extent (e.g., incremental product improvement in R&D departments), many important business innovation cannot be planned in advance (e.g., identifying new types of customers, customizing product offerings, altering marketing and distribution arrangements, etc.). Moreover, employees are often the first to identify these new opportunities. As a consequence, innovation in established organizations is typically the result of employees taking initiatives that go beyond their narrow job description. Burgelman (2002), for example, describes how Intel's strategic move from computer memory to semiconductors was initiated by employees developing new technologies without the knowledge of senior managers—the so-called "skunkworks" approach. More generally, Burgelman points to the importance of an autonomous employee-led innovation process that occurs outside of firms' main strategic decision processes.

If innovation could be planned, then managers could base incentive contracts on measurable performance objectives, and the analysis would resemble standard multi-task incentive models (Holmström and Milgrom, 1991). A central idea of this paper is to introduce innovation as an unplanned activity, and to examine how it interacts with planned activities. Our theoretical starting point is to note that there exists a parallel debate in the theory of the firm. The work of Grossman, Hart and Moore (see Grossman and Hart, 1984, Hart 1995, Hart and Moore, 1990, 2008) postulates contractual incompleteness, based on the argument that it is often too complex and costly to write incentive contracts. By contrast, the work of Holmström, Milgrom and Roberts (Holmström, 1979, 1999, Holmström and Milgrom, 1994, Holmström and Roberts, 1998) argues that incentive contracting is at the heart of the firm's internal organization. We would argue that one should not focus on which view is "right or wrong"—clearly, there is a considerable amount of incentive

compensation inside firms, but these incentive contracts remain incomplete in important ways—but focus on which view is appropriate for which type of activity. Indeed, our premise is that the standard tasks (i.e., the planned activities) are well captured by a model with performance measurement and ex-ante incentive contracts. However, when it comes to innovative tasks (i.e., the unplanned activities), a model of incomplete contracts and ex-post bargaining is more appropriate (see also Aghion and Tirole 1994).

We then examine the interactions between the planned and unplanned activities, and show how the optimal level of incentive compensation for planned tasks is influenced by the presence of unplanned innovation. Integrating both planned and unplanned activities also allows us to evaluate whether and when employees are innovating too much or too little, relative to a variety of benchmarks.

We examine a principal-agent model where the principal is a risk-neutral firm owner and the agent is a risk-neutral and wealth-constrained employee (Sappington, 1983). The employee faces a multi-tasking choice between performing a standard task, for which it is possible to specify ex-ante a performance measure and therefore also an incentive compensation, versus pursuing an innovation for which there is no ex-ante performance measure, and hence only the possibility of ex-post contracting. We assume that the employee receives a private signal about the attractiveness of a potential innovation, and that the principal cannot observe the employee's signal, as well as his effort and task choice. We use a model specification based on Demougin and Fluet (2001) where the optimal contract can be expressed as a base wage and a performance bonus. The principal's only tool of influencing the agent's behavior is the strength of incentives, i.e., the size of the performance bonus for the standard task.

While the firm can determine an optimal compensation for the standard task, the rewards for innovation are determined by ex-post bargaining. In practice, many innovative employees receive no or little rewards for their inventions. One industry survey (IPOA, 2004) found that 84% of all companies did not pay employees at all for disclosing inventions. While many companies pay small amounts for patents (in a range of \$500-\$1500), 76% of companies were found not to make adjustments for "highly valuable" patents. Yet, in other circumstances, employees are able to extract considerable rents from their innovation, such as in the financial service industry, where the threat of leaving typically results in higher retention pay (Groysberg, Nanda, and Prats, 2007). In general, the rents employees extract for their innovation depend on outside options, private benefits, the degree to which the ideas can be appropriated by the firm, and possibly the allocation of intellectual property rights within the firm (Talley, 1998). Therefore, our model allows for a flexible parameterization that enables us to characterize any division of ex-post bargaining rents.

A fundamental insight from the model is that there is a conflict between innovation and the strength of incentives provided for the standard task: the higher the performance bonus, the less the employee pursues unplanned innovation. We examine how the firm can modify its incentive compensation for standard tasks,

taking the employee's innovation choices into account. If innovation is highly firm-specific we find that the employee pursues too little innovation relative to the firm's preferred level of innovation. In equilibrium, the firm reduces incentives for the standard tasks, in order to encourage the employee to innovate more. This result provides an interesting interpretation for frequently-observed low-powered incentives (Williamson, 1975, Parent, 2002). They provide "breathing room" for an employee who might want to pursue unplanned activities that will not be reflected in his performance measure. However, we also find that if innovation has low levels of firm-specificity, then the firm wants to discourage innovation ex-ante. It can do so by increasing its performance bonus beyond normal levels. In this case, high-powered incentives provide greater effort for the measurable task, but curtail the pursuit of innovation.

Our model also allows for a critical assessment of the question whether there is too much or too little innovation. This is actually a matter of considerable debate in the applied management literature. On the one hand there are many who lament the lack of innovation inside large firms.¹ On the other hand, there is also large number of management scholars, including Michael Porter (1980), who advocate strategic focus, or what Rotemberg and Saloner (1994) call a "narrow business strategy." Our model allows us to evaluate those claims from a variety of vantage points. We first explain that the firm owner and the employee almost never agree on the optimal level of innovation. For high levels of firm-specificity of innovation, the employee pursues too little innovation from the owner's perspective, but this is always reversed for low levels of firm-specificity. From a perspective of joint efficiency, we identify two natural benchmarks, depending on whether a social planner can or cannot actively intervene to change the agent's decisions. We find that there is too little innovation for high firm-specificity of innovation, and too much for low levels of firm-specificity.

Our analysis relates to several strands of economic theory. Manso (2007) examines optimal incentives for experimentation and finds that standard incentive compensation is inappropriate for environments where experimentation is important. However, his optimal contract looks very different from ours because he assumes contractibility of the innovative task. Indeed, he shows that an optimal contract needs to have a complex dynamic structure that provides high-powered incentives in the long-term, but tolerance for failure in the short-term. In an extension of our model, we show that tolerance for failure can actually be a two-edged sword: while it may encourage innovation, it may also undermine incentives for standard tasks, an effect not considered by Manso.

Rotemberg and Saloner (1994) develop a model for the benefits of a narrow business strategy, which they contrast with a broader innovation strategy. In their model, the firm manager deliberately refuses to imple-

¹Consider, for example, the following quote from management guru Gary Hamel (1999):

"If you want your company to join the pantheons of wealth-creating superstars, you have to shift the balance of effort from stewardship to entrepreneurship in your organization."

ment innovation, because accepting one employee's innovation undermines the other employees' incentives. Our model only needs a single agent, whereas their model relies on interactions among employees. In their model, the firm also has an ex-post commitment problem, which it can only solve up-front with a narrow strategy. While the mechanics of the two models are very different, both share the premise that there can be a trade-off between focusing on the core business versus innovation. In a related vein, Subramanian (2005) examines a dynamic model of a firm where employees can choose to create value for the firm or develop innovation. His main focus is the role of complementarities between the innovation and the firms' current assets. In the empirical section of this paper, he finds a negative relationship between incentive compensation and intrapreneurial activities, a prediction that is also consistent with our model.

Following the work of Grossman, Hart and Moore, several authors have tried to either explain contractual incompleteness (Segal 1999, Hart and Moore 1999), or else explore how a principal can creatively augment standard contracts to include unforeseen contingencies (Kreps, 1992, Maskin and Tirole, 1999). Our approach is fundamentally different in that we take the contractual incompleteness around innovation as given, but focus on how the agent (as opposed to the principal) responds to unforeseen innovation opportunities. The principal's response in our model is limited to adapting incentive contracts, specifically performance bonuses, to indirectly take into account the contractual incompleteness.

The work of Bernardo, Cai and Luo (2009) and Inderst and Klein (2007) examines how firms can optimally provide incentives to employees when there is both an incentive problem and a problem of truthful reporting for resource allocation. These models do not consider any multi-tasking problems. Anton and Yao (1994, 1995), Gans and Stern (2000, 2003) and Hellmann and Perotti (2007) investigate how the appropriability of innovation affects their generation and use. There is also a closely related literature that compares incentives inside firms with market-based incentives, see e.g. Anand and Galetovic (2000), de Bettignies and Chemla (2008), Gambardella and Panico (2007), and Hellmann (2007).

Our paper naturally builds on the large literature on multi-tasking, see Holmström and Milgrom (1994), Feltham and Xie (1994), and Baker (2002). The work by Holmström and Milgrom is the closest to our paper in that they also allow for tasks which are not captured by (ex ante) contractible performance measures. However, the essential difference between Holmström and Milgrom (1994) and our paper is that we incorporate ex post bargaining over the division of surplus which stems from these unmeasurable, but profitable activities such as the pursuit of innovative ideas. In other words, while unmeasurable tasks do not generate any utility for the agent in Holmström and Milgrom (1994), our framework comprises a complete parameterization of the value captured by the agent.

The remainder of the paper is structured as follows. Section 2 introduces the basic model. Section 3 examines separately the provision of incentives for the standard task (planned activity) and innovation

(unplanned activity). Section 4 examines the agent’s task choice. Section 5 derives the principal’s optimal incentive provision. Section 6 addresses the question of whether there is too much or too little innovation. In section 7, we discuss the effectiveness of several policies to encourage innovation, such as tolerance for failure, investing in employee innovation, stock-based compensation, and the allocation of intellectual property rights. Section 8 discusses empirical prediction that can be derived from our framework. Section 9 summarizes our main results and concludes.

2 The Base Model

Consider an employment relationship between a risk-neutral principal and a risk-neutral agent. The agent is financially constrained and his reservation utility is normalized to zero.² There are three dates. At date 0, the principal offers the agent an employment contract. According to this contract, the agent is charged with conducting a set of tasks, which are henceforth referred to as the *standard task*. The agent’s effort for the standard task cannot directly be observed by the principal, but she receives a performance measure which can be utilized in a court-enforceable incentive contract. At date 1, the agent may privately observe an innovation opportunity (i.e., an innovative idea), which requires further processing to bring to fruition. Due to time constraints, the agent is confronted with a specialization decision: he needs to decide on how much effort (or time) he wants to allocate to the standard task and the development of the innovation opportunity. At date 2, the principal observes the agent’s standard task performance, and makes the corresponding incentive payments as specified in the employment contract. In case the agent generated an innovation, both parties bargain over the division of the created value.

Let us first elaborate on the agent’s main responsibility to perform the standard task. By implementing effort e_S on the standard task, the agent generates revenue $V(e_S)$, which accrues entirely to the principal. To ensure interior solutions, we impose the standard requirement that revenue is increasing and concave in e_S . Also, $V(0) = 0$ and $V'(0) = \infty$. Neither the agent’s effort e_S nor the revenue $V(e_S)$ are verifiable, and thus cannot be used in a court-enforceable contract. The principal, however, receives several information about the agent’s standard effort e_S which can be summarized in a verifiable and binary signal $S \in \{0, 1\}$.³ By implementing effort e_S , the agent affects the likelihood that the favorable signal $S = 1$ will be realized. Formally, let $\Pr\{S = 1|e_S\} = \min\{e_S, 1\}$ be the conditional probability that $S = 1$. We assume that the agent’s effort costs are such that $\Pr\{S = 1|e_S^*\} = e_S^* < 1$ for the second-best effort level e_S^* in order to ensure interior solutions. Finally, to motivate standard effort, the principal ties the agent’s compensation to

²The limited liability constraint ensures that transferring the firm to the agent is not feasible.

³Demougin and Fluet (2001) show that if the involved parties are risk-neutral, any verifiable performance measures can be summarized by a binary statistic without affecting the efficiency of the contract.

his performance measure S . More specifically, in addition to a fixed compensation α , the principal pays the agent a bonus β in the event that the favorable performance measure $S = 1$ is realized.⁴

While being charged with conducting the standard task, the agent may privately observe an innovation opportunity with potential value Y . Because innovation opportunities are not known ex-ante, it is reasonable to assume that no performance measure can be defined ex-ante. Hence the two parties will have to rely on ex-post bargaining. Moreover, to capture the notion that the employee is the one who spots the opportunity, we assume that the signal about the value of the innovation idea is privately observed by the agent. We assume that the potential value of his idea $Y \in [0, \infty)$ is drawn from a distribution $\Phi(Y)$ with p.d.f. $\phi(Y)$.⁵ We further assume that innovation with high firm-specific values are less likely to be discovered by the agent, which is reflected by the assumption that $\phi(Y)$ is monotonically decreasing in Y .⁶ Once the agent has observed a particular innovation idea with potential value Y , he needs to develop this innovation by implementing non-verifiable effort e_I . The innovation process will be successful with probability $\min\{\mu e_I, 1\}$, and fails with probability $1 - \min\{\mu e_I, 1\}$, where $0 < \mu < 1$. The parameter μ captures the idea that the pursuit of innovation is more challenging than the standard task in the sense that for a given effort level, success is more likely for the standard task than for the innovation. To ensure interior solutions, we assume that $\mu e_I^* < 1$ for the second-best effort level e_I^* . Contractible performance measures which capture the agent's innovation effort e_I are not available.

We assume that the innovation value Y is observed by both the principal and the agent, but remains unobservable to third parties such as courts. Due to the absence of contractible measures, the principal cannot provide the agent with a court-enforceable incentive contract to motivate the creation of innovation. Both parties can, however, bargain over the division of the created innovation value. The relative bargaining power depends on the two parties' outside options, which in turn can depend on a variety of factors. If the idea is highly specific to the employee, and requires his specialized human capital, then the employee's bargaining position may be strong (e.g., the employee may have an outside option of implementing the innovation outside of the firm). On the other hand, if the innovation can be easily appropriated inside the firm, then the employee has a weak bargaining position. Moreover, if the idea requires the use complementary assets that are controlled by the firm, then the firm has a significantly strong bargaining position. Yet another factor is to what extent the value created by the innovation consists of private benefits that cannot be taken away from the employee.

⁴Due to the binary structure of the information system, the optimal incentive contract must be a bonus contract.

⁵We could also impose an upper bound \bar{Y} for the potential innovation value Y without affecting the quality of our results. For sufficiently low values of \bar{Y} , however, it will become clear that innovation might never be beneficial from both the principal's and agent's perspective.

⁶For instance, the exponential distribution satisfies this assumption.

To model the ex-post bargaining game between the principal and the agent, we apply the symmetric Nash bargaining solution.⁷ Clearly, the principal's share is determined by the inside value of the innovation, Y , as well as by her outside option, which we denote by $z_P Y$. The parameter z_P reflects the appropriability of the innovation, i.e., z_P captures the extent to which the principal can profitably utilize the innovation without the support of the agent. In contrast, the agent's bargaining position is determined by his outside option $z_A Y$, where the parameter z_A —as discussed above—measures the agent's private benefits from the innovation which cannot be expropriated by the principal (e.g. because of specific human capital). Using the Nash bargaining solution, the principal obtains $(1 + z_P - z_A) Y/2$ of the value of an innovation. It is convenient to define $\lambda \equiv (1 + z_P - z_A) / 2$, so that the principal gets λY , and the agent obtains $(1 - \lambda) Y$. Throughout this paper, we refer to λ as the *appropriability* of innovation.⁸ Note also that because innovation are not contractible ex-ante, any promise on the part of the principal to award the agent a higher share is never credible in this model since λ is the result of an ex-post bargaining game.⁹

For the agent's disutility of effort associated with the implementation of standard effort e_S and innovation activity e_I , let $e \equiv e_S + e_I$ denote the agent's total effort. Implementing e imposes private cost $c(e)$ which is strictly convex increasing in total effort e . The agent's private effort cost function $c(e)$ thus captures the notion that the total effort (or time) he spends on both tasks matters, but not the relative effort allocation across these tasks.

3 The Agent's Effort Choice

We begin our analysis by investigating the agent's effort choice for the standard task and innovation separately. This facilitates the derivation of the efficient bonus contract for the standard task which constitutes an important benchmark throughout this paper.

Suppose for a moment that the agent focuses exclusively on the standard task.¹⁰ Given the bonus contract (α, β) , the agent chooses his effort e_S to maximize his expected utility¹¹

$$\max_{e_S} U^S(e_S) = \alpha + \beta e_S - c(e_S). \quad (1)$$

⁷Using alternative bargaining solutions would not affect the qualitative nature of our results.

⁸Throughout the paper the concept of appropriability refers to the ability of the firm to capture the value of the employee's innovation. We therefore focus on appropriation *within* firms, rather than appropriation *across* firms that compete in the market place.

⁹However, there are indirect ways for the principal to credibly commit to a different sharing rule, e.g. through the explicit allocation of intellectual property rights (which we discuss in section 7.3), or through repeated interactions with the agent.

¹⁰We show in section 4.1 that, for a given incentive contract, the agent either exerts the standard task (so that $e_I = 0$) or pursues the innovation (so that $e_S = 0$).

¹¹Note that $c(e) = c(e_S)$ as $e_I = 0$ if the agent exerts only the standard task. The reversed observation can be made if he pursues the innovation.

It can be deduced from the first-order condition that the agent's second-best effort $e_S^*(\beta)$ is implicitly characterized by

$$\beta = c'(e_S). \quad (2)$$

The agent chooses effort $e_S^*(\beta)$ which clearly ensures equality of his marginal expected bonus and marginal effort costs. From (2) we note that stronger incentives (higher β) lead to more effort.

Anticipating the agent's effort choice $e_S^*(\beta)$, the principal's problem is to find the bonus contract which maximizes her expected net payoff from the standard task:

$$\Pi^S(\alpha, \beta) = V(e_S^*(\beta)) - \alpha - \beta e_S^*(\beta). \quad (3)$$

Moreover, the optimal bonus contract needs to ensure the agent's participation, and to account for his liability limit, which requires that $\alpha \geq 0$ and $\alpha + \beta \geq 0$. The next lemma characterizes the efficient contract which maximizes the principal's expected net payoff from the standard task, in addition to its effect on the agent's expected utility.

Lemma 1

- (i) *The agent's expected utility from the standard task, $U^S(\alpha, \beta)$, is increasing in α and β .*
- (ii) *The principal's expected profit from the standard task, $\Pi^S(\alpha, \beta)$, is decreasing in α , and concave in β . It reaches its maximum at β_0 which is implicitly characterized by*

$$\frac{d\Pi^S(e_S^*(\beta), \beta)}{d\beta} = 0. \quad (4)$$

- (iii) *The optimal choice of α is $\alpha_0 = 0$.*

Proof All proofs are given in the Appendix.

