European Economic Review 53 (2009) 871–887

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

Innovation management in organizations ☆

Roman Inderst

University of Frankfurt and LSE, Germany

ARTICLE INFO

Article history:
Received 31 October 2007
Accepted 28 April 2009
Available online 22 May 2009

JEL classification:
D23
J33

Keywords:
Incentive pay
Innovation
Delegation

ABSTRACT

This paper poses the question of how a firm should optimally choose both its organization and its compensation in the pursuit of innovation. One key result is that incentive pay arises as a robust instrument of innovation management both with and without delegation, although in the present model its primary purpose is not to elicit more effort for the creation of new ideas, but to ensure that new ideas are implemented if and only if this is efficient. While without delegation, the firm may "underinvest" in innovation, with delegation the opposite bias may arise as new ideas may be implemented too often ("overinvestment"). The optimal organizational choice trades off these two biases.

1. Introduction

Practitioners regard the delegation of decisions down the hierarchy of an organization as a powerful tool both to motivate employees and to ensure that decisions are made by those who possess the most accurate information. Economists are more sanguine about the benefits of delegation, as they rightly point out that the interests of the employee may not always coincide with those of the organization as a whole.

For instance, in the seminal paper by Aghion and Tirole (1997), the drawback of giving "real authority" to managers is that sometimes they will prefer underperforming "pet projects", as these generate *private* benefits.[1] In the present paper, instead, the conflict of interest between the organization and its employees arises endogenously out of employees' compensation.

Our model comprises the two stages of innovation management: the creation of new ideas and the subsequent decision of whether or not to implement them. Innovative organizations must clearly perform well along both dimensions, ensuring that new ideas are created and that they are subsequently implemented if and only if it is indeed efficient to do so.

We contrast two organizational forms. Under delegation, the decision of whether to implement a new idea is made by the employee (agent) who generated the idea in the first place. Without delegation, this decision is made in the firm's interest—e.g., by the employee's superior (principal). We find that both delegation and non-delegation may lead to biased decision making, though the "sign" of the respective bias is the opposite. Without delegation, there may be "underinvestment"; that is, innovations are implemented too infrequently. In contrast, with delegation,"overinvestment" may occur, as an innovation is sometimes implemented even though it has negative value. While delegation provides a

☆ This paper is a fully revised version of an earlier draft entitled "Incentives to Generate and Implement New Ideas," which focused on the design of the incentive scheme only.

E-mail address: inderst@finance.uni-frankfurt.de

[1] Exogenous differences in preferences are also assumed, as in, in Dessein (2002) or Alonso and Matouschek (2008), for example.

0014-2921/$ - see front matter © 2009 Published by Elsevier B.V.
doi:10.1016/j.euroecorev.2009.04.011

commitment for the firm to implement new ideas more often, this may not always result in overall more-efficient decision making.

Interestingly, high-powered incentive compensation, which links the agent's pay to the success of the innovation, becomes optimal both with and without delegation. The role of incentive pay is, however, not to create higher incentives at the stage where new ideas are generated. It serves, instead, to reduce the bias at the implementation stage. The use of incentive pay is, thus, "robust" with respect to the organizational choice since it serves the same purpose both with and without delegation. Thus, our model would lead to the somewhat counterintuitive prediction that greater authority need not imply steeper incentives. As discussed later in more detail, this conforms to findings in Wulf (2007) that more-senior managers do not necessarily have more-high-powered contracts.

In our model, the agent is incentivized to generate an innovation since he receives a financial reward upon its implementation. Paying a reward is credible for the firm, as implementing the agent's *own* innovation makes him more valuable. If the decision of whether to implement the innovation lies with the firm, then, according to a seminal insight from Rotemberg and Saloner (1994), the firm may forego marginally profitable innovations so as to save on the promised reward. If, instead, the agent himself can decide whether the innovation is implemented, then the opposite bias arises, leading to "overinvestment" instead of "underinvestment." In either regime, the use of incentive pay, which links the agent's reward to the future performance of the innovation, can mitigate (or sometimes even resolve) the respective bias. The reason for this result is as follows.