Lemma 1 provides three important observations. First, the agent's expected utility is always increasing in the bonus compensation β . The fact that the agent's expected utility is strictly increasing in β will be a fundamental driver for his specialization decision.

The second important observation is that the principal's expected profit is non-monotonic in β , with an interior maximum at β_0 . For low values of β , the marginal benefit of paying the agent more outweighs the cost, but for high β , the opposite applies. The profit-maximizing value β_0 is a useful benchmark of the intensity of incentives, as it measures the principal's preferred level of incentives in the absence of any multi-tasking concerns. Hence we will refer to β_0 as the *unconstrained* incentive bonus.

Third, we note that any fixed compensation α only transfers rents from the principal to the agent. Since the agent's participation constraint is always satisfied for any α and β satisfying the conditions $\alpha \geq 0$ and $\alpha + \beta \geq 0$, it is immediate that the principal always chooses $\alpha_0 = 0$.

After identifying the agent's optimal standard effort and the associated bonus contract, we now turn to the agent's incentives to bring innovative ideas to fruition. Suppose the agent focuses exclusively on the development of the innovation. Given the appropriability of his idea (parameterized by λ), he chooses his innovation effort e_I to maximize his expected utility

$$\max_{e_I} U^I(e_I) = (1 - \lambda)Y\mu e_I - c(e_I). \quad (5)$$

The first-order condition reveals that the agent's second-best innovation effort $e_I^*(\lambda, Y)$ is implicitly characterized by

$$(1 - \lambda)Y\mu = c'(e_I). \quad (6)$$

As for the standard task, the agent chooses his innovation effort $e_I^*(\lambda, Y)$ such that the marginal compensation equals his marginal cost of effort. The principal, however, cannot directly control the agent's innovation effort as it is implicitly determined by the appropriability of his idea, λ , and the innovation value Y . It can be inferred from (6) that identifying an innovation opportunity with a lower appropriability (λ) or detecting a more valuable innovation (Y) motivates the agent to implement more effort.

For the subsequent analysis, we now investigate the relationship between the appropriability of innovation λ and the principal's expected net payoff from innovation. We start by briefly discussing the extreme cases: the innovation can entirely be appropriated by the principal ($\lambda = 1$) or the agent ($\lambda = 0$). It is clear that, if the principal can capture the entire surplus from the innovation ($\lambda = 1$), the agent refuses to develop any innovation ($e_I^*(1, Y) = 0$) due to the lack of prospective compensations. This in turn implies a zero payoff for the principal from innovation ($\Pi^I(e_I^*(1, Y), 1, Y) = 0$). If the agent can capture the entire surplus ($\lambda = 0$), he is motivated to implement the first-best effort level $e_I^*(0, Y)$. However, the principal does not benefit from the innovation ($\Pi^I(e_I^*(0, Y), 0, Y) = 0$). The previous observations imply that the principal can only benefit from the innovation if $\lambda \in (0, 1)$ such that both involved parties receive a strictly positive share on the corresponding surplus. The next lemma emphasizes the relationship between the appropriability of the innovation λ and the principal's expected net payoff from innovation.

Lemma 2 *There exists a threshold innovation appropriability level $\lambda^* \in (0, 1)$ such that the principal's expected net payoff from innovation $\Pi^I(e_I^*(\lambda, Y), \lambda, Y)$ is increasing in λ for $\lambda \leq \lambda^*$, and decreasing in λ for $\lambda > \lambda^*$.*

Lemma 2 exposes one of the fundamental forces in our framework. As discussed above, a more appropriate innovation (higher λ) diminishes the agent's incentives to develop an innovation opportunity. However, this concurrently enhances the principal's share on the created surplus. Thus, an increase in λ has two opposite effects on the principal's expected payoff from innovation. We thus encounter a familiar result from the incentive literature, namely that for low level of incentives ($\lambda > \lambda^*$) both the principal and the agent would benefit from stronger incentives. For higher levels of incentives ($\lambda \leq \lambda^*$), however, only the agent benefits from stronger incentives, while the principal's cost of giving up rents exceeds the benefits of stronger incentives.

4 The Agent's Specialization Decision

4.1 The Agent's Task Choice

The agent, while being charged with conducting the standard task, may privately observe a potential innovation with the firm-specific value Y . He is thus faced with a private specialization decision which cannot directly be controlled by the principal. Put differently, the agent needs to decide on how much effort he wants to allocate to the standard task versus the development of the innovation respectively.

The next lemma emphasizes the relationship between the innovation value Y , which is privately observed by the agent, and his optimal effort allocation.

Lemma 3 *There exists a threshold innovation value Y^A such that for $Y < Y^A$, the agent only conducts the standard task (i.e., $e_S^* > e_I^* = 0$), and for $Y > Y^A$ the agent only explores the innovation (i.e., $e_I^* > e_S^* = 0$).*

According to Lemma 3, the agent either conducts the standard task or pursues the innovation.¹² The threshold Y^A , which captures the agent's specialization decision, maximizes his expected utility across both tasks. Using the agent's optimal effort levels $e_I^*(\beta)$ and $e_I^*(\lambda, Y)$ as defined by (2) and (6), his expected utility can be stated as

$$U(\cdot) = \int_0^{Y^A} [\alpha + \beta e_S^* - c(e_S^*)] d\Phi(Y) + \int_{Y^A}^{\infty} [\alpha + (1 - \lambda)Y\mu e_I^* - c(e_I^*)] d\Phi(Y). \quad (7)$$

¹²Technically, this lemma is due to the fact that both tasks have a constant marginal expected benefit, and that the disutility of effort is convex in total effort. This is extremely useful, because it generates a unique threshold between the standard task and the innovation. However, we demonstrate in section 7.4 that the model can be extended by allowing for decreasing marginal benefits of effort, so that the agent divides his time between the two activities. We show that our main results—in particular the optimal adjustment of effort incentives—also hold for the more general framework.

Note that we have suppressed the arguments of e_S^* and e_I^* . For parsimony purposes, we stick to the same practice in the remainder of this paper. It follows directly from the first-order condition with respect to Y^A that the agent's threshold innovation value $Y^A(\beta, \lambda)$ is implicitly characterized by

$$\beta e_S^* - c(e_S^*) = (1 - \lambda)Y^A \mu e_I^* - c(e_I^*), \quad (8)$$

which is equivalent to $U^S(e_S^*, \beta) = U^I(e_I^*, \lambda, Y^A)$. We can immediately infer from (8) that the agent's choice between exploitation (standard task) and exploration (innovation) is determined by his respective rents for conducting these tasks. At first glance, the agent's specialization decision is beyond the principal's control. However, recall from Lemma 1 that the provision of incentives for the standard task directly affects the agent's associated rent extraction. Thus, we can infer from (8) that the principal has some latitude to indirectly control the agent's specialization decision by adjusting standard incentives (β). We now clarify the relationship between effort incentives for the standard task and the agent's preference for exploring potential innovation.

Proposition 1

- (i) *If the principal provides the agent with stronger incentives for the standard task, the agent innovates less frequently. Formally, the agent's threshold innovation value $Y^A(\beta, \lambda)$ is increasing in β .*
- (ii) *The greater the appropriability of the innovation λ , the less the agent innovates, i.e., $Y^A(\beta, \lambda)$ is increasing in λ .*

Proposition 1 establishes the fundamental mechanism by which the principal can indirectly control the agent's specialization decision. Accordingly, intensifying effort incentives for the standard task allows the principal to make the exploration of innovation less attractive to the agent. The reason is as follows. From Lemma 1 we know that the principal can affect the agent's rent from the standard task by adjusting the bonus payment β . This in turn determines the value of conducting the standard task relative to developing the innovation, and hence the agent's preferences, as reflected by the threshold Y^A . Generally speaking, it requires sufficiently more 'promising' ideas for the agent to be innovative. Note, however, that adjusting standard incentives with the aim of manipulating the agent's specialization decision comes at a cost for the principal: it impairs the profitability of the standard task.

The second part of Proposition 1 provides an important and intuitive comparative static result for the agent's innovation preference. The greater the appropriability of the innovation (λ), the weaker his preference for innovation, as reflected by a higher threshold value Y^A .

4.2 The Principal's Preference

In order to understand the principal's optimal compensation choices, it is useful to first understand the principal's preference for innovation, and to compare her preferences to the actual choices made by the agent. For now it suffices to examine the principal's innovation preference at the unconstrained level β_0 ; in Section 6 we will generalize it by accounting for the optimal level of β . We define $Y^P(\beta_0, \lambda)$ as the threshold value of Y , such that the principal prefers the agent to perform the standard task for all $Y \leq Y^P$, and the innovation otherwise.¹³

Given the agent's second-best effort levels e_S^* and e_I^* as characterized in Section 3, the principal's expected profit is

$$\Pi(\cdot) = \int_0^{Y^P} [V(e_S^*) - \alpha - \beta e_S^*] d\Phi(Y) + \int_{Y^P}^{\infty} [\lambda Y \mu e_I^* - \alpha] d\Phi(Y). \quad (9)$$

Note that by construction Y^P does not depend on β . As a consequence, $\Pi(\cdot)$ is always maximized at β_0 . The optimal choice of $Y^P(\beta_0, \lambda)$ is then implicitly characterized by

$$V(e_S^*) - \beta_0 e_S^* = \lambda Y^P \mu e_I^*, \quad (10)$$

which is equivalent to $\Pi^S(e_S^*, \beta_0) = \Pi^I(e_I^*, \lambda, Y^P)$. The principal's preference is simply dictated by her respective payoffs from either tasks. As for the agent, preferences for a specific task choice are not affected by the fixed wage α . The next lemma focuses on the relationship between the appropriability of the innovation λ and the principal's innovation preferences Y^P .

Lemma 4 *The principal's preferred threshold innovation value $Y^P(\beta_0, \lambda)$ is decreasing in λ for $\lambda \leq \lambda^*$, and increasing otherwise.*

Lemma 4 shows that the principal's innovation preferences, reflected by Y^P , are non-monotonic in the innovation appropriability λ . From Lemma 2 we know that the principal's profits from innovation first increases and then decreases in λ . This insight drives Lemma 4. For low λ (i.e., $\lambda \leq \lambda^*$), the principal benefits from a higher share on surplus, making innovation relatively more attractive. However, this result is reversed for higher levels of the appropriability of innovation (i.e., $\lambda > \lambda^*$).

We are now well equipped to contrast the principal's and agent's preferences for exploration versus exploitation, holding incentives constant at the unconstrained level β_0 .

¹³There also exists an alternative interpretation of the threshold value $Y^P(\beta_0, \lambda)$. Consider a model in which the principal, not the agent, makes the task decision. This could be either because the principal observed the signal Y , or because the agent requires some essential resources in order to pursue the innovation. In such a model, it is easy to see that the principal's optimal choices are always β_0 for the incentive compensation, and Y^P for the task allocation.




Figure 1: Task Choice: Principal's vs. Agent's Preferences

Proposition 2 *There exists a threshold level of the innovation appropriability λ^{**} , with $\lambda^{**} < \lambda^*$, such that the agent prefers more innovation than the principal if $\lambda \leq \lambda^{**}$ (i.e., $Y^A(\beta_0) \leq Y^P(\beta_0)$). Otherwise, the agent prefers less innovation (i.e., $Y^A(\beta_0) > Y^P(\beta_0)$).*

Figure 1 illustrates the principal's and agent's preferences, Y^P and Y^A , for different values of the innovation appropriability λ .¹⁴ Consider first the intersection of both curves. At λ^{**} , both the principal's and the agent's preferences towards exploration and exploitation are perfectly aligned. Put differently, we would observe the same effort allocation for all potential innovation values, irrespective of whether or not the agent's specialization decision can be dictated by the principal. The specific value λ^{**} of the innovation appropriability, where preferences are perfectly congruent, will play an important role in our subsequent analysis. For all other values of λ we can infer from Figure 1 that preferences differ. More specifically, if the agent is in a relatively better bargaining position than the principal because of a low appropriability level ($\lambda < \lambda^{**}$), he prefers to focus more frequently on innovation at the expense of the standard task. Yet, the agent has stronger preferences towards exploitation (the standard task) as compared to the principal if the innovation is highly appropriable ($\lambda > \lambda^{**}$). Figure 1 thus illustrates the conflict of interests between both the principal and the agent, which constitutes the foundation for our subsequent analysis.

¹⁴We show in the Appendix that, if the entire surplus goes to either the principal ($\lambda \rightarrow 1$) or the agent ($\lambda \rightarrow 0$), the principal's preferred threshold $Y^P(\cdot)$ goes to infinity. Because innovation lead to zero payoffs for the principal in both cases, she wants the agent to exert only the standard task, irrespectively of potential innovation values. The same observation applies to the agent's preferred threshold $Y^A(\cdot)$ in case the principal receives the entire surplus ($\lambda \rightarrow 1$).

5 Firm's Optimal Incentive Provision

5.1 Balancing Incentives and Task Choice

We can now investigate the principal's choice of contracts which aim at balancing incentives for the standard task and the efficiency of the agent's task choice. Because the agent has private information, the principal might benefit from offering a menu of contracts that allows the agent to self-select, and thereby to reveal some of his private information. Even though the variable of the information asymmetry, Y , is continuous, the analysis above already shows that we can classify agents into two types: the first type (for lower values of Y) wants to exclusively focus on the standard task, the other type (for higher values of Y) wants to pursue his innovative idea. A useful aspect of this model that greatly simplifies the analysis is that once an agent has chosen to be one of these two types, the exact value of Y turns out to be unimportant for the purpose of designing the optimal contract. Intuitively, if the agent chooses to focus on the standard task, the exact value of Y is irrelevant since the innovation is not pursued anyway. If, on the other hand, the agent does pursue the innovation, the exact value of Y obviously matters for the agent's effort decision. However, as we will see below, the principal's choice of contractual instruments is sufficiently limited so that she would always want to offer the same contract to all agents that pursue their innovative ideas. We can therefore limit the analysis to the case where the principal offers a menu of two contracts. To do so, let $\{\alpha^*, \beta^*\}$ denote the contract tailored to the agent who wants to exert the standard task, and $\{\alpha^{**}, \beta^{**}\}$ the contract for the innovative agent, respectively. For future reference, we denote the optimal menu of contracts by $\Omega^* = \{\{\alpha^*, \beta^*\}, \{\alpha^{**}, \beta^{**}\}\}$.

We immediately note two important properties of the optimal menu of contracts. First, we can deduce from Lemma 1 that the optimal contract for the non-innovative agent comprises $\alpha^* = 0$. Moreover, since the innovative agent would disregard the standard task, his optimal contract is characterized by $\beta^{**} = 0$. Thus, it remains to derive the optimal bonus β^* for the non-innovative agent, and the optimal fixed wage α^{**} for the innovative agent. To characterize the optimal menu of contracts, we proceed in two steps. First, we derive the remaining contract components for both types of agents separately. Second, we characterize the optimal menu of contracts for different values of λ which accounts for the principal's preference for innovation.

Consider the contract $\{\alpha^* = 0, \beta^*\}$ a non-innovative agent would choose. This contract maximizes the principal's expected profit, taking the agent's second-best effort levels e_S^* and e_I^* , and his threshold innovation value $Y^A(\beta, \lambda)$ into account. The principal's expected profit is thus given by

$$\Pi(e_S^*, e_I^*, \beta, Y^A) = \int_0^{Y^A} [V(e_S^*) - \beta e_S^*] d\Phi(Y) + \int_{Y^A}^{\infty} [\lambda Y \mu e_I^*] d\Phi(Y). \quad (11)$$

Incentives for the standard task affect the principal's expected profit in two ways. First, the bonus β indirectly determines the agent's standard effort e_S^* , which has a significant impact on the profitability of exploitation. Second, the bonus β influences the agent's specialization decision, as reflected by the threshold $Y^A(\beta, \lambda)$. From the first-order condition, the optimal bonus $\beta^*(\lambda)$ is implicitly characterized by

$$\frac{\partial \Pi^S}{\partial \beta} = [\Pi^I - \Pi^S] \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\beta}. \quad (12)$$

The central insight from this equation is as follows. Suppose that $\Pi^I > \Pi^S$. This means that at $Y^A(\beta, \lambda)$ (i.e., the agent's point of indifference between the two tasks), the principal prefers innovation over the standard task. Intuitively, this corresponds to the case where $\lambda > \lambda^{**}$ in Proposition 2 (see also Figure 1). To act upon her preference for innovation, the principal now adjusts the agent's compensation. The above first-order condition indicates that for $\Pi^I > \Pi^S$, the optimal compensation occurs at a point where $\partial \Pi^S / \partial \beta$ is positive. Given the concavity of Π^S (see Lemma 1), this means that $\beta^*(\lambda) < \beta_0$, i.e., the principal lowers the incentive bonus below the unconstrained level. The intuition is that lowering standard incentives reduces the agent's rent from the standard task, and therefore encourages him to pursue more innovation. For $\Pi^I < \Pi^S$, the opposite applies. The following proposition formalizes the intuition.

Proposition 3

- (i) For $\lambda < \lambda^{**}$, the optimal incentive bonus lies above the unconstrained level (i.e., $\beta^*(\lambda) > \beta_0$), and for $\lambda > \lambda^{**}$, the optimal incentive bonus lies below the unconstrained level (i.e., $\beta^*(\lambda) < \beta_0$).
- (ii) For $\lambda < \lambda^*$, the optimal incentive bonus $\beta^*(\lambda)$ is decreasing in λ (i.e., $d\beta^*(\lambda)/d\lambda < 0$), and for $\lambda > \lambda^*$, the optimal incentive bonus $\beta^*(\lambda)$ is increasing in λ (i.e., $d\beta^*(\lambda)/d\lambda > 0$).