If the agent receives a lump-sum (bonus) payment, as is the case in Rotemberg and Saloner (1994), then the firm incurs the same *additional* wage costs irrespective of the profitability of the new idea. Without delegation, the firm will then always want to forego marginally profitable ideas. It is important to note that what matters for the agent's incentives to undertake effort is only the *expected* reward that he receives from a new idea. With incentive pay, the same *expected* level of the reward can be achieved while paying most of it in states where very profitable innovations are implemented. Instead, when a less profitable innovation is implemented, then the agent's reward will be also lower. For marginally profitable innovations, this can tilt the firm's decision towards implementation, thereby mitigating its "underinvestment" bias.

In the case of delegation, the mechanism by which incentive pay mitigates the bias (now of "overinvestment") is analogous. With incentive pay, the agent participates in the outcome of a newly implemented innovation. His incentives at the decision stage are, then, more aligned with those of the firm. More formally, the same level of the necessary *expected* reward can again be achieved in a way that leads to more-efficient decision making, now tilting the agent's decision *against* implementing some unprofitable innovations.

Our analysis of the optimal compensation scheme applies whether or not the firm has a choice to delegate decision making or not. In fact, the firm often may not be able to exert "real authority" over decisions—e.g., if the agent has valuable private information (cf. Aghion and Tirole, 1997). On the other hand, sometimes the firm may also find it difficult to credibly cede decision rights to an employee, in particular in the analyzed case of implementing new ideas that are, by their very nature, hard to describe *ex-ante.*[2] That being said, especially when if private information provides the main obstacle for or against (real) delegation, this could possibly be remedied, thereby giving the firm a choice between delegation or non-delegation.[3] For this case, we obtain some additional implications.

The optimal choice between delegation or non-delegation, then, depends on how the respective levels of inefficiency compare. Our core (robust) insight here is that delegation is less likely to be optimal the more important the decision is to the firm, in terms of value destruction that would result from inefficient implementation. There are two reasons for this. First, the higher the firm's stake, the more the firm is already committed to make *ex post* an efficient decision. Hence, the problem of "underinvestment" becomes smaller. Second, as the agent does not internalize the firm's higher stake unless the contract is changed, the inefficiency that results from the agent's "overinvestment" bias will increase. In addition, at least as long as inefficiencies are still small, we can show generally that inefficiencies are lower with delegation if the agent's effort costs are high compared to the value of the innovation.

These implications suggest that the firm should delegate "small decisions." Note that this finding is not driven by simply assuming that the agent has better information with regards to these decisions or, likewise, by assuming that the firm (or superior) has fixed costs to acquire information, implying that it would be worthwhile to incur these costs only for sufficiently "large" decisions.

Moreover, these implications also suggest that delegation, while being associated with more innovation, implies a lower return on the average innovation that is undertaken. This holds both as delegation (inefficiently) increases the range of undertaken innovations, thus lowering the hurdle, and as delegated decisions are associated with a worse *ex ante* relationship between payoff and (effort) cost.

Though we do not obtain a generally monotonic effect of the informativeness of the performance measure, we find, also, that the "overinvestment" bias is particularly severe if the performance measure is very noisy. While the firm's expected residual payoff conditional on the innovation's profitability is always sufficiently "steep" to make the firm discriminate between good and bad innovations, the agent's payoff becomes almost flat if the performance measure is very noisy.

[2] Hence, as the firm clearly would not like to grant an employee the right to spend corporate resources on whatever he likes, it may be hard to draw a clear boundary between expenditures over which the agent has uncontested authority and those decisions over which the firm retains ultimate control.

[3] Likewise, holding back information from the agent's superiors may ensure that delegation is truly credible.

As already noted above, our analysis relates to that in Rotemberg and Saloner (1994). One of our results is to show that the problem of "underinvestment" in new ideas, which is at the core of their analysis, can be mitigated or even fully overcome if the presence of a sufficiently informative performance measure allows the agent to be rewarded with steep incentive pay. In addition, we show that while delegation can also solve the problem of "underinvestment," it may not lead to overall more-efficient decision making.