Next, consider the contract $\{\alpha^{**}, \beta^{**} = 0\}$ that an innovative agent would choose. Although the fixed payment α^{**} does not affect the agent's effort choice, it can be adjusted to manipulate his specialization decision. This can be shown as follows: by accounting for the menu of contracts Ω^* , the agent's preferences—reflected by Y^A —are implicitly characterized by $U^S(\alpha^* = 0, \beta^*(\lambda)) = U^I(\alpha^{**}, \beta^{**} = 0, Y^A)$. Implicitly differentiating Y^A with respect to α^{**} yields

$$\frac{dY^A}{d\alpha^{**}} = -\frac{1}{(1-\lambda)\mu e_I^*} < 0. \quad (13)$$

Clearly, providing the agent with a higher fixed wage α^{**} as part of a menu of contracts encourages him to pursue more innovation. The optimal fixed payment α^{**} maximizes the principal's expected profit given by

$$\Pi(\cdot) = \int_0^{Y^A} [V(e_S^*) - \beta^* e_S^*] d\Phi(Y) + \int_{Y^A}^{\infty} [\lambda Y \mu e_I^* - \alpha^{**}] d\Phi(Y). \quad (14)$$




Figure 2: Innovation Appropriability and the Optimal Menu of Contracts

The optimal fixed wage $\alpha^{**}(\lambda)$ is implicitly characterized by the following first-order condition:

$$[\Pi^S - \Pi^I] \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\alpha^{**}} = 1. \quad (15)$$

This equation offers an important insight into the optimal design of incentive contracts. Clearly, (15) can only be satisfied if $\Pi^I > \Pi^S$. In this case, the principal strictly prefers innovation over the standard task, and is thus willing to offer the agent a contract comprising a strictly positive base wage $\alpha^{**}(\lambda)$. Otherwise, innovation are of relatively less importance for the principal such that she refrains from offering the agent a strictly positive fixed wage. To get a more complete picture of the optimal base wage offered to innovative agents, the next proposition elaborates on the relationship between $\alpha^{**}(\lambda)$ and the innovation appropriability λ .

Proposition 4

- (i) For $\lambda \leq \lambda^{**}$, the optimal fixed wage $\alpha^{**}(\lambda)$ is identical to the unconstrained level (i.e., $\alpha^{**}(\lambda) = \alpha_0 = 0$), and for $\lambda > \lambda^{**}$, the optimal fixed wage $\alpha^{**}(\lambda)$ lies above the unconstrained level (i.e., $\alpha^{**}(\lambda) > \alpha_0 = 0$).
- (ii) For $\lambda^{**} < \lambda < \lambda^*$, the optimal fixed wage $\alpha^{**}(\lambda)$ is increasing in λ (i.e., $d\alpha^{**}(\lambda)/d\lambda > 0$), and for $\lambda > \lambda^*$, the optimal fixed wage $\alpha^{**}(\lambda)$ is decreasing in λ (i.e., $d\alpha^{**}(\lambda)/d\lambda < 0$).

Figure 2 illustrates the fundamental insights from Propositions 3 and 4. Consider first the optimal incentive bonus $\beta^*(\lambda)$ offered to non-innovative agents. For low levels of λ (i.e., $\lambda \leq \lambda^*$), $\beta^*(\lambda)$ is decreasing in λ , reflecting the above-mentioned desire of the principal to encourage innovation by reducing the agent's rent from the standard task. This desire reaches its peak at λ^* , where we know from Lemma 2 that the principal obtains maximal profits from innovation. Observe further that at λ^{**} , the optimal bonus $\beta^*(\lambda)$ becomes

identical to the unconstrained level β_0 . Indeed, we know from Proposition 2 (see also Figure 1) that at λ^{**} , the principal and the agent have congruent preferences towards the task choice, reflected by identical innovation thresholds Y^P and Y^A . Finally, for $\lambda < \lambda^{**}$, the principal wants to actively discourage innovation and refocus the agent on the standard task. For this she raises the incentive bonus above the unconstrained level.

Since the optimal fixed payment $\alpha^{**}(\lambda)$ for innovative agents does not interfere with effort incentives, its exclusive objective is to manipulate the agent's specialization decision. We know from Proposition 2 that the agent is too innovative from the principal's perspective for $\lambda < \lambda^{**}$. To discourage innovation, the optimal menu of contracts comprises higher-powered incentives (i.e., $\beta^*(\lambda) > \beta_0$), and a zero fixed payment (i.e., $\alpha^* = 0$). For $\lambda > \lambda^{**}$, however, the principal wants the agent to pursue more innovation, which can be achieved by lowering standard incentives (as discussed above), and offering a strictly positive base wage for innovative agents. Because the principal's benefit from innovation reaches its peak when $\lambda = \lambda^*$, she minimizes standard incentives and maximizes the base wage for innovative agents.

From the previous discussion it becomes clear that the optimal menu of contracts $\Omega^*(\lambda)$ crucially hinges on how much of the surplus from innovation can be captured by the principal, which is dictated by the innovation appropriability level λ . Because a contractual choice on the side of the agent encourages more innovation, offering a menu of contracts can only be optimal whenever the agent is not innovative enough, which is clearly the case for $\lambda > \lambda^{**}$. The next proposition characterizes the optimal menu of contracts for all values of λ based upon our previous observations.

Proposition 5 *To discourage innovation for $\lambda \leq \lambda^{**}$, the principal offers the agent the contract $\{\alpha^* = 0, \beta^*(\lambda)\}$. In contrast, to encourage more innovation for $\lambda > \lambda^{**}$, the principal offers the agent a menu of contracts $\Omega^* = \{\{\alpha^* = 0, \beta^*(\lambda)\}, \{\alpha^{**}(\lambda), \beta^{**} = 0\}\}$.*

Proposition 5 in conjunction with Proposition 3 addresses an important economic puzzle around low-powered incentives. Many formal and informal empirical studies (e.g. Williamson, 1975, Parent, 2002) have commented on the prevalence of low-powered incentives in many (but not all) organizations. If one only focuses on the formal job description of an employee, it might seem that higher-powered incentives would generate more efficient outcomes. In fact, this is clearly apparent in the current model: for $\lambda > \lambda^{**}$ we observe low-powered incentives (i.e., $\beta^*(\lambda) < \beta_0$); by focusing only on the standard task, it would appear as if there are opportunities for increasing (Pareto-) efficiency by enhancing incentives to β_0 . However, a central insight of this paper is that focusing on the planned and measurable part of the job description is too narrow a focus, because employees may also have opportunities to pursue unplanned innovation. Proposition 3 thus shows under what conditions the presence of these unplanned innovation can explain the prevalence

of these low-powered incentives. The condition is simply that most of the surplus from innovation can be appropriated by firms. Empirically, this is likely to be true in many but not all circumstances.

5.2 Equilibrium Levels of Innovation

In the analysis of section 4.2, we evaluated preferences at the unconstrained performance bonus β_0 . We now complete the analysis by explaining how innovation preferences behave at the optimal menu of contracts $\Omega^*(\lambda) \in \{\{0, \beta^*(\lambda)\}, \{\alpha^{**}(\lambda), 0\}\}$. Proposition 3 explains the optimal choice of incentives, and Proposition 1 shows how incentives affect the agent's specialization decision. We now characterize how the agent's equilibrium choice $Y^A(\Omega^*(\lambda), \lambda)$ depends on λ . Moreover, we examine whether in equilibrium there remains any discrepancy between the principal's and agent's innovation preferences, i.e., we compare $Y^P(\Omega^*(\lambda), \lambda)$ with $Y^A(\Omega^*(\lambda), \lambda)$.

Proposition 6

- (i) For $\lambda < \lambda^{**}$, the optimal menu of contracts $\Omega^*(\lambda)$ discourages innovation by raising the agent's innovation threshold Y^A (i.e., $Y^A(\Omega^*(\lambda), \lambda) > Y^A(\beta_0, \lambda)$). For $\lambda > \lambda^{**}$, the optimal menu of contracts encourages innovation by lowering the agent's innovation threshold Y^A (i.e., $Y^A(\Omega^*(\lambda), \lambda) < Y^A(\beta_0, \lambda)$).
- (ii) For $\lambda < \lambda^{**}$, the agent chooses too many innovation relative to the principal's preference (i.e., $Y^A(\Omega^*(\lambda), \lambda) < Y^P(\Omega^*(\lambda), \lambda)$). For $\lambda > \lambda^{**}$, the agent chooses too few innovation relative to the principal's preference (i.e., $Y^A(\Omega^*(\lambda), \lambda) > Y^P(\Omega^*(\lambda), \lambda)$).

Figure 3 illustrates the results from Proposition 6. The agent's equilibrium response to incentives is very intuitive. For $\lambda > \lambda^{**}$, the optimal menu of contracts generates fewer rents for the standard task, which encourages the agent to focus relatively more on innovation, thus resulting in a lower threshold Y^A . The opposite is true for $\lambda < \lambda^{**}$.

The principal's preferred innovation threshold Y^P also depends on the level of incentives.¹⁵ Interestingly, we find that the principal's innovation threshold under the contract $\{0, \beta^*\}$, $Y^P(\beta^*, \lambda)$, always lies below $Y^P(\beta_0, \lambda)$, the threshold based on the unconstrained performance bonus (except for $\lambda = \lambda^{**}$ where $\beta^* = \beta_0$). The reason is that β_0 maximizes the principal's net value from the standard task. Any other choice of β , irrespective of whether it is higher or lower than β_0 , generates lower profits from the standard task, which in turn raises the principal's preference for innovation (implying a lower innovation threshold Y^P). However, the effect of the alternative contract $\{\alpha^{**}, 0\}$ on the principal's preference Y^P is ambiguous.

¹⁵Refer to Proof of Proposition 6 in the Appendix for the analytical results discussed below.

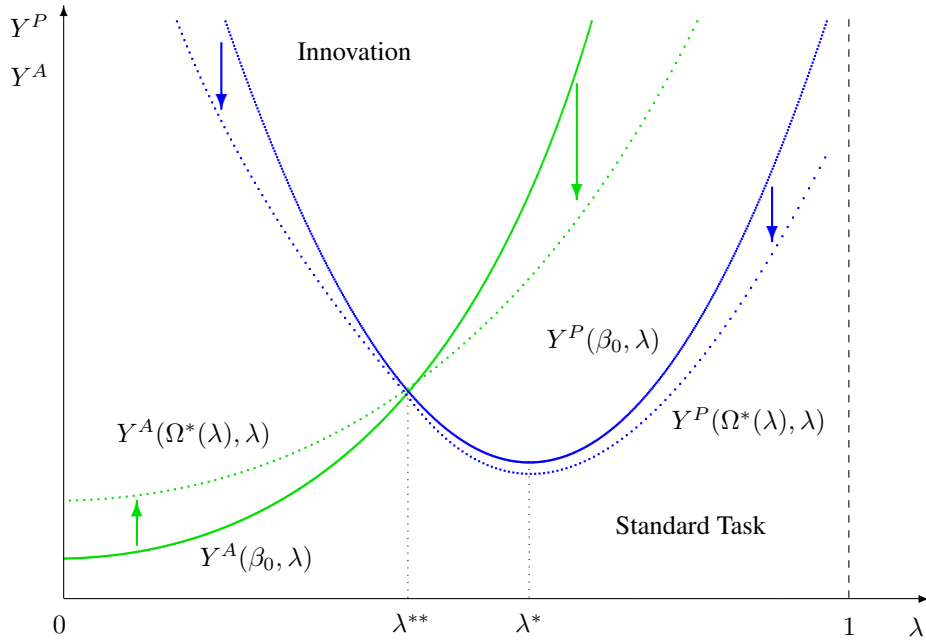


Figure 3: Menu of Contracts and the Adjustment of Preferences

First, this contract comprises a fixed wage which lies above the unconstrained level for $\lambda^{**} < \lambda < 1$, which in turn compromises the efficiency of innovation. Thus, one could expect Y^P to increase. On the other hand, the principal sets $\beta^{**} = 0$, rendering the standard task unprofitable. As can be inferred from our previous discussion, this would clearly lower Y^P . The net effect on Y^P , however, depends in particular on the appropriability level λ . Finally, comparing the agent's equilibrium innovation choice $Y^A(\Omega^*(\lambda), \lambda)$ with the principal's equilibrium preference $Y^P(\Omega^*(\lambda), \lambda)$, we note that except at λ^{**} , there always exists a discrepancy. Despite providing optimal incentives, for high appropriability levels of innovation the principal still finds that the agent pursues too few innovation (i.e., for $\lambda > \lambda^{**}$, we have $Y^A(\Omega^*(\lambda), \lambda) > Y^P(\Omega^*(\lambda), \lambda)$).¹⁶ The interval $(Y^P(\Omega^*(\lambda), \lambda), Y^A(\Omega^*(\lambda), \lambda))$ has an interesting interpretation: every time a perceived innovation value Y falls into this interval, the agent privately forgoes the innovation opportunity, even though the principal would have liked to see him pursue. For low levels of appropriability, we encounter the opposite situation. The agent pursues more innovation than the principal wants him to. Every time the potential innovation value Y falls into the range $(Y^A(\Omega^*(\lambda), \lambda), Y^P(\Omega^*(\lambda), \lambda))$, the agent pursues the innovation, even though the principal would have preferred the agent to stick to the standard task.

As an aside, one may wonder whether the trade-off between standard tasks and innovation could not be solved by hiring different groups of employees, one that focuses solely on standard tasks, the other on

¹⁶It is interesting to note that while standard incentives could be tailored to ensure perfect alignment of preferences, it is not optimal for the principal to do so. This is because adjusting the menu of contracts is costly to the principal: raising the fixed wage α imposes additional costs without affecting the effort intensities, and changing the incentive bonus β lowers the efficiency of the standard task.

innovation. The above analysis clarifies that this is never optimal in this setting, for two reasons. First, a purely innovative agent would pursue every innovation as long as $Y > 0$ and $\lambda < 1$, and thus, would always be too innovative from the perspective of the principal. Second, agents ‘exclusively’ employed for standard tasks might still come up with valuable innovative ideas, and will thus still be exposed to the trade-off between sticking to the standard task versus pursuing the innovation.

6 Evaluating The Agent’s Specialization Decision

In this section, we characterize several benchmarks which allow us to shed more light on the efficiency of the agent’s private task choice from a social perspective. There are a variety of benchmarks against which one may evaluate the agent’s task choice between innovation and the standard task, as characterized by $Y^A(\Omega^*(\lambda), \lambda)$. We have already encountered one such benchmark, namely the principal’s preference for innovation, as characterized by $Y^P(\Omega^*(\lambda), \lambda)$. We now explore two additional benchmarks that are based on the joint utility of the principal and the agent, using the sum of utilities as a standard social welfare criterion. Our first benchmark assumes an *active* social planner who can implement a first-best contract. Our second benchmark assumes a *passive* social planner who cannot influence the equilibrium outcome, but evaluates whether there is too much or too little innovation from a social perspective.

Our first social benchmark is the first-best outcome, which is defined as the outcome chosen by a social planner who can control all of the principal’s and the agent’s choices. In our model, the social planner would need to be able to control the agent’s task choice—the corresponding threshold is denoted by \hat{Y} —and his effort choices e_S and e_I in order to maximize the total surplus, given as

$$S(e_S, e_I, \hat{Y}) = \int_0^{\hat{Y}} [V(e_S) - c(e_S)] d\Phi(Y) + \int_{\hat{Y}}^{\infty} [Y\mu e_I - c(e_I)] d\Phi(Y). \quad (16)$$

The first-best effort choices e_S^{FB} and e_I^{FB} are implicitly characterized by the first-order conditions $V'(e_S) = c'(e_S)$ and $Y\mu = c'(e_I)$. Moreover, the first-best innovation threshold \hat{Y}^{FB} satisfies

$$V(e_S^{FB}) - c(e_S^{FB}) = \hat{Y}^{FB} \mu e_I^{FB} - c(e_I^{FB}). \quad (17)$$

Notice that \hat{Y}^{FB} is not affected by the innovation appropriability λ . This can be observed because λ determines only the division of surplus from the innovation, and not the total surplus.




Figure 4: Social Efficiency of the Agent's Task Choice

Our second social benchmark takes the equilibrium outcome— $\Omega^*(\lambda)$, $Y^A(\Omega^*(\lambda))$, e_S^* , and e_I^* —as given, and considers the preferences of a (powerless) social planner that cares about the total surplus. One can easily verify that the second-best social benchmark $\hat{Y}^{SB}(\Omega^*(\lambda), \lambda)$ satisfies

$$V(e_S^*) - c(e_S^*) = \hat{Y}^{SB} \mu e_I^* - c(e_I^*). \quad (18)$$

We now state the main proposition for this section.

Proposition 7 Consider the first-best social benchmark \hat{Y}^{FB} and the second-best benchmark $\hat{Y}^{SB}(\Omega^*(\lambda), \lambda)$.

- (i) Relative to \hat{Y}^{FB} , there is too much innovation for sufficiently low λ (i.e., $Y^A(\Omega^*(\lambda), \lambda) < \hat{Y}^{FB}$), and too little innovation for sufficiently high λ (i.e., $Y^A(\Omega^*(\lambda), \lambda) > \hat{Y}^{FB}$).
- (ii) Relative to $\hat{Y}^{SB}(\Omega^*(\lambda), \lambda)$, there is too much innovation for $\lambda < \lambda^{**}$ (i.e., $Y^A(\Omega^*(\lambda), \lambda) < \hat{Y}^{SB}(\Omega^*(\lambda), \lambda)$), and too little innovation for $\lambda > \lambda^{**}$ (i.e., $Y^A(\Omega^*(\lambda), \lambda) > \hat{Y}^{SB}(\Omega^*(\lambda), \lambda)$).

Figure 4 illustrates this result. Consider the first-best social benchmark \hat{Y}^{FB} . Clearly, \hat{Y}^{FB} always lies above $Y^A(\Omega^*(\lambda), \lambda)$ at $\lambda = 0$. We briefly sketch the main intuition. At $\lambda = 0$, the agent always provides efficient effort e_I^{FB} , because he is the full residual claimant of the innovation value. For \hat{Y}^{FB} , the agent would also provide efficient effort e_S^{FB} , but this is not true for $Y^A(\Omega^*(0), 0)$, since the agent is not the full residual claimant of the value of the standard task. This could only be achieved if the principal 'sells' the

firm to the agent by setting $\alpha < 0$.¹⁷ This, however, would violate the agent's liability limit, and is thus not feasible in this environment.

Finally, the second-best social benchmark $\widehat{Y}^{SB}(\Omega^*(\lambda), \lambda)$ is essentially a weighted average of the principal's and the agent's preferences. It therefore behaves in a straightforward manner: whenever the principal wants more or less innovation than the agent, then the same is true for the social planner.

7 Extensions

7.1 Tolerance of Failure

In a recent paper, Manso (2007) investigates how firms can encourage more innovation in a dynamic environment. He shows that the optimal dynamic contract exhibits tolerance for failure in the short run, but high-powered incentives in the long run. While Manso focuses on the case where the agent exerts only innovative activities, we consider how incentives for standard tasks interfere with incentives to innovate in a multi-tasking setting. It is therefore not only logical, but also of high practical relevance to examine whether adopting a tolerance policy can be optimal for firms when exploitation and exploration are two competing objectives.