In our model, the agent's compensation is the source of the "overinvestment" or "underinvestment" bias. Our analysis ties into the large extant literature on incentive pay. In a different context, Holmstrom (1989) has argued that incentive pay, which may be needed to stimulate innovation, could create a bias in a multi-task setting. The case with delegation, where the agent effectively faces a multi-task problem, is similar to that in Levitt and Snyder (1997) and Inderst and Klein (2007). Besides the comparison between different organizational forms, we also conduct a comparative statics analysis in the optimal compensation scheme.

The rest of the paper is organized as follows. Section 2 introduces the model. Sections 3 and 4 analyze the cases with and without delegation. Section 4 presents additional comparative results on the two cases. Section 5 discusses some of the modeling assumptions, and Section 6 concludes.

2. The model

As our model applies more generally to different levels of a firm's hierarchy, we adopt in what follows a more abstract language. In our model, a principal, who represents the interests of the whole firm, contracts with an agent. The principal could be the agent's superior manager but also the firm's owner or its board. The agent could represent any subordinate line manager or, possibly, an engineer in the firm's research department. We now lay out the basic features of our model.

The model has three periods: $t=0,1,2$. In the first period, $t=0$, the agent is hired. In the final period, $t=2$, the game ends and payoffs are realized. In period $t=1$, the two stages of innovation management take place. First, the agent can exert effort to generate a new idea. Second, a decision is made on whether an idea is actually implemented. Next, we provide more details.

If the agent does not succeed in generating an innovation, we suppose that his continued work realizes a strictly positive payoff Y with probability $0<\widehat{p}<1$, while with the residual probability $1-\widehat{p}$ the payoff is zero.[4] Clearly, $\widehat{p}$ still applies if an innovation was generated but ultimately not implemented. In both cases, the firm continues to conduct its business as before. Instead, an innovation may represent a new product or technology developed by the firm's research unit. It could also represent a different way of (re-)structuring the firm's processes or, on a smaller scale, the organization of the employee's own work.

Generating an innovation comes at private disutility $c>0$ to the agent. At the outset, the success probability p that could be achieved if an innovation were implemented is still uncertain and drawn from the distribution $F(p)$, which has support $p\in P:=[\underline{p},\overline{p}]$ and everywhere strictly positive density $f(p)$. Thus, as $\widehat{p}Y$ represents the firm's (true) opportunity costs from implementing the innovation, it is efficient to do so if and only if $p\geq\widehat{p}$. As for $\widehat{p}\geq\overline{p}$ there would be no benefits from generating an innovation, we stipulate that $\widehat{p}<\overline{p}$.

Contracts and organization: The worker is hired under a contract that can be based on a performance measure. It stipulates a base wage $\alpha\geq 0$ and a bonus $\beta\geq 0$. The bonus is paid only if "high performance" in $t=2$. For this, we specify that while the high payoff Y itself is not verifiable, it results with probability $0.5<\pi\leq 1$ in a "high" outcome of the (for $\pi<1$ noisy) performance measure. Likewise, in case of a failure (i.e., a payoff of zero), the performance measure is informative (i.e., "low") only with probability π.[5]

If the true success probability is p, the probability with which the bonus will be paid is, thus, given by

$$\phi(p):=p\pi+(1-p)(1-\pi). \tag{1}$$

Note that the case with $\pi=1$ corresponds to that in which Y itself is verifiable.

Another feature of the employment contract is that it is at-will. This has the following implications. Without a new, implemented innovation, the agent can be simply replaced by an equally qualified worker from outside. As the latter is willing to work at the (market) reservation value of w_m, it follows for the case in which no innovation was implemented that, irrespective of the initial contract (α,β), the agent's continuation payoff is equal to w_m.[6]

Instead, if the agent's innovation is implemented, he becomes less dispensable. For simplicity, we suppose that, in this case, the innovation has value to the firm only if the agent is retained: The likelihood of success when the agent leaves drops to $p=0$.