We will think of the tolerance policy as a choice of organizational culture, and assume that the organization either tolerates failure or not. We also use a very simple specification, where tolerance means that the firm can soften the impact of an employee's failure to either produce an innovation, or to generate a positive signal for the standard task.¹⁸ Specifically, we assume that tolerance of failure means that the agent obtains $\varepsilon > 0$ instead of 0 in case of failure. For simplicity, we further assume that the provision of ε is costless to the principal. What we have in mind is a corporate culture, where unsuccessful agents are protected and e.g. not humiliated in front of coworkers, which in turn provides them with a higher utility level as compared to the most severe punishment strategy.¹⁹ By allowing for tolerance, the agent's first-order conditions (2) and (6) change to

$$\beta - \varepsilon = c'(e_S), \quad (19)$$

$$(1 - \lambda)Y\mu - \mu\varepsilon = c'(e_I). \quad (20)$$

¹⁷Alternatively, the principal could set $\alpha = 0$ and adjust β such that the agent obtains the entire surplus from the standard tasks. In this case, however, the principal's net payoff from the standard task is zero, implying that this solution cannot be optimal.

¹⁸An alternative modeling strategy would be to assume that the firm has fine-grained control over tolerance for failure, so that it could distinguish between the failures of employees that attempted to innovate, versus the failures of employees that worked on their standard tasks. Allowing the firm to have such a fine-grained control over tolerance for failure would make the policy clearly more powerful. We focus on the case where tolerance for failure occurs at a coarser level of corporate culture, which appears to be closer to the empirical phenomenon we are trying to model here.

¹⁹A large sociology literature discusses how corporate culture can encourage or discourage innovation. See, for instance, the influential work of March (1991).

For simplicity, suppose that ε is sufficiently small such that the agent still wants to implement effort under either tasks, i.e., $e_S^*, e_I^* > 0$. Clearly, tolerance for failure curbs effort incentives for both tasks, which constitutes indirect costs of a tolerance policy for the principal. However, the benefit of adopting a tolerance policy lies in the motivation of innovative activities, as exposed by the next lemma.

Lemma 5 *Whenever the agent is sufficiently less likely to succeed in innovation than in standard tasks, a tolerance policy induces more innovation. Formally, there exists a threshold $\mu^* > 0$ such that tolerance for failure leads to more innovation if $\mu < \mu^*$,*

Tolerance for failure encourages more exploration whenever the agent is less likely to succeed in innovation as compared to standard tasks. In this case, a tolerance policy is a viable means to promote innovative activities inside firms. Indeed, Tian and Wang (2009) provide empirical evidence that tolerance for failure leads to more innovation inside firm. However, a tolerance policy is clearly a two-edged sword: while it may promote innovation, it also impairs the agent's incentives to implement effort for exploitation and exploration. The next proposition clarifies that although adopting a tolerance policy can induce the agent to be more innovative, it is not always optimal for the principal to do so.

Proposition 8 *Suppose that tolerance for failure encourages more innovation (i.e., $\mu < \mu^*$). Then, there exists a threshold innovation appropriability level $\hat{\lambda}$, with $\hat{\lambda} \geq \lambda^{**}$, such that the principal should never be tolerant towards failure for $\lambda \leq \hat{\lambda}$, but should always be tolerant for $\lambda > \hat{\lambda}$.*

Our previous observations indicate that the agent is already too innovative for low levels of innovation appropriability. A tolerance policy would only induce more innovation at the expense of the standard task, and can thus not be optimal. For $\lambda > \lambda^{**}$, however, the principal clearly wants the agent to be more innovative (see Proposition 6). Then, the principal faces a trade-off between encouraging more innovation and, concurrently, mitigating effort incentives for both tasks. According to Proposition 8, adopting a tolerance policy is only optimal for the principal if innovation are highly firm-specific (i.e., $\lambda > \hat{\lambda}$). Only in this case, the pursuit of more innovative ideas compensates the principal for lower effort.

7.2 Investing in Employee Innovation

Employment contracts are not the only tool for the principal to influence the agent's choice between exploitation and exploration. For instance, the principal can invest in human capital, in order to improve the agent's ability to successfully develop innovation, thus making exploration more attractive. In the same vein, the principal can provide access to physical assets and infrastructure in order to facilitate the pursuit of promis-

ing ideas. For instance, Google cultivates a creative environment by allowing its employees to dedicate one day every week to pursuing their potentially innovative ideas.

To capture potential investments in employee innovation, we now endogenize the likelihood of successfully bringing an innovative idea to fruition. Technically, suppose the principal can affect the innovation process by choosing $\mu \in [\underline{\mu}, \bar{\mu}]$, where $0 < \underline{\mu} < \bar{\mu} < 1$. Improving the agent's ability to successfully develop an innovation imposes convex costs $c(\mu)$, where $c(\underline{\mu}) = 0$ and $c(\bar{\mu}) = \infty$. Intuitively, investing in employee innovation has two effects (see the Appendix for a formal proof). First, it makes exploration more attractive to the agent such that he becomes more innovative (*specialization effect*). Technically, the agent's threshold Y^A decreases in μ . Second, once pursuing a promising idea, the agent is motivated to work harder in order to turn this idea into an exploitable innovation (*incentive effect*). Technically, the agent's innovation effort e_I^* increases in μ for all $\lambda < 1$. Thus, investing in employee innovation is a viable means to promote exploration activities within firms. Clearly, according to our previous observations, this might not always be desirable from the principal's perspective. The next proposition exposes a condition for investing in employee innovation being profitable for the principal.

Proposition 9 *There exists a threshold level of the innovation appropriability λ^{***} , with $\lambda^{***} < \lambda^{**}$, such that it is optimal for the principal to invest in employee innovation for $\lambda^{***} < \lambda < 1$ (i.e., $\mu^*(\lambda) > \underline{\mu}$). Otherwise, the principal never invests in employee innovation (i.e., $\mu^*(\lambda) = \underline{\mu}$).*

The most important insight is that for sufficiently low appropriability levels ($\lambda \leq \lambda^{***}$), the principal refuses to invest in innovation, even though innovation does occur in equilibrium. Put differently, the principal deliberately keeps the return to innovation down, in order to discourage the agent from pursuing innovation. Such a behavior is clearly socially inefficient. This insight further explains why some companies lack innovation: their strategy is to focus on their core business where tasks are measurable, so that innovation becomes a distraction.

For high levels of innovation appropriability (i.e., $\lambda^{**} < \lambda < 1$), we know from Proposition 6 that the agent is not innovative enough. In this case, it is intuitively clear that the principal finds it optimal to invest in employee innovation, in order to encourage more innovation. An interesting observation can be made for intermediate values of λ (i.e., $\lambda^{***} < \lambda < \lambda^{**}$). Although the agent is already too innovative, the principal still finds it optimal to invest in employee innovation. This observation is rooted in the above emphasized *incentive effect*, which has a positive effect on the principal's expected payoff from innovation. Here, the *incentive effect* clearly outweighs the *specialization effect*, such that it is optimal for the principal to support innovative activities despite distracting the agent even more from exploitation.




Figure 5: Menu of Contracts and Expected Profits

7.3 Allocation of Intellectual Property Rights

One interesting implication from our model concerns the allocation of intellectual property rights (IPR henceforth). If one models the allocation of IPR—see Aghion and Tirole (1994) or Hellmann (2007) for two examples—then one invariably obtains the result that allocating IPR to the agent increases his share of the surplus from innovation. In other words, if λ^P (λ^A) denotes the division of rents when the principal (agent) owns the IPR, then one can expect that $\lambda^P > \lambda^A$ ($\lambda^P < \lambda^A$). Figure 5 illustrates the principal's expected profit for the optimal menu of contracts $\Omega^*(\lambda) \in \{\{0, \beta^*(\lambda)\}, \{\alpha^{**}(\lambda), 0\}\}$, and thus provides a simple visualization for evaluating the principal's decision on whether to relinquish IPR to the agent. For example, for $\lambda^P > \lambda^*$, the principal might want to relinquish IPR, whereas for $\lambda^P < \lambda^*$, she would never want to do so. Another interesting special case is where IPR convey absolute bargaining power, so that either $\lambda^P = 1$ or $\lambda^P = 0$. In this case, Figure 5 reveals that the principal would never relinquish IPR.

7.4 Continuous Task Choice

The fundamental insight from our discrete task choice model is that the incentive bonus for the standard task lies above the unconstrained level (i.e., $\beta^*(\lambda) > \beta_0$) for a low innovation appropriability ($0 \leq \lambda < \lambda^{**}$), and below the unconstrained level (i.e., $\beta^*(\lambda) < \beta_0$) for high appropriability levels ($\lambda^{**} < \lambda < 1$). In this section, we demonstrate that this result even holds in a more general setting where the agent's task choice is continuous.

In an environment with a continuous task choice and private information about Y , it is optimal for the principal to offer the agent a menu consisting of a continuum of contracts. To keep the subsequent analysis tractable, however, assume that Y can also be observed by the principal, but she needs the agent to bring

the idea to fruition. This in turn allows us to focus exclusively on a single contract which aims at balancing effort incentives and the agent's effort allocation.

To ensure that the agent makes a continuous task choice, we relax our initial assumption of constant marginal benefits for the standard task and innovation. In particular, let $\text{Prob}[S = 1|e_S] = \rho(e_S)$ be the conditional probability that $S = 1$, with $\rho'(e_S) > 0$ and $\rho''(e_S) < 0$. Moreover, the conditional probability of the innovation process being successful is $\text{Prob}[Y > 0|e_I] = \mu(e_I)$, with $\mu'(e_I) > 0$ and $\mu''(e_I) < 0$. To ensure interior solutions, we also assume that $\rho'(0) = \mu'(0) = \infty$.

Given a bonus contract (α, β) , the agent chooses his effort levels e_S and e_I to maximize his expected utility

$$U(e_S, e_I) = \alpha + \beta\rho(e_S) + (1 - \lambda)\mu(e_I) - c(e_S + e_I). \quad (21)$$

It can be deduced from the first-order conditions that the agent's second-best effort levels e_S^* and e_I^* are implicitly characterized by

$$\beta\rho'(e_S) = c'(e), \quad (22)$$

$$(1 - \lambda)\mu'(e_I) = c'(e). \quad (23)$$

The next lemma summarizes important comparative statics results.

Lemma 6 *In the continuous task choice model, the agent's*

- (i) *effort for the standard task e_S^* is increasing in the incentive bonus β and in the innovation appropriability λ (i.e., $de_S^*/d\beta > 0$ and $de_S^*/d\lambda > 0$);*
- (ii) *innovation effort e_I^* is decreasing in the incentive bonus β and in the innovation appropriability λ (i.e., $de_I^*/d\beta < 0$ and $de_I^*/d\lambda < 0$);*
- (iii) *total effort $e^* = e_S^* + e_I^*$ is increasing in the incentive bonus β and decreasing in the innovation appropriability λ (i.e., $de^*/d\beta > 0$ and $de^*/d\lambda < 0$).*

It is intuitively clear that more standard incentives alter the agent's effort allocation towards a stronger focus on the standard task. The opposite observation can be made for the agent's incentives to innovative, which is (reversely) captured by the appropriability parameter λ .

To illustrate potential conflicts of interests with respect to the agent's effort allocation, we now identify the effort allocation desired by the principal. For a given innovation appropriability λ , the principal prefers the effort levels e_S and e_I which maximize

$$\Pi(e_S, e_I) = V(e_S) + \lambda Y \mu(e_I) - c(e_S + e_I). \quad (24)$$

The effort levels desired by the principal, denoted e_S^P and e_I^P , are implicitly characterized by the following first-order conditions:

$$V'(e_S) = c'(e), \quad (25)$$

$$\lambda Y \mu'(e_I) = c'(e). \quad (26)$$

The next lemma contrasts the principal's with the agent's preferences for innovation, and thus provides an intuitive rationale for the subsequently investigated contract adjustments.

Lemma 7 *There exists a threshold level of the innovation appropriability λ^{**} such that the agent focuses too much on innovation for $\lambda \leq \lambda^{**}$ (i.e., $e_I^* \geq e_I^P$), and too little for $\lambda > \lambda^{**}$ (i.e., $e_I^* < e_I^P$).*

As for the discrete task choice model, the conflict of interests between the principal and the agent is determined by the appropriability of innovation opportunities. In particular, low levels of innovation appropriability induce the agent to be too innovative from the principal's perspective, while the opposite is true for high appropriability levels.

It remains to characterize the optimal bonus contract (α^*, β^*) . The principal's problem can be stated as follows:

$$\max_{\alpha, \beta} \Pi(\alpha, \beta) = V(e_S^*) - \alpha - \rho(e_S^*)\beta + \lambda Y \mu(e_I^*), \quad (27)$$

where $\alpha \geq 0$ and $\alpha + \beta \geq 0$. Note first that the agent's liability limit implies $\alpha^* = 0$.²⁰ For convenience, let $\Pi^S \equiv V(e_S^*) - \rho(e_S^*)\beta$ and $\Pi^I \equiv \lambda Y \mu(e_I^*)$. Accordingly, the optimal bonus $\beta^*(\lambda)$ is implicitly characterized by the following first-order condition:

$$\frac{d\Pi(\beta)}{d\beta} = \frac{d\Pi^S}{d\beta} + \frac{d\Pi^I}{d\beta} = 0 \quad (28)$$

$$= [V'(e_S^*) - \beta \rho'(e_S^*)] \frac{de_S^*}{d\beta} - \rho(e_S^*) + \lambda Y \mu'(e_I^*) \frac{de_I^*}{d\beta} = 0. \quad (29)$$

The first term on the rhs of (28) represents the effect of increasing β on the principal's net payoff from the standard task. The second term captures the decline in her expected net payoff from innovation when enhancing incentives for the standard task, and hence, distracting the agent from pursuing identified innovation opportunities. As in the previous discrete task choice model, the optimal bonus $\beta^*(\lambda)$ thus needs to balance incentives for the standard task and the agent's motivation to develop innovation. To evaluate the optimal bonus $\beta^*(\lambda)$ for different values of λ , it is essential to first characterize the optimal bonus in the absence

²⁰See Proof of Proposition 1 for a formal derivation.

of innovation opportunities, i.e., the unconstrained bonus β_0 . To do so, suppose for a moment that $Y = 0$. Then, the first-order condition (28) simplifies to

$$\frac{d\Pi^S}{d\beta} = 0. \quad (30)$$

Clearly, in the absence of innovation opportunities, we obtain the same unconstrained incentive payment β_0 as in the discrete task choice model. We are now equipped to characterize the optimal bonus $\beta^*(\lambda)$ for different levels of the innovation appropriability λ .

Proposition 10 *For sufficiently low levels of innovation appropriability λ , the optimal bonus $\beta^*(\lambda)$ lies above the unconstrained level (i.e., $\beta^*(\lambda) > \beta_0$). In contrast, for sufficiently high levels of innovation appropriability λ , with $\lambda < 1$, the optimal bonus $\beta^*(\lambda)$ lies below the unconstrained level (i.e., $\beta^*(\lambda) < \beta_0$).*

According to Proposition 10, our main insight from the discrete task choice model—namely that low innovation appropriability levels are accompanied with higher-powered incentives for the standard task, and high innovation appropriability levels with lower-powered incentives—also holds in a more general continuous task choice framework.

7.5 Stock-based Compensation

Our analysis rests on a fundamental assumption that innovation are not contractible. This is an intuitive assumption: the outcome of certain tasks cannot be specified in advance, and we define these tasks as innovation. One interesting objection to this line of reasoning is that even if innovation cannot be specified in advance, it may still be possible to design incentive contracts for them. On a purely theoretical basis, Maskin and Tirole (1999) propose a mechanism design solution to this problem, showing how the problem of unspecified contingencies can be solved by playing a revelation game that is based on utilities implied by these unspecified contingencies. The mechanism proposed by Maskin and Tirole, however, has been criticized for being too abstract to be of any practical use. In addition we note that their mechanism relies on risk-averse agents and the absence of any wealth constraints.

A more practical solution to the non-contractibility problem seems to be the use of stocks or stock options. Indeed, it is often argued that employee stock options are a useful way of motivating employees in innovative firms. In this section, we examine whether stock options can solve the multi-tasking problem between planned and unplanned activities.

Stock options clearly rely on a company either being listed on a stock exchange, or else having clear expectations of going public or being acquired in the foreseeable future. Stock options are therefore not

a solution for the vast majority of privately held firms. Even if a company is publicly listed, there are many problems with using stock options to compensate unplanned activities. First, the firm must have a good understanding of the timing of the innovation in order to set an appropriate maturity date for the stock option. Second, the firm must be able to prevent the employee from hedging his financial position in the company. Third, it is rarely possible to price employee stock options in a manner that actually provides efficient incentives.

Let us elaborate on the latter point. If the only uncertainty about the firm's value is due to one employee's choice of the assigned versus the unplanned task, then it is theoretically possible to set a strike price of the option such that this employee captures enough of the appreciation value to face efficient incentives. However, as soon as we allow for other uncertainty in the value of the firm, setting a strike price to provide efficient incentives becomes more problematic. With a stochastic firm value, stock options are either in the money, in which case employees' incentives to create innovation are severely diluted; or stock options are under water, in which case they have no incentive effect at all. Moreover, awarding stock options is expensive for the firm as it requires a transfer of rents to employees.

To illustrate our last argument, let us consider a simple extension of our base model which allows the principal to award stock options to the agent. For simplicity, suppose the firm exists only for two periods, and profits are paid out in the last period. The firm's total profit in the last period, denoted Π , is

$$\Pi = \tilde{\Pi} + \Pi^A(e_S, e_I, Y, \lambda), \quad (31)$$

where $\Pi^A(e_S, e_I, Y, \lambda)$ denotes the agent's contribution to firm value, and $\tilde{\Pi}$ as the value which stems from other assets and agents. Assuming a perfect capital market, the stock price p in $t = 1$ is

$$p = \frac{1}{n}\Pi = \frac{1}{n} \left[\tilde{\Pi} + \Pi^A(e_S, e_I, Y, \lambda) \right], \quad (32)$$

where n denotes the number of stocks.

From an ex ante perspective, the firm's total profit—and thus the stock price p —is a continuous random variable with $\Pi \in [0, \infty)$ and standard deviation σ . Assuming zero discounting and letting \bar{p} denote the strike price, the expected value φ of one stock option is

$$\varphi(\bar{p}, n) = \text{E}[\max\{p - \bar{p}, 0\}] = \text{E} \left[\max \left\{ \frac{1}{n}\Pi - \bar{p}, 0 \right\} \right]. \quad (33)$$

Recall that $\Pi = \tilde{\Pi} + \Pi^A(e_S, e_I, Y, \lambda)$, where $\tilde{\Pi}$ and Π^A are assumed to be independent random variables. Let $\Psi(\Pi|e_S, e_I, Y)$ denote the cumulative distribution function of Π conditional on the agent effort levels e_S and e_I , and the (privately) observed innovation value Y . The value of one stock option therefore is

$$\varphi(e_S, e_I, Y, \lambda, \bar{p}, n) = \int_{n\bar{p}}^{\infty} \left(\frac{1}{n} \Pi - \bar{p} \right) d\Psi(\Pi|e_S, e_I, Y), \quad (34)$$

where the lower bound ($n\bar{p}$) follows from the fact that the stock option will only be exercised if $\Pi/n \geq \bar{p}$.