In what follows, we distinguish between two regimes. In one regime, the decision of whether or not to implement the innovation lies with the principal. In the second regime, the decision is delegated to the agent. In either regime, the party

[4] It is straightforward to extend the model, though at the cost of complicating all expressions, to the case in which both payoffs are strictly positive. To ensure that the firm always has enough cash to pay the promised compensation, we may suppose that the firm generates additional and always strictly positive cash flow from other units.

[5] We could imagine that the payoff from the agent's work (unit) cannot be made perfectly verifiable as, it is part of the firm's overall performance.

[6] This would clearly be the case if the agent left, i.e., if he quit or was fired. It also holds when the initial compensation contract is renegotiated. The case in which the firm grants the agent some additional job security (e.g., in the form of severance pay) is further discussed below.

with the decision right is assumed to privately observe p. For the time being, we abstract from the possibility of renegotiations before the decision is made. Also, we assume that the firm can credibly commit to delegate the decision.

3. The case without delegation

3.1. Preliminary analysis

We solve the game backwards. Suppose, first, that the agent has generated an idea with success probability p. When deciding whether or not to implement the idea, the firm (i.e., the principal acting in the firm's interest) weighs off the following two payoffs. If the innovation is not implemented, we know that the firm's expected payoff net of the agent's wage is

$$\widehat{p}Y - w_m. \tag{2}$$

Instead, if for a given p the innovation is implemented, the firm's expected payoff becomes

$$Yp - [\alpha + \beta\phi(p)],$$

where the term in rectangular brackets captures the agent's expected incentive pay. Recall, also, that $\phi(p)$, as defined in (1), captures the firm's performance measure (with $\phi(p) = p$ if the performance measure is perfectly informative).

It is immediate that we can restrict consideration to contracts that satisfy $\beta < Y$. The firm's payoff from the innovation is, thus, strictly increasing in p.[7] This implies that the principal's optimal decision corresponds to a simple cutoff rule: The principal realizes the innovation if and only if $p \geq p^*$. If p^* is interior, it must satisfy

$$Yp^* - [\alpha + \beta\phi(p^*)] = \widehat{p}Y - w_m. \tag{3}$$

In words, at the critical value p^*, the firm's expected profits *net* of the employee's expected compensation is just equal with and without the innovation. There may, however, not always exist an interior cutoff. If

$$Y\underline{p} - [\alpha + \beta\phi(\underline{p})] \geq \widehat{p}Y - w_m \tag{4}$$

holds, the principal will always want to implement the innovation. We capture this case by stipulating that the principal then applies the cutoff $p^* = \underline{p}$. At the other extreme, the principal prefers to never implement the innovation, which is the case if

$$Y\overline{p} - [\alpha + \beta\phi(\overline{p})] \leq \widehat{p}Y - w_m. \tag{5}$$

As $p^* = \overline{p}$ is a zero-probability event, we capture this case by specifying the cutoff $p^* = \overline{p}$.

Lemma 1. *In the case without delegation, depending on the compensation contract* (α, β), *the firm takes the following decision. If*

$$Y\underline{p} - [\alpha + \beta\phi(\underline{p})] < \widehat{p}Y - w_m < Y\overline{p} - [\alpha + \beta\phi(\overline{p})]$$

holds, then there is a unique interior cutoff $\underline{p} < p^* < \overline{p}$ *as given by* (3) *such that the innovation is undertaken if and only if* $p \geq p^*$. *Otherwise, the innovation is always undertaken if* (4) *holds and never undertaken if* (5) *holds.*

So far, we have assumed that an innovation is generated. Lemma 1 then describes the decision whether it is implemented or not. An innovation is, however, generated only if the agent exerts effort at the beginning of period $t = 1$. To set up the agent's respective incentive constraint, recall, first, that the agent can ensure himself an expected compensation above w_m *only* if he generates an innovation *and* if this innovation is subsequently implemented. As he must exert effort at cost c, his incentive compatibility constraint becomes

$$\int_{p^*}^{\overline{p}} [\alpha + \beta\phi(p) - w_m] f(p)\, dp \geq c. \tag{6}$$

The agent's expected compensation, thus, must lie sufficiently above his market wage w_m.