To illustrate the rationale behind the subsequent analysis, suppose there are several homogenous agents in the firm, all conducting the same standard and innovative tasks as in our base model. From the perspective of one agent, the future stock price p has a low variability (i.e., σ is low) in case very few agents are employed. If, on the other hand, many agents are employed, individual agents perceive the future stock price as highly uncertain as their relative impact is only very small (i.e., σ is large). We can thus utilize the standard deviation σ as a natural measure of an individual agent's relative contribution to firm value.²¹ Moreover, as well known, the value of a stock option $\varphi(\cdot)$ is increasing in the standard deviation σ for a given exercise price \bar{p} , i.e., $\partial\varphi(\cdot)/\partial\sigma > 0$ with $\lim_{\sigma \rightarrow \infty} \varphi(\cdot) = \infty$. This is also very intuitive as the agent is protected from downside risks (due to the lower bound $n\bar{p}$), but can clearly benefit from potentially higher stock prices.

We can now analyze how the agent's specialization decision responds to the provision of stock options for a given strike price \bar{p} . To do so, suppose the agent receives m stock options, where m is measured in infinitesimal small units so that it becomes a continuous variable. For the sake of consistency with our base model, we focus on the provision of stock options in addition to the incentive bonus β .²² The next lemma emphasizes the *task choice effect* which stems from the provision of stock options.²³

Lemma 8 *Providing the agent with a small number of stock options makes him less innovative for $\lambda \leq \lambda^{**}$, and more innovative for $\lambda > \lambda^{**}$ (i.e., $Y^A(\cdot, m)$ is increasing in m , with $m \rightarrow 0$, for $\lambda \leq \lambda^{**}$; and decreasing in m for $\lambda > \lambda^{**}$).*

Clearly, stock options can be used to mitigate the task choice conflict between the principal and the agent. This is a very intuitive result: by awarding stock options to the agent, his compensation becomes more sensitive to the firm's actual profits, which in turn aligns the agent's with the principal's preferences for the choice between the standard task and innovation.

²¹Alternatively, one can use σ as an indicator of an agent's relative position in the corporate hierarchy, where a low σ corresponds to agents on lower hierarchy levels, and vice versa.

²²By letting $\beta = 0$ throughout the subsequent analysis, one could easily obtain the case where the principal awards only stock options to the agent. As will become clear, however, this would not effect the qualitative nature of our results in this section.

²³Besides this *task choice effect*, we demonstrate in the Appendix that the provision of stock options has also an *incentive effect*. In particular, for a given incentive bonus β , with $\beta \leq \beta^*$, providing the agent with more stock options induces more effort for the standard task as well as for the innovative task (i.e., e_S^* and e_I^* increase).

So far, we have ignored the costs for the principal associated with stock options. We complete our analysis by characterizing the optimal provision of stock options. Let $C(m, \cdot)$ denote the principal's costs of awarding m stock options to the agent, with $C(m, \cdot) = m\varphi(\cdot)$. By accounting for these costs, and using the agent's adjusted optimal effort levels $e_S^*(\beta, m)$ and $e_I^*(\lambda, Y, m)$, the principal's expected profit becomes

$$\Pi(\cdot, m) = \int_0^{Y^A} [V(e_S^*) - \beta e_S^*] d\Phi(Y) + \int_{Y^A}^{\infty} [\lambda Y \mu e_I^*] d\Phi(Y) - C(m, \cdot). \quad (35)$$

We are now interested to see whether the optimal provision of stock options can indeed eliminate the task choice conflict between the principal and the agent. From the first-order condition, the optimal number of stock options m^* is implicitly characterized by

$$[\Pi^S - \Pi^I] \phi(Y^A) \frac{dY^A}{dm} = \frac{\partial C(\cdot)}{\partial m}. \quad (36)$$

Consider first the principal's marginal benefit of using stock options reflected by the lhs of (36).²⁴ By drawing on Lemma 8, one can immediately infer that the marginal benefit is positive because (i) $dY^A/dm > 0$ whenever $\Pi^S > \Pi^I$ (which is true for $\lambda < \lambda^{**}$); and (ii), $dY^A/dm < 0$ whenever $\Pi^S < \Pi^I$ (which is true for $\lambda > \lambda^{**}$). Moreover, the marginal benefit is clearly zero if $\Pi^S = \Pi^I$ (which is true for $\lambda = \lambda^{**}$). In the latter case, the principal and the agent already have congruent preferences, and thus, there is no need to award stock options to the agent with the aim of improving his task choice.

The next proposition emphasizes that—although it can improve his task choice as pointed out by Lemma 8—providing the agent with stock options is not always beneficial for the principal.

Proposition 11 *If the agent's relative contribution to firm value is sufficiently small (i.e., $\sigma_{\bar{\Pi}}$ is sufficiently large), it is optimal for the principal to refrain from providing the agent with stock options (i.e., $m^* = 0$).*

The main insight from Proposition 11 is that while stock options can be used to manipulate the agent's innovation decision, there are considerable costs to the principal. Indeed, if stock prices are influenced by many other factors, it is optimal for the principal to forgo the use of stock options entirely. Even if the principal uses stock options, the optimal level of stock options never eliminates the task choice conflict between the principal and the agent. Thus, in equilibrium, the agent is still either too innovative (for $\lambda < \lambda^{**}$), or not innovative enough (for $\lambda > \lambda^{**}$) from the principal's perspective.

While stock options might seem like a theoretically attractive solution to the conflict between standard tasks and innovation, our analysis suggests that they are likely to be of limited practical use. For many

²⁴Note that the first-order condition (36) for m^* implicitly interacts with the first-order condition for the optimal incentive bonus β^* , which we have omitted for parsimony purposes. To derive our main results for this section, however, it is sufficient to restrict our attention to (36).

smaller firms that never plan to go public or get acquired, there is no stock price or related firm performance metric that can be used to base the option on. For fast-growing or larger companies that either are public or expect to be sold or go public in the foreseeable future, our analysis shows that stock options are an expensive and imperfect solution. Indeed, if the impact of an individual employee on the overall firm performance is small, as we would expect for line-of-business employees in these types of companies, we showed that the optimal firm policy may well be to not issue any stock options at all. Moreover, if firms still issue stock options, we should not expect them to fully resolve the conflict between standard tasks and innovation.

8 Empirical Implications

Our theoretical analysis already provides an explanation of several commonly observed firm behaviors with respect to employee innovation. We now turn to a more systematic evaluation of the empirical predictions that emanate from our theory.

One central prediction concerns the relationship between the appropriability of innovation and employees' compensations. In particular, our model predicts a U-shaped relationship between the appropriability and the strength of incentive payments for (non-innovative) core tasks. At the same time, it also predicts an inverted U-shaped relationship between the appropriability of innovation and the base pay for employees. To the best of our knowledge, these non-linear predictions are unique in the literature.

The model also generates a number of auxiliary empirical predictions. It predicts that tolerance for failure should be seen mostly in environments where employees have difficulty obtaining rents from their innovation. Moreover, firms devote more resources to supporting their employees to innovate if these innovation are easily appropriable by the firms.

These predictions are quite intuitive. Examples of industries with low levels of appropriability might include software and design, where much of the value of an innovation is tied to the human capital of the employee. In such industries we predict high-powered incentives for core tasks, little support by firms for innovation, no tolerance for failure, and also a significant number of spin-offs. For intermediate levels of innovation appropriability, we expect both lower-powered incentives and also fewer spin-offs. This might describe industries such as hardware or simple manufacturing. A high level of innovation appropriability is typically due to the presence of strong complementary assets, such as in the airline industry, or in branded consumer goods industries. In such industries we would expect very few spin-offs, yet some intermediate use of performance pay.

These predictions naturally await careful econometric testing. We recognize several empirical challenges. There is a non-trivial measurement problem around the appropriability of employee innovation.

The seminal work of Levin, Klevorick, Nelson, and Winter (1987) uses survey-based measures of industry appropriability. The work of Klepper and Sleeper (2005) suggests using spin-offs as an alternative measure, although one should note that this measure might be partially affected by endogenous firm choices (Hellmann, 2007). Another empirical challenge concerns correctly identifying the direction of causal relationships, and controlling for unobserved heterogeneity. Subramanian (2005), for example, finds evidence that stock options and spin-offs are positively correlated in a cross-section. However, such a correlation does not prove causality. Probably the most promising identification strategy would be to use panel data and identify an exogenous shock to appropriability of innovation, such as a change in the allocation of intellectual property rights between employers and employees (Merges, 1999).

9 Conclusion

We provide a theory that integrates complete and incomplete contracting in the context of employment relationships. Incentive contracts are feasible for those tasks that are well understood and measurable ex-ante. Innovation, however, involve tasks that—by definition—cannot be anticipated ahead of time. This paper analyzes how the presence of such unplanned activities interacts with the provision of optimal incentives for planned activities.

The appropriability of employee innovation is found to play a pivotal role. When firms can appropriate most of the value of their employees' innovation, employees are discouraged from innovation. In such situations firms want to reduce incentives for planned activities, in order to create some 'elbow room' for innovation. This changes dramatically if employees can capture most of the value of their innovation. In this case, firms use stronger performance pay to keep their employees focused on their assigned core tasks.

Our theory is framed within a principal-agent framework with symmetric ex-ante information. It abstracts from questions such as which kind of employees are attracted to more or less innovative firms. One fruitful area of future research therefore concerns the interaction between the incentive effects discussed in this paper, and the self-selection process by which innovative employees choose their employers. Closely related to this, we note that our model takes a somewhat simplified view of the innovation process, where employees generate innovation in isolation. Yet, innovation appear to rely on interactions among employees, often with diverse backgrounds. Thinking about how employees interact in their planned and unplanned activities is therefore another topic worthy of future research.

Appendix

Proof of Lemma 1.

From an inspection of the agent's utility function $U^S(\cdot)$, it becomes clear that $U^S(\cdot)$ is strictly increasing in α . Next, note that $\beta = 0$ leads to $U(e_S(0), 0) = 0$. Moreover, applying the Envelope Theorem yields $dU(\cdot)/d\beta = e_S^*(\beta) > 0$. Hence, the agent extracts a rent for all $\beta > 0$, which is increasing in β .

Observe further that $\Pi^S(\alpha, \beta)$ is strictly decreasing in α . Moreover, due to concavity of $V(e_S)$ and convexity of $c(e)$, $\Pi^S(\alpha, \beta)$ must be concave in β . To derive the optimal incentive contract (α_0, β_0) , note first that $e_S > 0$ requires $\beta > 0$. Thus, $\alpha + \beta \geq 0$ is satisfied if $\alpha \geq 0$ holds. Consequently, we can ignore the former limited liability constraint. Let γ and ξ denote Lagrange multipliers for the agent's participation constraint, and for his liability limit constraint respectively. Then, the Lagrangian is

$$\mathcal{L}(\alpha, \beta) = V(e_S^*) - \alpha - \beta e_S^* + \gamma [\alpha + \beta e_S^* - c(e_S^*)] + \xi \alpha. \quad (37)$$

The first-order conditions with respect to α and β are

$$-1 + \gamma + \xi = 0, \quad (38)$$

$$\left[V'(e_S^*) - \beta \right] \frac{de_S^*}{d\beta} - e_S^* + \gamma \left[e_S^* + (\beta - c'(e_S^*)) \frac{de_S^*}{d\beta} \right] = 0. \quad (39)$$

Suppose $\gamma > 0$. Then, $\alpha + \beta e_S^* - c(e_S^*) = 0$ due to complementary slackness. Since $\alpha \geq 0$, this would imply that $\beta e_S^* \leq c(e_S^*)$, and hence $e_S = 0$. Thus, $\gamma > 0$ cannot be a solution such that $\gamma = 0$. We can then infer from (38) that $\xi = 1$. Consequently, $\alpha_0 = 0$ due to complementary slackness. Since $\gamma = 0$, it follows from (39) that β^* solves

$$\left[V'(e_S^*) - \beta \right] \frac{de_S^*}{d\beta} - e_S^* = 0, \quad (40)$$

which is equivalent to $d\Pi^S(\cdot)/d\beta = 0$. □

Proof of Lemma 2.

Given the agent's effort choice e_I^* , the principal's expected profit from innovation is

$$\Pi^I(e_I^*, \lambda, Y) = \lambda Y \mu e_I^*. \quad (41)$$

Taking the first derivative with respect to λ yields

$$\frac{d\Pi^I(\cdot)}{d\lambda} = Y\mu e_I^* + \lambda Y\mu \frac{de_I^*}{d\lambda}. \quad (42)$$

Recall that $de_I^*/d\lambda < 0$. Thus, for $\lambda \in (0, 1)$, the first term on the right side of (42) is always positive while the second term is always negative. Moreover, it can be easily shown that the first term is decreasing in λ , and that the second term is increasing in λ . This in turn implies that there exists a unique threshold λ^* such that $d\Pi^I(\cdot)/d\lambda|_{\lambda=\lambda^*} = 0$. To identify the sign of (42) for all $\lambda \neq \lambda^*$, suppose for a moment that $\lambda \rightarrow 0$. Then, the second term goes to zero while the first term is strictly positive. Hence, $d\Pi^I(\cdot)/d\lambda \geq 0$ for $\lambda \leq \lambda^*$. For $\lambda \rightarrow 1$, the first term goes to zero (because $e_I^* = 0$ for $\lambda = 1$) while the second term is strictly negative. Thus, $d\Pi^I(\cdot)/d\lambda < 0$ for $\lambda > \lambda^*$. \square

Proof of Lemma 3.

Note first that the agent's marginal expected benefit from standard task is β , and from innovation $(1 - \lambda)Y\mu$. Thus, both tasks are characterized by constant marginal expected benefits. Moreover, recall that the agent's cost of effort is convex in total effort e . These observations imply that the agent maximizes his expected utility by choosing only one of the two tasks, i.e., he either implements $\{e_S^* > 0, e_I^* = 0\}$ or $\{e_S^* = 0, e_I^* > 0\}$.

Next, observe that the agent's expected payoff from standard task is strictly positive for $\beta > 0$ and constant in Y . In contrast, one can show by applying the Envelope Theorem that the agent's expected payoff from innovation is zero for $Y = 0$ and strictly increasing in Y . Thus, we can infer that there exists a unique threshold innovation value Y^A such that for $Y < Y^A$, the agent implements $\{e_S^* > 0, e_I^* = 0\}$, and $\{e_S^* = 0, e_I^* > 0\}$ for $Y > Y^A$, respectively. \square

Proof of Proposition 1.

Recall that $Y^A(\beta, \lambda)$ is implicitly defined by $U^S(e_S^*, \beta) = U^I(e_I^*, \lambda, Y^A)$. First, we show that $dY^A/d\beta > 0$. Implicit differentiating Y^A with respect to β yields

$$\frac{dY^A}{d\beta} = \frac{\frac{dU^S(\cdot)}{d\beta}}{\frac{dU^I(\cdot)}{dY}\Big|_{Y=Y^A}}. \quad (43)$$

As previously shown, $dU^S(\cdot)/d\beta > 0$. Moreover, applying the Envelope Theorem gives

$$\frac{dU^I(\cdot)}{dY}\Big|_{Y=Y^A} = (1 - \lambda)\mu e_I^* > 0. \quad (44)$$

Thus, $dY^A/d\beta > 0$. Next, we demonstrate that $dY^A/d\lambda > 0$. By implicitly differentiating Y^A with respect to λ one get

$$\frac{dY^A}{d\lambda} = -\frac{\frac{dU^I(\cdot)}{d\lambda}}{\frac{dU^I(\cdot)}{dY}\Big|_{Y=Y^A}}. \quad (45)$$

Again, $dU^I(\cdot)/dY|_{Y=Y^A} > 0$. Furthermore, applying the Envelope Theorem yields

$$\frac{dU^I(\cdot)}{d\lambda} = -\mu e_I^* < 0. \quad (46)$$

Thus, $dY^A/d\lambda > 0$. □

Proof of Lemma 4.

Recall that $Y^P(\beta_0, \lambda)$ is implicitly characterized by $\Pi^S(e_S^*, \beta) = \Pi^I(e_I^*, \lambda, Y^P)$. Implicit differentiating Y^P with respect to λ yields

$$\frac{dY^P}{d\lambda} = -\frac{\frac{d\Pi^I(\cdot)}{d\lambda}}{\frac{d\Pi^I(\cdot)}{dY}\Big|_{Y=Y^P}}. \quad (47)$$

Observe that

$$\frac{d\Pi^I(\cdot)}{dY}\Big|_{Y=Y^P} = \lambda\mu e_I^* + \lambda Y^P \mu \frac{de_I^*}{dY}\Big|_{Y=Y^P}. \quad (48)$$

Since $de_I^*/dY|_{Y=Y^P} > 0$, we have $d\Pi^I(\cdot)/dY|_{Y=Y^P} > 0$. Moreover, recall from Lemma 2 that $d\Pi^I(\cdot)/d\lambda \geq 0$ for $\lambda \leq \lambda^*$, and $d\Pi^I(\cdot)/d\lambda < 0$ for $\lambda > \lambda^*$, respectively. Consequently, $dY^P/d\lambda \leq 0$ for $\lambda \leq \lambda^*$, and $dY^P/d\lambda > 0$ for $\lambda > \lambda^*$. □

Proof of Proposition 2.