3.2. Optimal contract design

It is now instructive to suppose for the moment that the performance measure is fully uninformative as $\pi = 0.5$. In this case, the employee's expected wage after an innovation was implemented, $\alpha + \beta\phi(p)$, is independent of p and given by $\alpha + \beta/2$. This is analogous to making a lump-sum payment $b = (\alpha + \beta/2) - w_m$ in addition to the market wage w_m. We now make several observations with respect to this particular case of our model.

First, suppose for the moment that it was feasible to make the implementation decision contingent on p, in which case setting $p^* = \widehat{p}$ would be efficient and clearly also optimal for the firm. Substituting $p^* = \widehat{p}$ into (6) and solving with $\pi = 0.5$ for the required lump-sum payment b, we find that paying $b = c/[1 - F(\widehat{p})]$ is just sufficient to incentivize the agent.

[7] Note that $\phi'(p) = 2\pi - 1 \leq 1$.

The crux, however, is that the principal's decision is not *ex ante* contractible. Instead, the choice of the respective cutoff p^* is made such that it is *ex post* optimal for the principal (cf. Lemma 1). Suppose, now, that the efficient cutoff was interior: $\widehat{p} \in (\underline{p}, \overline{p})$. For any strictly positive lump-sum payment $b > 0$, this implies, however, that at $p = \widehat{p}$, the principal would strictly prefer *not* to implement the innovation. Consequently, for given $b > 0$, the resulting cutoff would be strictly above $\widehat{p}$. To be precise, it would be equal either to $p^* < \overline{p}$ satisfying $(p^* - \widehat{p})Y = b$ (cf. condition (3)) or to $p^* = \overline{p}$ if $(\overline{p} - \widehat{p})Y \leq b$. Hence, with a lump-sum payment, the principal will *always* implement the innovation too infrequently.

If the performance measure is completely uninformative ($\pi = 0.5$), the principal cannot design a better contract and, thereby, implement a more-efficient decision rule. We show next that this is different when the performance measure is informative as $\pi > 0.5$. In this case, a more-efficient decision rule can be achieved by linking the agent's compensation to the success of the implemented innovation. Crucially, the purpose of the agent's incentive pay, then, is not to elicit more effort, but to ensure more-efficient decision making by the principal.

The role of incentive pay: The intuition for the optimality of incentive pay is as follows. Note, first, that with $\beta > 0$ and with an informative performance measure, the agent's expected compensation under an innovation, $\alpha + \beta\phi(p)$, becomes strictly increasing in p. Thus, the agent's expected reward, which is necessary to elicit effort, is provided in a way such that *ex post*, it is higher for high values of p and lower for low values of p. Note, now, that for high values of p, the additional compensation above w_m is inconsequential for the principal's decision, as he still prefers to implement the innovation. Instead, it is for lower values of p that the wedge between the expected compensation with the innovation and the market wage w_m could tilt his decision, making him forego an innovation even though it is overall profitable. The more the agent's expected reward is now shifted into high realizations of p, which is the case if the incentive pay $\alpha + \beta\phi(p)$ becomes steeper, the further can p^* be pushed down. In fact, as we show next, it may then even be possible to implement the first-best decision.

If it is possible to implement the first-best decision, then it is clearly optimal to do so. This follows immediately from the observation that, by optimality for the principal, the agent's incentive constraint (6) will be binding, implying that the principal becomes the full residual claimant. From an *ex ante* perspective, the principal's expected payoff is, thus, maximized if he chooses *ex post* the efficient cutoff $p^* = \widehat{p}$. To derive the condition for when this is feasible, there are two cases to distinguish: that with $\widehat{p} \leq \underline{p}$ and that with $\underline{p} < \widehat{p} < \overline{p}$.[8]