First, to prove existence of a unique threshold λ^{**} such that $Y^A(\beta_0) = Y^P(\beta_0)$ for $\lambda = \lambda^{**}$, it is necessary to characterize Y^A and Y^P for the extreme values of λ , i.e., $\lambda = 0, 1$. Recall that Y^P is implicitly characterized by $\Pi^S(e_S^*, \beta) = \Pi^I(e_I^*, \lambda, Y^P)$, which is equivalent to

$$V(e_S^*) - \beta_0 e_S^* = \lambda Y^P \mu e_I^*. \quad (49)$$

Suppose that $\lambda \rightarrow 0$. In this case, the rhs of (49) goes to zero while the lhs is strictly positive. Thus, (49) requires that $Y^P \rightarrow \infty$ for $\lambda \rightarrow 0$. Now suppose that $\lambda \rightarrow 1$. In this case, $e_I(\cdot) \rightarrow 0$, implying that the rhs

of (49) goes to zero while the lhs is strictly positive. Again, to ensure that (49) is satisfied, $Y^P \rightarrow \infty$. Next, recall that Y^A is implicitly characterized by $U^S(e_S^*, \beta) = U^I(e_I^*, \lambda, Y^A)$, which is equivalent to

$$\beta_0 e_S^* - c(e_S^*) = (1 - \lambda) Y^A \mu e_I^* - c(e_I^*). \quad (50)$$

Again, suppose for a moment that $\lambda \rightarrow 0$. In this case, both sides of (50) are strictly positive, implying that $Y^A \in (0, \infty)$. Now suppose that $\lambda \rightarrow 1$. Then, the rhs of (50) goes to zero while the lhs is strictly positive. This in turn requires that $Y^A \rightarrow \infty$. Next, we need verify that $Y^A > Y^P$ for $\lambda \rightarrow 1$. We prove that $dY^A/d\lambda > dY^P/\lambda$ for all λ . We can infer from Proof of Proposition 1 and Proof of Lemma 4 that $dY^A/d\lambda > dY^P/\lambda$ is equivalent to

$$\begin{aligned} & -\frac{\frac{dU^I(\cdot)}{d\lambda}}{(1-\lambda)\mu e_I^*} > -\frac{\frac{d\Pi^I(\cdot)}{d\lambda}}{\lambda\mu e_I^*} \\ \Leftrightarrow & \frac{-Y\mu e_I^* + (1-\lambda)Y\mu \frac{de_I^*}{d\lambda}}{(1-\lambda)\mu e_I^*} < \frac{Y\mu e_I^* + \lambda Y\mu \frac{de_I^*}{d\lambda}}{\lambda\mu e_I^*} \\ \Leftrightarrow & 0 < Y\mu e_I^*, \end{aligned}$$

which is clearly satisfied for all $\lambda > 0$. Hence, $Y^A > Y^P$ for $\lambda \rightarrow 1$. By combining the previous observations, we can infer that there exists a unique threshold λ^{**} such that $Y^A(\beta_0, \lambda) \leq Y^P(\beta_0, \lambda)$ for $\lambda \leq \lambda^{**}$, and $Y^A(\beta_0) > Y^P(\beta_0)$ otherwise.

Finally, it remains to demonstrate that $\lambda^{**} < \lambda^*$. For this purpose, we need to draw on the optimal bonus $\beta^*(\lambda)$ which will be derived in section 5.1. As will be shown in Proof of Proposition 3, $d\beta^*(\lambda)/d\lambda \leq 0$ for $\lambda \leq \lambda^*$, where $\beta^*(0) > \beta_0$ and $\beta^*(\lambda^*) < \beta_0$. Since $Y^A(\cdot) = Y^P(\cdot)$ at $\lambda = \lambda^{**}$, it must hold that $\beta^*(\lambda^{**}) = \beta_0$. This in turn implies that $\lambda^{**} < \lambda^*$. \square

Proof of Proposition 3.

First, we characterize the optimal bonus $\beta^*(\lambda)$ for the extreme values of λ : $\lambda = 0, 1$. Suppose for a moment that $\lambda = 0$. Since $\Pi^I(\lambda = 0) = 0$, the first-order condition (12) then becomes

$$\frac{d\Pi^S}{d\beta} = -\Pi^S \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\beta}. \quad (51)$$

Note that the rhs of (51) is strictly negative, implying that $d\Pi^S/d\beta|_{\beta=\beta^*} < 0$. Since $\Pi^S(\cdot)$ is concave in β with its maximum at $\beta = \beta_0$, we have $d\Pi^S(\cdot)/d\beta|_{\beta < \beta_0} > 0$, and $d\Pi^S(\cdot)/d\beta|_{\beta > \beta_0} < 0$. Thus, we can infer

that $\beta^*(0) > \beta_0$. Now suppose that $\lambda = 1$. Then, the agent chooses $e_I^*(\lambda = 1) = 0$, and the first-order condition (12) simplifies to

$$\frac{d\Pi^S}{d\beta} = 0, \quad (52)$$

which is equivalent to the first-order condition characterizing β_0 , see Lemma 1. Hence, $\beta^*(1) = \beta_0$. Now consider the case $\lambda = \lambda^{**}$. Recall from Proposition 2 that $\lambda = \lambda^{**}$ implies $Y^A = Y^P$, i.e., preferences are perfectly aligned. In this case, the first-order condition (12) simplifies to (52). Consequently, $\beta^*(\lambda = \lambda^{**}) = \beta_0$. Next, we investigate whether $\beta^*(\lambda)$ lies below or above the unconstrained level β_0 for $\lambda \neq 0, 1$. According to Proposition 2, $Y^P(\beta_0) > Y^A(\beta_0)$ for $\lambda < \lambda^{**}$. Since $dY^A(\beta)/d\beta > 0$, it must hold that $\beta^*(\lambda) > \beta_0$ for $\lambda < \lambda^{**}$. The opposite observation can be made for $\lambda > \lambda^{**}$. Next, we identify how $\beta^*(\lambda)$ responds to a change of λ . First define

$$F_{\beta^*} \equiv U^I(\cdot) - U^S(\cdot) \quad (53)$$

$$G_{\beta^*} \equiv \frac{d\Pi^S}{d\beta} - [\Pi^I - \Pi^S] \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\beta}. \quad (54)$$

Applying Cramer's rule yields

$$\frac{d\beta^*(\lambda)}{d\lambda} = \frac{\det(A)}{\det(B)}, \quad (55)$$

where

$$A \equiv \begin{pmatrix} -\frac{\partial F_{\beta^*}}{\partial \lambda} & \frac{\partial F_{\beta^*}}{\partial Y^A} \\ -\frac{\partial G_{\beta^*}}{\partial \lambda} & \frac{\partial G_{\beta^*}}{\partial Y^A} \end{pmatrix} \quad B \equiv \begin{pmatrix} \frac{\partial F_{\beta^*}}{\partial \beta} & \frac{\partial F_{\beta^*}}{\partial Y^A} \\ \frac{\partial G_{\beta^*}}{\partial \beta} & \frac{\partial G_{\beta^*}}{\partial Y^A} \end{pmatrix}.$$

Since $\Pi(\cdot)$ is concave, B must be negative definite. Hence, $\det(B) > 0$. Therefore, $d\beta^*(\lambda)/d\lambda > 0$ if

$$\det(A) = -\frac{\partial F_{\beta^*}}{\partial \lambda} \frac{\partial G_{\beta^*}}{\partial Y^A} + \frac{\partial G_{\beta^*}}{\partial \lambda} \frac{\partial F_{\beta^*}}{\partial Y^A} > 0. \quad (56)$$

By applying the Envelope Theorem, one get

$$\frac{\partial F_{\beta^*}}{\partial \lambda} = -Y^A \mu e_I^* \quad (57)$$

$$\frac{\partial F_{\beta^*}}{\partial Y^A} = (1 - \lambda) \mu e_I^*. \quad (58)$$

Moreover,

$$\frac{\partial G_{\beta^*}}{\partial Y^A} = -[\Pi^I - \Pi^S] \frac{\partial}{\partial Y^A} \left(\frac{\phi(Y^A)}{\Phi(Y^A)} \right) \frac{dY^A}{d\beta} \quad (59)$$

$$\frac{\partial G_{\beta^*}}{\partial \lambda} = -\frac{d\Pi^I}{d\lambda} \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\beta}. \quad (60)$$

Therefore, $d\beta^*(\lambda)/d\lambda > 0$ if $\det(A) > 0$, which—after eliminating μe_I^* —is equivalent to

$$Y^A [\Pi^I - \Pi^S] \frac{\partial}{\partial Y^A} \left(\frac{\phi(Y^A)}{\Phi(Y^A)} \right) \frac{dY^A}{d\beta} + (1 - \lambda) \frac{d\Pi^I}{d\lambda} \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\beta} < 0. \quad (61)$$

Now we need to consider three different intervals: (i) $0 < \lambda < \lambda^{**}$, (ii) $\lambda^{**} < \lambda < \lambda^*$; and, (iii) $\lambda^* < \lambda < 1$. Consider first the interval $0 < \lambda < \lambda^{**}$. As shown, $\beta^*(\lambda) > \beta_0$ for this interval. Observe that $\beta^*(\lambda) > \beta_0$ implies that the rhs of the first-order condition (12) must be strictly negative. Hence, $\Pi^I - \Pi^S < 0$. Next, we need to identify the sign of $\det(A)$. It can be easily verified that $\frac{\partial}{\partial Y^A} \left(\frac{\phi(Y^A)}{\Phi(Y^A)} \right) < 0$. Moreover, recall from Lemma 2 that $d\Pi^I/d\lambda > 0$ for $0 < \lambda < \lambda^*$, where $\lambda^{**} < \lambda^*$. Consequently, both terms on the lhs of (61) are strictly positive, implying that $\det(A) < 0$. Thus, $d\beta^*(\lambda)/d\lambda < 0$ for $0 < \lambda < \lambda^{**}$. Next, consider the interval $\lambda^{**} < \lambda < \lambda^*$. Here, $\beta^*(\lambda) > \beta_0$ would be a contradiction to Proposition 2. Therefore, we can infer from Proposition 2 that $\beta^*(\lambda) < \beta_0$ for $\lambda^{**} < \lambda < \lambda^*$. Moreover, recall that $d\Pi^I/d\lambda > 0$ for $\lambda \leq \lambda^*$, and $Y^P(\cdot) < Y^A(\cdot)$ with $dY^P(\cdot)/d\lambda < 0$ for $\lambda^{**} < \lambda < \lambda^*$ (see Lemmas 2 and 4, and Proposition 2). Thus, $d\beta^*(\lambda)/d\lambda > 0$ would be a contradiction. Consequently, it must hold that $d\beta^*(\lambda)/d\lambda < 0$ for $\lambda^{**} < \lambda < \lambda^*$. At $\lambda = \lambda^*$, Π^I is maximized (see Lemma 2). As will be subsequently shown, $d\beta^*(\lambda)/d\lambda > 0$ for $\lambda > \lambda^*$. This in turn implies that $dY^A/d\beta|_{\lambda=\lambda^*} = 0$. Moreover, $d\Pi^I/d\lambda|_{\lambda=\lambda^*} = 0$. Thus, both terms on the lhs of (61) become zero, implying that $\beta^*(\lambda)$ is minimized at $\lambda = \lambda^*$. Finally, consider the interval $\lambda^* < \lambda < 1$. Here again, $\beta^*(\lambda) > \beta_0$ would contradict Proposition 2. Therefore, we can infer from Proposition 2 that $\beta^*(\lambda) < \beta_0$ for $\lambda^* < \lambda < 1$. This implicates that the rhs of (12) must be strictly positive, which requires that $\Pi^I - \Pi^S > 0$. Next, we need to identify the sign $\det(A)$. Recall that $d\Pi^I/d\lambda < 0$ for $\lambda^* < \lambda < 1$ (see Lemma 2). Hence, both terms on the lhs of (61) are strictly negative, implying that $\det(A) > 0$. Consequently, $d\beta^*(\lambda)/d\lambda > 0$ for $\lambda^* < \lambda < 1$. \square

Proof of Proposition 4.

First, we need to investigate how $Y^A(\alpha^{**}, \beta^*(\lambda), \lambda)$ responds to a change of α^{**} , everything else equal. Recall that Y^A is implicitly characterized by $U^S(e_S^*, \beta^*(\lambda)) = U^I(e_I^*, \alpha^{**}, \lambda, Y^A)$. Implicit differentiating Y^A with respect α^{**} yields

$$\frac{dY^A}{d\alpha^{**}} = -\frac{\frac{dU^I(\cdot)}{d\alpha^{**}}}{\frac{dU^I(\cdot)}{dY^A}\Big|_{Y=Y^A}}. \quad (62)$$

As previously shown, $dU^I(\cdot)/dY|_{Y=Y^A} > 0$ (see Proof of Proposition 1). Moreover, observe that $dU^I(\cdot)/d\alpha^{**} = 1 > 0$. Consequently, $dY^A/d\alpha^{**} < 0$. Next, recall from Proposition 2 that $Y^A(\beta_0, \lambda) \leq Y^P(\beta_0, \lambda)$ for $\lambda \leq \lambda^{**}$. Because $dY^A/d\alpha^{**} < 0$, it would be optimal to set $\alpha^{**}(\lambda) \leq 0$ to distract the agent from innovation. However, since this would violate the agent's liability limit, the principal sets $\alpha^{**}(\lambda) = 0$ for $\lambda \leq \lambda^{**}$. Moreover, recall from Proposition 2 that $Y^A(\beta_0, \lambda) > Y^P(\beta_0, \lambda)$ for $\lambda > \lambda^{**}$. Therefore, to encourage more innovation, it is optimal to set $\alpha^{**}(\lambda) > 0$ for all $\lambda^{**} < \lambda < 1$. For $\lambda = 1$, note that the principal's objective function is equivalent to (3) in section 3 because $e_I^*(\lambda = 1) = 0$. Hence, according to Lemma 1, $\alpha^{**}(0) = 0$.

Finally, it remains to identify the relationship between $\alpha^{**}(\lambda)$ and λ for $\lambda^{**} < \lambda < 1$. To do so, define

$$F_{\alpha^{**}} \equiv U^I(\cdot) - U^S(\cdot) \quad (63)$$

$$G_{\alpha^{**}} \equiv -[\Pi^I - \Pi^S] \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\alpha^{**}} - 1. \quad (64)$$

Applying Cramer's rule yields

$$\frac{d\alpha^{**}(\lambda)}{d\lambda} = \frac{\det(A)}{\det(B)}, \quad (65)$$

where

$$A \equiv \begin{pmatrix} -\frac{\partial F_{\alpha^{**}}}{\partial \lambda} & \frac{\partial F_{\alpha^{**}}}{\partial Y^A} \\ -\frac{\partial G_{\alpha^{**}}}{\partial \lambda} & \frac{\partial G_{\alpha^{**}}}{\partial Y^A} \end{pmatrix} \quad B \equiv \begin{pmatrix} \frac{\partial F_{\alpha^{**}}}{\partial \alpha^{**}} & \frac{\partial F_{\alpha^{**}}}{\partial Y^A} \\ \frac{\partial G_{\alpha^{**}}}{\partial \alpha^{**}} & \frac{\partial G_{\alpha^{**}}}{\partial Y^A} \end{pmatrix}.$$

Since $\Pi(\cdot)$ is concave, B must be negative definite. Hence, $\det(B) > 0$. Therefore, $d\alpha^{**}(\lambda)/d\lambda > 0$ if

$$\det(A) = -\frac{\partial F_{\alpha^{**}}}{\partial \lambda} \frac{\partial G_{\alpha^{**}}}{\partial Y^A} + \frac{\partial G_{\alpha^{**}}}{\partial \lambda} \frac{\partial F_{\alpha^{**}}}{\partial Y^A} > 0. \quad (66)$$

Now observe that $\partial F_{\alpha^{**}}/\partial \lambda = \partial F_{\beta^*}/\partial \lambda$ and $\partial F_{\alpha^{**}}/\partial Y^A = \partial F_{\beta^*}/\partial Y^A$, see Proof of Proposition 3.

Moreover,

$$\frac{\partial G_{\alpha^{**}}}{\partial Y^A} = -[\Pi^I - \Pi^S] \frac{\partial}{\partial Y^A} \left(\frac{\phi(Y^A)}{\Phi(Y^A)} \right) \frac{dY^A}{d\alpha^{**}} \quad (67)$$

$$\frac{\partial G_{\alpha^{**}}}{\partial \lambda} = -\frac{\partial \Pi^I}{\partial \lambda} \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\alpha^{**}}. \quad (68)$$

Hence, $\text{sign}(\partial G_{\alpha^{**}}/\partial Y^A) = -\text{sign}(\partial G_{\beta^*}/\partial Y^A)$ and $\text{sign}(\partial G_{\alpha^{**}}/\partial \lambda) = -\text{sign}(\partial G_{\beta^*}/\partial \lambda)$. By combining the previous observations, we can infer that

$$\text{sign} \left(\frac{d\alpha^{**}(\lambda)}{d\lambda} \right) = -\text{sign} \left(\frac{d\beta^*(\lambda)}{d\lambda} \right) \quad \forall \lambda^{**} < \lambda < 1. \quad (69)$$

Therefore, we can infer from Proposition 3 that $d\alpha^{**}(\lambda)/d\lambda > 0$ for $\lambda^{**} < \lambda < \lambda^*$, and $d\alpha^{**}(\lambda)/d\lambda < 0$ for $\lambda^* < \lambda < 1$. \square

Proof of Proposition 6.

First, we demonstrate in general how the menu of contracts affects the agent's specialization decision. Recall that $dY^A/d\alpha^{**} < 0$ and $dY^A/d\beta^* > 0$. Since $\beta^* > \beta_0$ for $0 < \lambda < \lambda^{**}$, we can infer that $Y^A(\Omega^*(\lambda), \lambda) > Y^A(\beta_0, \lambda)$ for $\lambda < \lambda^{**}$. Moreover, since $\alpha^{**} > \alpha_0 = 0$ and $\beta^* < \beta_0$ for $\lambda > \lambda^{**}$, it must hold that $Y^A(\Omega^*(\lambda), \lambda) < Y^A(\beta_0, \lambda)$ for $\lambda > \lambda^{**}$, independent of which contract the agent chooses.