Obtaining the first-best decision: Suppose, first, that $\underline{p} < \widehat{p} < \overline{p}$. Hence, under the efficient cutoff $p^* = \widehat{p}$, the innovation should, with strictly positive probability, sometimes be implemented and sometimes not. In this case, the uniquely optimal contract is, then, pinned down by substituting $p^* = \widehat{p}$ into both the decision rule (3) and the binding incentive constraint (6). This obtains

$$\beta = \frac{c}{2\pi - 1} \frac{1}{\int_{\widehat{p}}^{\overline{p}} [p - \widehat{p}] f(p)\, dp} \tag{7}$$

and

$$\alpha = w_m - \frac{c}{2\pi - 1} \frac{(1-\pi) + \widehat{p}(2\pi - 1)}{\int_{\widehat{p}}^{\overline{p}} [p - \widehat{p}] f(p)\, dp}. \tag{8}$$

This contract is clearly feasible only if the respective base wage in (8) satisfies $\alpha \geq 0$. If this is not the case, then it is optimal to set α as low as possible: $\alpha = 0$. The decision is then distorted with $p^* > \widehat{p}$.

In the second case, we have that $\widehat{p} \leq \underline{p}$: the innovation should always be implemented. In this case, there may be some leeway to specify the optimal contract. While, by optimality, the agent's incentive constraint (6) must still be satisfied with equality, the decision rule of the principal is no longer pinned down by an indifference condition, as was previously the case for $\underline{p} < \widehat{p} < \overline{p}$. Instead, for p^* to be chosen efficiently, the inequality requirement (4) must now be satisfied. As is intuitive, this puts only a *lower* bound on the steepness β of the incentive scheme. Whether it is feasible in this case to implement the efficient decision rule depends, again, on whether for this value of β, the corresponding base wage α is still non-negative.

To make this formal, define first $E[p] := \int_{\underline{p}}^{\overline{p}} p f(p)\, dp$ and $E[\phi(p)] := \int_{\underline{p}}^{\overline{p}} \phi(p) f(p)\, dp$. From the binding incentive constraint (6) and condition (4), we have that, in the present case, the first best $p^* = \underline{p}$ can be obtained by specifying

$$\beta \geq \frac{c}{2\pi - 1} \frac{1}{E[p] - \widehat{p}}, \tag{9}$$

which can then be substituted into $\alpha = w_m + c - \beta E[\phi(p)]$. This is indeed feasible if

$$w_m - \frac{c}{2\pi - 1} \frac{E[\phi(p)]}{E[p] - \widehat{p}} \geq 0.$$

Optimal incentive compensation: In what follows, we want to streamline the exposition of results by focusing without loss of insights on the case in which $\underline{p} < \widehat{p} < \overline{p}$.

[8] Recall that for $\widehat{p} \geq \overline{p}$, it would clearly not be optimal to incentivize the agent to exert effort for the creation of an innovation that is always *ex post* unprofitable.

Assumption. For what follows, we restrict consideration to the case in which it is sometimes, though not always, efficient to implement the innovation: $\underline{p} < \widehat{p} < \overline{p}$.

The fact that the contract is pinned down uniquely in this case facilitates the derivation of comparative statics results.[9]

Proposition 1. *If the decision is not delegated, then the optimal contract prescribes the following compensation scheme. If*

$$c\left[\frac{1-\pi}{2\pi-1}+\widehat{p}\right] \leq w_m\left[\int_{\widehat{p}}^{\overline{p}}[p-\widehat{p}]f(p)\,dp\right] \tag{10}$$

holds, then incentive compensation allows the first-best outcome to be achieved, where the unique optimal contract (α, β) *is characterized by* (7) *and* (8). *Otherwise, i.e., if* (10) *does not hold, the innovation will be implemented too infrequently compared to the first-best benchmark as* $p^* > \widehat{p}$. *In this case, it is uniquely optimal to set* $\alpha = 0$ *and*

$$\beta = \frac{1}{2\pi-1}\frac{c-[1-F(p^*)]Y(p^*-\widehat{p})}{\int_{p^*}^{\overline{p}}[p-p^*]f(p)\,dp}. \tag{11}$$

Proof. We first fully specify the principal's contract design problem in $t = 0$. The objective function, taking the subsequent decision rule p^* as given, is to maximize expected profits net of wages:

$$F(p^*)[\widehat{p}Y - w_m] + \int_{p^*}^{\overline{p}}[pY - \beta\phi(p) - \alpha]f(p)\,dp. \tag{12}$$