Next, we show in general how the menu of contracts affects the principal's preference Y^P . To do so, we first need to identify how Y^P responds to a change of α^{**} and β^* respectively. Recall that $Y^P(\alpha, \beta, \lambda)$ is implicitly characterized by $\Pi^S(e_S^*, \beta) = \Pi^I(e_I^*, \alpha, \lambda, Y^P)$. We first implicitly differentiate Y^P with respect to β and get

$$\frac{dY^P}{d\beta} = \frac{\frac{d\Pi^S(\cdot)}{d\beta}}{\frac{d\Pi^I(\cdot)}{dY} \Big|_{Y=Y^P}}. \quad (70)$$

Since $\Pi^S(\cdot)$ is concave in β with its maximum at $\beta = \beta_0$, we have $d\Pi^S(\cdot)/d\beta > 0$ for $\beta < \beta_0$, and $d\Pi^S(\cdot)/d\beta < 0$ for $\beta > \beta_0$. Using our results from Proposition 3, we can infer that $dY^P/d\beta < 0$ for $\lambda < \lambda^{**}$, and $dY^P/d\beta > 0$ for $\lambda > \lambda^{**}$. Since $\beta^* > \beta_0$ for $0 < \lambda < \lambda^{**}$, it must hold that $Y^P(\beta^*, \lambda) < Y^P(\beta_0, \lambda)$. For $\lambda^{**} < \lambda < 1$ however, we have $\beta^* < \beta_0$. Consequently, $Y^P(\beta^*, \lambda) < Y^P(\beta_0, \lambda)$ for $\lambda^{**} < \lambda < 1$, provided the agent has chosen the contract $\{0, \beta^*\}$. Implicit differentiating Y^P with respect to α yields

$$\frac{dY^P}{d\alpha} = -\frac{\frac{d\Pi^I(\cdot)}{d\alpha}}{\frac{d\Pi^I(\cdot)}{dY} \Big|_{Y=Y^P}}. \quad (71)$$

As previously shown, $d\Pi^I(\cdot)/dY|_{Y=Y^P} > 0$ (see Proof of Lemma 4). Furthermore, observe that $d\Pi^I(\cdot)/d\alpha = -1 < 0$. Consequently, $dY^P/d\alpha > 0$. Since $\alpha^{**} > 0$ for $\lambda^{**} < \lambda < 1$, this would imply that $Y^P(\alpha^{**}, \beta_0, \lambda) > Y^P(0, \beta_0, \lambda)$ at $\beta = \beta_0$. Note however, that the contract $\{\alpha^{**}, 0\}$ consists of a zero incentive payment, which—according to the previous observations—would lower Y^P . Thus, the net effect of $\{\alpha^{**}, 0\}$ on Y^P is ambiguous.

Finally note that the principal can always find a menu of contracts, denoted by $\Omega'(\lambda)$, which ensures perfect alignment of preferences, i.e., $Y^A(\Omega'(\lambda), \lambda) = Y^P(\Omega'(\lambda), \lambda)$. However, since the required contract adjustments would sacrifice the efficiency of standard task (adjusting β^*) and innovation (adjusting α^{**}), it cannot be optimal for the principal to do so. Thus, $Y^A(\Omega^*(\lambda), \lambda) \neq Y^P(\Omega^*(\lambda), \lambda)$ for all $\lambda \neq \lambda^{**}$. \square

Proof of Proposition 7.

First, recall that the first-best social benchmark Y^{FB} is implicitly characterized by

$$\underbrace{V(e_S^{FB}) - c(e_S^{FB})}_{\equiv S^S} = \underbrace{Y^{FB} \mu e_I^{FB} - c(e_I^{FB})}_{\equiv S^I}. \quad (72)$$

Now suppose that $\lambda \rightarrow 1$. As shown, $Y^A(\cdot), \hat{Y}(\cdot), Y^P(\cdot) \rightarrow \infty$ for $\lambda \rightarrow 1$, where $Y^P(\cdot) < \hat{Y}(\cdot) < Y^A(\cdot)$. Since $Y^{FB} \in (0, \infty)$, it follows that $Y^{FB} < Y^P(\cdot) < \hat{Y}(\cdot) < Y^A(\cdot)$. Next, suppose that $\lambda \rightarrow 0$. As demonstrated, $Y^P(\cdot), \hat{Y}(\cdot) \rightarrow \infty$. Hence, it remains to compare Y^{FB} and $Y^A(\cdot)$ for $\lambda \rightarrow 0$. Note first that $e_I^* = e_I^{FB}$ for $\lambda = 0$. Therefore, $U^I = S^I$ for $\lambda = 0$. Moreover, the first-best surplus S^S constitutes the upper bound for all potential surpluses from standard tasks. As well known, $S^S = U^S$ can only be achieved in a principal-agent relationship if the agent is the residual claimant and thus implements e_S^{FB} . This however, would require that $\alpha^* < 0$, which violates the agent's liability limit. Because under any second-best incentive scheme with $\alpha^* \geq 0$ the agent implements $e_S^* \neq e_S^{FB}$ and total surplus is split between principal and agent, it must hold that $S^S > U^S$. Hence, we can infer from (73) and (72) that $Y^{FB} > Y^A(\cdot)$ at $\lambda = 0$. In sum, this implies that $Y^A(\cdot) < \hat{Y}^{FB}$ for sufficiently low λ , and $Y^A(\cdot) > \hat{Y}^{FB}$ otherwise.

Next, we investigate the second-best benchmark $\hat{Y}^{SB}(\cdot)$. To do so, recall that $Y^A(\cdot)$ and $Y^P(\cdot)$ are implicitly characterized by

$$\underbrace{\beta^* e_S^* - c(e_S^*)}_{\equiv U^S} = \underbrace{(1 - \lambda) Y^A \mu e_I^* - c(e_I^*)}_{\equiv U^I}, \quad (73)$$

$$\underbrace{V(e_S^*) - \beta^* e_S^*}_{\equiv \Pi^S} = \underbrace{\lambda Y^P \mu e_I^*}_{\equiv \Pi^I}. \quad (74)$$

For $\lambda = \lambda^{**}$, we have $Y^A(\cdot) = Y^P(\cdot)$ since $\beta^* = \beta_0$, see Propositions 2 and 3. Therefore, we can add (73) and (74) and obtain

$$V(e_S^*) - c(e_S^*) = \widehat{Y}^{SB} \mu e_I^*. \quad (75)$$

Note that this is the condition which implicitly characterizes $\widehat{Y}^{SB}(\cdot)$. Consequently, $\widehat{Y}^{SB}(\cdot) = Y^A(\cdot) = Y^P(\cdot)$ for $\lambda = \lambda^{**}$. Next, we characterize $\widehat{Y}^{SB}(\cdot)$ for the extreme values of λ : $\lambda = 0, 1$. First suppose that $\lambda \rightarrow 0$. For an arbitrary Y , the rhs of (75) is identical to the rhs of (73). Note however, that under any second-best incentive scheme it must hold that $S^S > U^S$. Hence, the lhs of (75) must be strictly greater than the lhs of (73). This in turn implies that $\widehat{Y}^{SB}(\cdot) > Y^A(\cdot)$ for $\lambda \rightarrow 0$. Moreover, since $\widehat{Y}^{SB}(\cdot) \in (0, \infty)$ and $Y^P(\cdot) \rightarrow \infty$ for $\lambda \rightarrow 0$, we have $\widehat{Y}^{SB}(\cdot) < Y^P(\cdot)$ for $\lambda \rightarrow 0$. Now suppose that $\lambda \rightarrow 1$. Because in this case $e_I^* \rightarrow 0$, the rhs of (75) goes to zero for an arbitrary Y , while the lhs remains strictly positive. This in turn requires that $\widehat{Y}^{SB}(\cdot) \rightarrow \infty$ for $\lambda \rightarrow 1$. Furthermore, note that for $\lambda \rightarrow 1$, the rhs of both (75) and (73) go to zero. Nevertheless, the lhs of (75) is strictly greater than the lhs of (74). Thus, for $\lambda \rightarrow 1$, we have $\widehat{Y}^{SB}(\cdot) > Y^P(\cdot)$. Finally, because $\widehat{Y}^{SB}(\cdot)$ is a weighted average of $Y^P(\cdot)$ and $Y^A(\cdot)$, we can infer that $\widehat{Y}^{SB}(\cdot) < Y^A(\cdot)$ for $\lambda \rightarrow 1$. Combing our previous observations, we have $Y^A(\cdot) < \widehat{Y}^{SB}(\cdot)$ for $\lambda < \lambda^{**}$, and $Y^A(\cdot) > \widehat{Y}^{SB}(\cdot)$ for $\lambda > \lambda^{**}$. \square

Proof of Lemma 5.

Recall that Y^A is implicitly defined by $U^S(e_S^*, \beta, \varepsilon) = U^I(e_I^*, \lambda, Y^A, \varepsilon)$. Implicit differentiating Y^A with respect to ε yields

$$\frac{dY^A}{d\varepsilon} = \frac{\mu e_I^* - e_S^*}{(1 - \lambda)\mu e_I^*}. \quad (76)$$

Thus, $dY^A/d\varepsilon < 0$ if $\mu e_I^* < e_S^*$, which is equivalent to $\text{Prob}[Y > 0|e_I^*] < \text{Prob}[S = 1|e_S^*]$. Recall that $de_I^*/d\mu > 0$. Hence, the fraction on the rhs of (76) is negative for sufficiently low values of μ , and positive for sufficiently high values of μ . Moreover, one can show that

$$\frac{d}{d\mu} \left(\frac{\mu e_I^* - e_S^*}{(1 - \lambda)\mu e_I^*} \right) = \frac{e_S^*(1 - \lambda) \left[e_I^* + \mu \frac{de_I^*}{d\mu} \right]}{[(1 - \lambda)\mu e_I^*]^2} > 0 \quad (77)$$

Accordingly, the rhs of (76) is increasing in μ . Consequently, there exists a threshold $\mu^* > 0$ satisfying $dY^A/d\varepsilon|_{\mu=\mu^*} = 0$ such that $dY^A/d\varepsilon < 0$ for $\mu < \mu^*$, and $dY^A/d\varepsilon \geq 0$ otherwise. \square

Proof of Proposition 8.

First, recall from Lemma 5 that $dY^A/d\varepsilon < 0$ if $\mu < \mu^*$. Moreover, Proposition 6 implies that $Y^P(\cdot) > Y^A(\cdot)$ for $\lambda < \lambda^{**}$. Thus, offering the agent $\varepsilon > 0$ cannot be optimal for $\lambda < \lambda^{**}$.

Next, we demonstrate that tolerance for failure is optimal for sufficiently high values of λ . Recall that offering the agent $\varepsilon > 0$ does not impose direct costs on the principal, but indirectly affects e_S^* , e_I^* , and Y^A . By applying the Envelope Theorem, one get

$$\frac{d\Pi(\cdot)}{d\varepsilon} = \frac{\partial\Pi(\cdot)}{\partial e_S^*} \frac{de_S^*}{d\varepsilon} + \frac{\partial\Pi(\cdot)}{\partial e_I^*} \frac{de_I^*}{d\varepsilon} + \frac{\partial\Pi(\cdot)}{\partial Y^A} \frac{dY^A}{d\varepsilon}. \quad (78)$$

To demonstrate that $d\Pi(\cdot)/d\varepsilon > 0$ for sufficiently high values of λ , suppose for a moment that $\lambda \rightarrow 1$. Because in this case $e_I^* \rightarrow 0$, (78) simplifies to

$$\frac{d\Pi(\cdot)}{d\varepsilon} = \underbrace{\frac{\partial\Pi(\cdot)}{\partial e_S^*} \frac{de_S^*}{d\varepsilon}}_{\equiv A} + \underbrace{\frac{\partial\Pi(\cdot)}{\partial Y^A} \frac{dY^A}{d\varepsilon}}_{\equiv B}. \quad (79)$$

Since $\partial\Pi(\cdot)/\partial e_S^* > 0$ and $de_S^*/d\varepsilon < 0$, A is strictly negative and clearly finite. Moreover, Proposition 6 implies that $d\Pi(\cdot)/dY^A < 0$ for $\lambda \rightarrow 1$. From Proof of Lemma 5 we know that

$$\frac{dY^A}{d\varepsilon} = \frac{\mu e_I^* - e_S^*}{(1 - \lambda)\mu e_I^*}. \quad (80)$$

Since $e_I^* \rightarrow 0$ for $\lambda \rightarrow 1$, it follows that $\lim_{\lambda \rightarrow 1} \frac{dY^A}{d\varepsilon} = -\infty$, implying that $B \rightarrow \infty$. Hence, $B > -A$ for $\lambda \rightarrow 1$. This in turn implies that $d\Pi(\cdot)/d\varepsilon > 0$ for sufficiently high values of λ . To summarize the previous observations, there exists a threshold $\hat{\lambda}$, with $1 > \hat{\lambda} \geq \lambda^{**}$, such that a tolerance policy is never optimal for $\lambda \leq \hat{\lambda}$, and optimal for $\lambda > \hat{\lambda}$. \square

Comparative Statics Investing in Employee Innovation.

First, consider the agent's innovation effort e_I^* which is implicitly characterized by

$$(1 - \lambda)Y\mu = c'(e_I). \quad (81)$$

Implicit differentiating e_I^* with respect to μ gives

$$\frac{de_I^*}{d\mu} = \frac{(1 - \lambda)Y}{c''(e_I)}, \quad (82)$$

which is strictly positive. Next, recall that Y^A is implicitly defined by $U^S(e_S^*, \beta) = U^I(e_I^*, \lambda, Y^A, \mu)$.

Implicit differentiating Y^A with respect to μ yields

$$\frac{dY^A}{d\mu} = -\frac{(1 - \lambda)Y^A e_I^*}{(1 - \lambda)\mu e_I^*} = -\frac{Y^A}{\mu}, \quad (83)$$

which is strictly negative. □

Proof of Proposition 9.

By endogenizing μ , the principal's expected profit can be written as

$$\Pi(e_S^*, e_I^*, \beta, Y^A, \mu) = \int_0^{Y^A} \Pi^S d\Phi(Y) + \int_{Y^A}^{\infty} \Pi^I d\Phi(Y) - c(\mu). \quad (84)$$

The optimal investment level $\mu^*(\lambda)$ is implicitly characterized by the following first-order-condition:

$$\frac{d\Pi^I}{d\mu} + [\Pi^S - \Pi^I] \frac{\phi(Y^A)}{\Phi(Y^A)} \frac{dY^A}{d\mu} = \frac{c'(\mu)}{\Phi(Y^A)}. \quad (85)$$

Notice that the first term on the lhs captures the incentive effect, and the second term the specialization effect. Suppose for a moment that either $\lambda = 0$ or $\lambda = 1$. Recall that in both cases $\Pi^I = 0$. Hence, (85) simplifies to

$$\Pi^S \phi(Y^A) \frac{dY^A}{d\mu} = c'(\mu). \quad (86)$$

Since $dY^A/d\mu < 0$ (see Appendix: Comparative Statics Investing in Employee Innovation), the lhs is strictly negative. Hence, the principal sets $\mu^*(\lambda) = \underline{\mu}$ for $\lambda = 0, 1$. Next, consider the interval $\lambda^{**} < \lambda < 1$. Proposition 2 implies that the second term on the lhs of (85) is strictly positive. Moreover, it can be easily shown that $d\Pi^I/d\mu > 0$. Consequently, $\mu^*(\lambda) > \underline{\mu}$ for $\lambda^{**} < \lambda < 1$. Now consider the threshold λ^{**} . Recall from Proposition 2 that $Y^A = Y^P$ for $\lambda = \lambda^{**}$. This implies that under the optimal menu of contracts $\Omega^*(\lambda^{**})$, $\Pi^S = \Pi^I$. Therefore, (85) simplifies to

$$\frac{d\Pi^I}{d\mu} \Phi(Y^A) = c'(\mu). \quad (87)$$

Because the lhs is strictly positive, we have $\mu^*(\lambda^{**}) > \underline{\mu}$. Finally, consider the interval $0 < \lambda < \lambda^{**}$. Note that $d\Pi^I/d\mu > 0$. Moreover, Proposition 6, in conjunction with Proposition 1 and Lemma 4, implies that the second term on the lhs of (85) is strictly negative, and its absolute value is diminishing in λ . Since $\mu^*(0) = \underline{\mu}$ and $\mu^*(\lambda^{**}) > \underline{\mu}$, we can thus infer that there exists a threshold λ^{***} such that $\mu^*(\lambda) = \underline{\mu}$ for all $\lambda \leq \lambda^{***}$, and $\mu^*(\lambda) > \underline{\mu}$ for all $\lambda^{***} < \lambda < 1$. □

Proof of Lemma 6.

First, define

$$F \equiv \beta\rho'(e_S) - c'(e), \quad (88)$$

$$G \equiv (1 - \lambda)\mu'(e_I) - c'(e). \quad (89)$$

Totally differentiating yields

$$\begin{pmatrix} de_S \\ de_I \end{pmatrix} = \frac{1}{-\det(\Delta)} \begin{pmatrix} (1 - \lambda)\mu''(e_I) - c''(e) & c''(e) \\ c''(e) & \beta\rho''(e_S) - c''(e) \end{pmatrix} \begin{pmatrix} \rho'(e_S)d\beta \\ -\mu'(e_I)d\lambda \end{pmatrix},$$

where

$$\Delta \equiv \begin{pmatrix} \frac{\partial F}{\partial e_S} & \frac{\partial F}{\partial e_I} \\ \frac{\partial G}{\partial e_S} & \frac{\partial G}{\partial e_I} \end{pmatrix}. \quad (90)$$

Since $U(e_S, e_I)$ is concave in e_S and e_I , Δ must be negative definite. Hence, $\det(\Delta) > 0$. Thus, we obtain

$$\begin{aligned} \frac{de_S^*}{d\beta} &= -\frac{1}{\det(\Delta)}\rho'(e_S) \underbrace{[(1 - \lambda)Y\mu''(e_I) - c''(e)]}_{<0 \text{ due to s.o.c.}} > 0 \\ \frac{de_S^*}{d\lambda} &= \frac{1}{\det(\Delta)}c''(e)Y\mu'(e_I) > 0 \\ \frac{de_I^*}{d\beta} &= -\frac{1}{\det(\Delta)}\rho'(e_S)c''(e) < 0 \\ \frac{de_I^*}{d\lambda} &= \frac{1}{\det(\Delta)}Y\mu'(e_I) \underbrace{[\beta\rho''(e_S) - c''(e)]}_{<0 \text{ due to s.o.c.}} < 0 \\ \frac{de^*}{d\beta} &= \frac{de_S^*}{d\beta} + \frac{de_I^*}{d\beta} = -\frac{1}{\det(\Delta)}[\rho'(e_S)(1 - \lambda)Y\mu''(e_I)] > 0 \\ \frac{de^*}{d\lambda} &= \frac{de_S^*}{d\lambda} + \frac{de_I^*}{d\lambda} = \frac{1}{\det(\Delta)}[Y\mu'(e_I)\beta\rho''(e_S)] < 0. \end{aligned}$$

Note that the signs of the latter two derivatives follow from $\mu''(e_I) < 0$ and $\rho''(e_S) < 0$. □

Proof of Lemma 7.

Recall from Proposition 6 that $de_I^*/d\lambda < 0$. Moreover, one can show that $de_I^P/d\lambda > 0$, where $e_I^P(\lambda = 0) = 0$. Thus, e_I^* is decreasing, while e_I^P is increasing in λ . Hence, there exists a threshold λ^{**} satisfying

$e_I^*(\lambda^{**}) = e_I^P(\lambda^{**})$, such that $e_I^* \geq e_I^P$ for $\lambda \leq \lambda^{**}$, and $e_I^* < e_I^P$ otherwise. \square

Proof of Proposition 10.