Substituting the binding incentive constraint (6) into (12),[10] we have

$$Y\left[\widehat{p}F(p^*) + \int_{p^*}^{\overline{p}} pf(p)\,dp\right] - w_m - c,$$

which is strictly quasiconcave in p^* and is maximized at $p^* = \widehat{p}$. If (10) holds, then the characterization of the unique optimal contract follows from setting $p^* = \widehat{p}$ in (3) and (6). Furthermore, (10) is obtained from requiring that $\alpha \geq 0$, where α is given by (8).

Suppose, next, that (10) does not hold. We argue, first, that the optimal contract must specify $\alpha = 0$. Suppose that this was not the case, such that $\alpha > 0$ while also $p^* > \widehat{p}$. Consider now a marginal adjustment $d\alpha < 0$ and $d\beta = |d\alpha|\phi(p^*)$, such that from (3) p^* remains unchanged. As $\phi(p)$ is strictly increasing, we have from (6) that the incentive constraint is relaxed. The final step is to reduce $d\beta$ until the incentive constraint again becomes binding, which also reduces p^* and, thus, increases the objective function (12). (To ensure that the incentive constraint becomes binding, we use here from (3) that p^* is continuous and strictly increasing in β.)

Substituting $\alpha = 0$ and from (3) that

$$\beta = \frac{Y(p^*-\widehat{p}) + w_m}{\phi(p^*)},$$

we obtain for the binding incentive constraint

$$w_m\int_{p^*}^{\overline{p}}\left[\frac{\phi(p)}{\phi(p^*)}-1\right]f(p)\,dp + Y(p^*-\widehat{p})\int_{p^*}^{\overline{p}}\frac{\phi(p)}{\phi(p^*)}f(p)\,dp = c. \tag{13}$$

While the left-hand side of (13) is not monotonic in p^*, the optimal contract clearly implements the lowest value of p^* that solves (13). Finally, for $\alpha = 0$ and given p^* from (13), the unique optimal incentive component β is again obtained from the binding incentive constraint (6), which together with the principal's decision rule in (3) leads to (11). □

Discussion: Before we continue with some comments on Proposition 1, we need to first spell out more explicitly some assumptions that we have made implicitly. Suppose, first, that (10) holds such that it is feasible to implement the efficient decision rule. We assumed that it is optimal for the principal to elicit the agent's effort in the first place, which, in this case, holds if

$$Y\int_{\widehat{p}}^{\overline{p}}[p-\widehat{p}]f(p)\,dp \geq c. \tag{14}$$

If (10) is not satisfied, an assumption analogous to (14) must hold, where now, however, the efficient decision rule must be substituted by the inefficiently high cutoff $p^* > \widehat{p}$.[11]

[9] This restriction can also be given the following additional motivation. Otherwise—i.e., if $\widehat{p} \leq \underline{p}$ holds—it may be possible to commit to the efficient decision rule simply by stipulating that *any* innovation of the agent will be implemented. (Admittedly, this may create agency problems of its own, as the agent may then come up with ideas that are "outside" those captured by the set $p \in P$.) Also, in this case, it will be straightforward from the following analysis that it is optimal (though sometimes not uniquely so) to simply delegate the decision to the agent, provided that such a choice is feasible.

[10] It is standard to show that the incentive constraint binds by optimality.

[11] As shown in the proof of Proposition 1, we no longer obtain an explicit characterization of the resulting cutoff p^* for this case. As a consequence, we can no longer express explicitly, and only in terms of the model's primitives, the respective condition in analogy to (14).

When is it more likely that the first-best decision rule can be obtained? Inspection of condition (10) obtains the following results.