First, we characterize $\beta^*(\lambda)$ for $\lambda = 1$. Since $e_I^*(\lambda = 1) = 0$, the first-order condition (28) simplifies to $d\Pi^S/d\beta|_{\lambda=1} = 0$. Hence, $\beta^*(1) = \beta_0$. Next, we demonstrate that $\beta^*(0) > \beta^*(1) = \beta_0$. To do so, we need to express the first-order condition (29) for $\lambda = 0$ and $\lambda = 1$:

$$\left. \frac{d\Pi^S}{d\beta} \right|_{\lambda=0} = \underbrace{[V'(e_S^*(\beta_0, 0)) - \beta_0 \rho'(e_S^*(\beta_0, 0))]}_{\equiv A1} \left. \frac{de_S^*}{d\beta} \right|_{\lambda=0} - \underbrace{\rho(e_S^*(\beta_0, 0))}_{\equiv B1} = 0 \quad (91)$$

$$\left. \frac{d\Pi^S}{d\beta} \right|_{\lambda=1} = [V'(e_S^*(\beta^*, 1)) - \beta^* \rho'(e_S^*(\beta^*, 1))] \left. \frac{de_S^*}{d\beta} \right|_{\lambda=1} - \rho(e_S^*(\beta^*, 1)) = 0. \quad (92)$$

To prove that $\beta^*(0) > \beta^*(1) = \beta_0$, we need to verify that at $\beta = \beta_0$,

$$\left. \frac{d\Pi^S}{d\beta} \right|_{\lambda=1, \beta=\beta_0} = \underbrace{[V'(e_S^*(\beta_0, 1)) - \beta_0 \rho'(e_S^*(\beta_0, 1))]}_{\equiv A2} \left. \frac{de_S^*}{d\beta} \right|_{\lambda=1, \beta=\beta_0} - \underbrace{\rho(e_S^*(\beta_0, 1))}_{\equiv B2} > 0. \quad (93)$$

We can infer from Proposition 6 that $e_S^*(\beta_0, \lambda = 0) > e_S^*(\beta_0, \lambda = 1)$. Hence, $B1 > B2$. Because $de_S^*/d\beta > 0$ and $de_I^*/d\beta < 0$ (see Proposition 6), we can deduce from (29) that $V'(e_S^*(\beta_0, 0)) - \beta_0 \rho'(e_S^*(\beta_0, 0)) > 0$. Moreover, since $e_S^*(\beta_0, \lambda = 0) > e_S^*(\beta_0, \lambda = 1)$, we can conclude that $A1 < A2$. Finally, it remains to demonstrate that

$$\left. \frac{de_S^*}{d\beta} \right|_{\lambda=0} < \left. \frac{de_S^*}{d\beta} \right|_{\lambda=1}. \quad (94)$$

It can be shown that e_S^* is concave increasing in β . From concavity of e_S^* in β and $e_S^*(\beta_0, \lambda = 0) > e_S^*(\beta_0, \lambda = 1)$, it follows that (94) must hold. Consequently, $d\Pi^S/d\beta|_{\lambda=1, \beta=\beta_0} > 0$, which implies that $\beta^*(0) > \beta^*(1) = \beta_0$. Next, we demonstrate that $\beta^*(\lambda) < \beta_0$ for sufficiently high values of λ , with $\lambda < 1$.

To do so, we express the first-order condition (28) in a more general fashion:

$$\frac{d\Pi}{d\beta} = \underbrace{\frac{\partial \Pi^S}{\partial e_S^*} \frac{\partial e_S^*}{\partial \beta}}_{\equiv A} + \frac{\partial \Pi^S}{\partial \beta} + \underbrace{\frac{\partial \Pi^I}{\partial e_I^*} \frac{\partial e_I^*}{\partial \beta}}_{\equiv B}. \quad (95)$$

First, observe that $\partial \Pi / \partial e_I^* > 0$, and recall from Proposition 6 that $\partial e_I^* / \partial \beta < 0$. Thus, (95) implies that $A > 0$. Now suppose for a moment that $\lambda \rightarrow 1$. Then, $e_I^* \rightarrow 0$, and as a result $B \rightarrow 0$. Hence, $d\Pi/d\beta|_{\beta=\beta^*(\lambda), \lambda \rightarrow 1} > 0$, implying that $\beta^*(\lambda) < \beta_0$ for $\lambda \rightarrow 1$. To summarize, $\beta^*(\lambda) > \beta_0$ for sufficiently low λ , and $\beta^*(\lambda) < \beta_0$ for sufficiently high λ . \square

Proof of Lemma 8.

Let $e_S^*(\beta, m)$ and $e_I^*(\lambda, Y, m)$ denote the agent's optimal effort levels for a given incentive bonus β and a given number of stock options m . The agent's expected utility then becomes

$$U(\cdot) = \int_0^{Y^A} [\beta e_S^* + m\varphi(\cdot, e_S^*, 0) - c(e_S^*)] d\Phi(Y) \quad (96)$$

$$+ \int_{Y^A}^{\infty} [(1 - \lambda)Y\mu e_I^* + m\varphi(\cdot, 0, e_I^*) - c(e_I^*)] d\Phi(Y). \quad (97)$$

From the first-order condition with respect to Y^A , the agent's threshold innovation value $Y^A(\beta, \lambda, m)$ is now implicitly characterized by

$$\beta e_S^* + m\varphi(\cdot, e_S^*, 0) - c(e_S^*) = (1 - \lambda)Y^A\mu e_I^* + m\varphi(\cdot, 0, e_I^*) - c(e_I^*), \quad (98)$$

which is equivalent to $U^S(e_S^*, \beta, m) = U^I(e_I^*, \lambda, m, Y^A)$. Implicit differentiating Y^A with respect to m yields

$$\frac{dY^A}{dm} = \frac{\frac{d}{dm} [U^S(\cdot) - U^I(\cdot)]}{\left. \frac{dU^I(\cdot)}{dY} \right|_{Y=Y^A}}. \quad (99)$$

We know from Proposition 1 that $dU^I(\cdot)/dY|_{Y=Y^A} > 0$. Moreover, applying the Envelope Theorem and letting $m \rightarrow 0$ yield

$$\left. \frac{d}{dm} [U^S(\cdot) - U^I(\cdot)] \right|_{m=0} = \varphi(\cdot, e_S^*, 0) - \varphi(\cdot, 0, e_I^*). \quad (100)$$

Consider first the interval $\lambda \in (0, \lambda^{**}]$. For any $\lambda \in (0, \lambda^{**}]$ and $Y \in (Y^A(\cdot), Y^P(\cdot)]$, we can infer from Proposition 2 that $\Pi^A(\cdot, e_S^*, 0) \geq \Pi^A(\cdot, 0, e_I^*)$. Since $\partial\varphi(\cdot)/\partial\Pi_A > 0$, this implies that $\varphi(\cdot, e_S^*, 0) \geq \varphi(\cdot, 0, e_I^*)$ for $m > 0$, with $m \rightarrow 0$, and $\lambda \leq \lambda^{**}$. Thus, $dY^A/dm|_{m=0} \geq 0$ for $\lambda \leq \lambda^{**}$. Likewise, one can show that $dY^A/dm|_{m=0} < 0$ for $\lambda > \lambda^{**}$. \square

Incentive Effect of Stock Options.

Let $\Theta(\cdot) \equiv m\varphi(\cdot)$ denote the agent's expected compensation from stock options. First, suppose that the agent focuses exclusively on the standard task. For a given number of stock options m , the first-order condition characterizing the agent's optimal effort level $e_S^*(\beta, m)$ for the standard task becomes

$$\beta + \frac{\partial\Theta(\cdot)}{\partial e^S} = c'(e_S). \quad (101)$$

Implicit differentiating $e_S^*(\beta, m)$ with respect to m yields

$$\frac{de_S^*}{dm} = -\frac{\frac{\partial^2 \Theta(\cdot)}{\partial e_S \partial m}}{\frac{\partial}{\partial e_S} \left[\beta + \frac{\partial \Theta(\Pi|\cdot)}{\partial e^S} - c'(e_S) \right]} = -\frac{\frac{\partial \varphi(\cdot)}{\partial e_S}}{\frac{\partial}{\partial e_S} \left[\beta + \frac{\partial \Theta(\Pi|\cdot)}{\partial e^S} - c'(e_S) \right]}, \quad (102)$$

where the denominator is strictly negative due to the second-order condition. Moreover, notice that $\partial E[\Pi^A]/\partial e_S > 0$ for all $\beta \leq \beta^*$ and $m \leq m^*$, implying that $\partial \varphi(\cdot)/\partial e_S > 0$. Thus, $de_S^*/dm > 0$.

Next, suppose the agent focuses on innovation. For a given number of stock options m , the first-order condition, which characterizes the agent's optimal effort level $e_I^*(\lambda, Y, m)$, becomes

$$(1 - \lambda)Y\mu + \frac{\partial \Theta(\cdot)}{\partial e_I} = c'(e_I). \quad (103)$$

Implicit differentiating $e_I^*(\lambda, Y, m)$ with respect to m gives

$$\frac{de_I^*}{dm} = -\frac{\frac{\partial^2 \Theta(\cdot)}{\partial e_I \partial m}}{\frac{d}{dm} \left[(1 - \lambda)Y\mu + \frac{\partial \Theta(\cdot)}{\partial e_I} - c'(e_I) \right]} = -\frac{\frac{\partial \varphi(\cdot)}{\partial e_I}}{\frac{d}{dm} \left[(1 - \lambda)Y\mu + \frac{\partial \Theta(\cdot)}{\partial e_I} - c'(e_I) \right]}, \quad (104)$$

where the denominator is strictly negative due to the second-order condition. Furthermore, $\partial E[\Pi^A]/\partial e_I > 0$ for all $m \leq m^*$ and $\lambda > 0$, implying that $\partial \varphi(\cdot)/\partial e_I > 0$. As a result, $de_I^*/dm > 0$. \square

Proof of Proposition 11.

To show that $m^* = 0$ for $\sigma \rightarrow \infty$, it is sufficient to demonstrate that for $\sigma \rightarrow \infty$ with $m \rightarrow 0$, (i) the marginal costs of providing the agent with stock options go to infinity; and (ii), the marginal benefit goes to zero. In this case, two essential Inada conditions are violated, which in turn leads to a corner solution with $m^* = 0$. First, one can easily verify that

$$\lim_{\sigma \rightarrow \infty} \frac{\partial C(\cdot)}{\partial m} \Big|_{m=0} = \lim_{\sigma \rightarrow \infty} \varphi(\cdot) = \infty. \quad (105)$$

Next, consider the principal's marginal benefit—denoted B_m —of providing the agent with stock options:

$$B_m = [\Pi^S - \Pi^I] \phi(Y^A) \frac{dY^A}{dm}. \quad (106)$$

We can infer from (100) (see Proof of Lemma 8) that

$$\frac{dY^A}{dm} \Big|_{m=0} = \frac{\varphi(\cdot, e_S^*, 0) - \varphi(\cdot, 0, e_I^*)}{\frac{dU^I(\cdot)}{dY} \Big|_{Y=Y^A}}, \quad (107)$$

where $dU^I(\cdot)/dY|_{Y=Y^A} > 0$ (see Proof of Proposition 1). Because $\lim_{\sigma \rightarrow \infty} \varphi(\cdot) = \infty$, it follows that

$$\lim_{\sigma \rightarrow \infty} \left. \frac{dY^A}{dm} \right|_{m=0} = 0, \quad (108)$$

and hence, $\lim_{\sigma \rightarrow \infty} B_m|_{m=0} = 0$. Thus, $m^* = 0$ for $\sigma \rightarrow \infty$. \square

References

- Aghion, P. and J. Tirole (1994). The Management of Innovation. *Quarterly Journal of Economics*, 109, 1185-1209.
- Anand, B. and A. Galetovic (2000). Weak Property Rights and Hold-up in R&D. *Journal of Economics, Management, & Strategy*, 9(4), 615-642.
- Anton, J. and D. Yao (1994). Expropriations and Inventions: Appropriable Rents in the Absence of Property Rights. *American Economic Review*, 84, 190-209.
- Anton, J. and D. Yao (1995). Start-ups, Spin-offs, and Internal Projects. *Journal of Law, Economics & Organization*, 11(2), October, 362-378.
- Baker, G. P. (2002). Distortion and Risk in Optimal Incentive Contracts. *Journal of Human Resources*, 37(4), 728-751.
- Bernardo, A., H. Cai and J. Luo (2009). Motivating Entrepreneurial Activity in a Firm. *Review of Financial Studies*, 22(2), 1089-1118.
- Burgelman, R. (2002). Strategy as Vector and the Inertia of Coevolutionary Lock-in. *Administrative Science Quarterly*, 47(2), 325-357.
- de Bettignies, J. and G. Chemla (2008), Corporate Venturing, Allocation of Talent, and Competition for Star Managers. *Management Science*. 54(3), 505-522.
- Demougin, D. and C. Fluet (2001). Monitoring versus Incentives. *European Economic Review*, 45(9), 1741-1764.
- Feltham, G. A. and J. Xie (1994). Performance Measure Congruity and Diversity in Multi-Task Principal/Agent Relations. *Accounting Review*, 69(3), 429-453.
- Gambardella A. and C. Panico (2007). Optimal Governance Mechanism for Knowledge-Intensive Activities. Mimeo, Bocconi University.
- Gans, J. and S. Stern (2000). Incumbency and R&D Incentives: Licensing in the Gale of Creative Destruction. *Journal of Economics, & Management Strategy*, 9(4), 485-511.
- Gans, J. and S. Stern (2003). The Product Market and the 'Market for Ideas': Commercialization Strategies for Technology Entrepreneurs. *Research Policy*, 32(2), 333-350.
- Gibbons, R. (1998). Incentives in Organizations. *Journal of Economic Perspectives*, 12, 115-32.
- Grossman, S. and O. Hart (1986). The Costs and Benefits of Ownership: A Theory of Lateral and Vertical Integration. *Journal of Political Economy*, 94(4), 691-719.
- Groysberg, B., A. Nanda, and M. J. Prats (2007). Does Individual Performance Affect Entrepreneurial Mobility? Empirical Evidence from the Financial Analysis Market. NBER Working Paper Series, No. 13633.

- Inderst, R. and M. Klein (2007). Innovation, Endogenous Overinvestment, and Incentive Pay. *Rand Journal of Economics*, 38(4), 881-904.
- Hamel, G. (1999). Bringing Silicon Valley Inside. *Harvard Business Review*, 77(5), 71-84.
- Hart, O. (1995). *Firms, Contracts and Financial Structure*. Clarendon Press. Oxford.
- Hart, O. and J. Moore (1990). Property Rights and the Nature of the Firm. *Journal of Political Economy*, 98(6), 1119-1158.
- Hart, O. and J. Moore (1999). Foundations of Incomplete Contracts. *Review of Economic Studies*, 66(1), 115-138.
- Hart, O. and J. Moore (2008). Contracts as Reference Points. *Quarterly Journal of Economics*, 123(1), 1-48.
- Hellmann, T. (2007). When Do Employees Become Entrepreneurs? *Management Science*, 53(6), 919-933.
- Hellmann, T. and E. Perotti (2007). The Circulation of Ideas in Firms and Markets. Mimeo, University of British Columbia.
- Holmström, B. (1979). Moral Hazard and Observability. *Bell Journal of Economics*, 10(1), 74-91.
- Holmström, B. (1999). The Firm as a Subeconomy. *Journal of Law, Economics, & Organization*, 15(1), 74-102.
- Holmström, B. and P. Milgrom (1991). Multitask Principal-Agent Analyses: Incentives Contracts, Asset Ownership, and Job Design. *Journal of Law, Economics, & Organization*, 7, 24-52.
- Holmström, B. and P. Milgrom (1994). The Firm as an Incentive System. *American Economic Review*, 84(4), 972-991.
- Holmström, B. and J. Roberts (1998). The Boundaries of the Firm Revisited. *Journal of Economic Perspectives*, 12 (4), 73-94.
- IPOA (2004). Employee Inventor Compensation Practices Survey. Intellectual Property Owners Association, www.ipoa.org.
- Klepper, S. and S. Sleeper (2005). Entry by Spinoffs. *Management Science*, 51(8), 1291-1306.
- Kreps, D. (1992). Static Choice and Unforeseen Contingencies, in: P. Dasgupta, D. Gale, O. Hart, and E. Maskin (Eds), *Economic Analysis of Markets and Games*. MIT Press, Cambridge, 259-281.
- Levin, R. C., A. K. Klevorick, R. R. Nelson, and S. G. Winter (1987). Appropriating the Returns from Industrial Research and Development. *Brookings Papers on Economic Activity*, 18(3), 783-832.
- Manso, G. (2007). Motivating Innovation. Mimeo, MIT.
- March, J. G. (1991). Exploration and Exploitation. *Organization Science*, 2(1), February, 71-87.

- March, J. and J. Olsen (1991). Exploration and Exploitation in Organizational Learning. *Organizational Science*, 2(1), 71-78.
- Maskin, E. and J. Tirole (1999). Unforeseen Contingencies and Incomplete Contracts. *Review of Economic Studies*, 66(1), 83-114.
- Merges R. (1999). The Law and Economics of Employee Inventions. *Harvard Journal of Law and Technology*, 13(1), 1-63.
- Parent, D. (2002). Incentive Pay in the United States: Its Determinants and its Effects, in: M. Brown and J. Heywood (Eds.), *Paying for Performance: An International Comparison*, M.E. Sharpe, 17-51.
- Porter, M. (1980). *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. MacMillan, New York, New York.
- Rotemberg, J. and G. Saloner (1994). Benefits of Narrow Business Strategy. *American Economic Review*, 84(5), 1330-1349.
- Sappington, D. (1983). Limited Liability Contracts Between Principal and Agent. *Journal of Economic Theory*, 29, 1-21.
- Segal, I. (1999). Complexity and Renegotiation: A Foundation for Incomplete Contracts. *Review of Economic Studies*, 66(1), 57-82.
- Subramanian, N. (2005). The Economics of Intrapreneurial Innovation. *Journal of Economic Behavior and Organization*, 58(4), 487-510.
- Talley, E. (1998). Turning Servile Opportunities to Gold: A Strategic Analysis of the Corporate Opportunities Doctrine. *Yale Law Journal*, 108, 277-375.
- Tian, X. and T. Y. Wang (2009). Tolerance for Failure and Corporate Innovation. Working Paper.
- Williamson, O. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. The Free Press, New York.