Corollary 1. *Without delegation, it is (ceteris paribus) more likely that the first-best outcome can be obtained. That is, condition (10) is satisfied, if:*

(i) *it is less costly to elicit the agent's effort (lower c);*
(ii) *the innovation is ex ante relatively more profitable as either the distribution $F(p)$ improves in the sense of first-order stochastic dominance (FOSD) or as $\widehat{p}$ decreases;*
(iii) *the performance measure is more precise (higher π) or*
(iv) *the agent's market wage is higher (higher w_m).*

Assertion (i) is immediate, as a lower value of c reduces the expected reward that is necessary to elicit effort from the agent, which, in turn, makes the principal less biased towards rejecting the innovation. The comparative change in assertion (ii) makes an innovation relatively more successful (from an *ex ante* perspective) compared to the existing business. If this is the case, as $\widehat{p}$ decreases, then the first-best cutoff $p^* = \widehat{p}$ also decreases. This increases the probability with which the innovation will be implemented, which allows to incentivize the agent by promising him a already a lower reward if the innovation is implemented. Again, this reduces the principal's bias. Likewise, if $\widehat{p}$ stays constant, but more probability mass is shifted into values $p \geq \widehat{p}$, then implementing the innovation again becomes more likely.[12] In addition, with incentive pay, such a shift of $F(p)$ allows to make the agent's expected compensation *conditional on p* steeper.

Furthermore, an increase in the informativeness of the performance measure, as in assertion (iii), also allows to shift more of the agent's expected reward into higher p, thereby reducing the principal's bias at lower values of p and, thus, pushing down p^*. Finally, the comparative analysis of the agent's market wage (cf. assertion (iv)) is similar to that in other agency problems. As seen formally in (8), for given β, a higher value of w_m makes it necessary to pay the agent a larger base wage, which, in turn, ensures that his limited liability constraint, $\alpha \geq 0$, is more likely to still be satisfied.

As noted in the Introduction, our analysis for the case with $\pi = 0.5$ is analogous to that in Rotemberg and Saloner (1994), who pointed out that a lump-sum reward for the agent distorts the principal's decision of whether or not to implement the innovation. Our analysis shows that if the agent's pay can be linked to the success of the innovation, then the principal's decision may no longer be distorted under the optimal incentive contract (cf. the comparative analysis in Corollary 1).

Illustrative example: To illustrate our results from Proposition 1 and Corollary 1, we take the case in which p is distributed uniformly over $[0, \overline{p}]$, with $0.5 < \overline{p} \leq 1$. Set, also, $\widehat{p} = 0.5$. For this example, condition (10) becomes

$$y := \frac{c}{w_m} \frac{1}{2\pi - 1} \leq \frac{1/4 - \overline{p}(1 - \overline{p})}{\overline{p}}. \tag{15}$$

Here, the left-hand side of (15), y, provides an aggregate measure for the severity of the agency problem, which increases in c and decreases in π and w_m. If we make the innovation *ex ante* as promising as possible by increasing the upper boundary of the support to $\overline{p} = 1$, then this condition simplifies to $y \leq \frac{1}{4}$. If, in addition, $\pi = 1$ holds, then we have, further, that $c \leq w_m/4$. Hence, the additional compensation that the agent must receive as a reward for his innovative activity must not be too large compared to his market wage. (Recall, also, that the agent's total expected wage equals the sum of his market wage w_m and his effort cost c, for which he has to be additionally compensated.)

If the first best is feasible, then from (7), the optimal (incentive) bonus equals[13]

$$\beta = \frac{c}{2\pi - 1} \frac{8\overline{p}}{4\overline{p}(\overline{p} - 1) + 1}.$$

In the "best case," where $\pi = 1$ and $\overline{p} = 1$, this simplifies to $\beta = 8c$, which is then matched with a base wage of $\alpha = w_m - 4c$.

3.3. Comparative analysis

By taking the limit as the performance measure becomes increasingly uninformative with $\pi \to 0.5$, we show in the proof of Corollary 2 that, in this case, the equilibrium cutoff p^* is the lowest value satisfying

$$Y(p^* - \widehat{p})[1 - F(p^*)] = c. \tag{16}$$

We denote this value by $\widetilde{p} > \widehat{p}$.

From inspection of (16), it is immediate that the lowest value $p^* < \overline{p}$ that satisfies this condition (provided it exists) is smaller as c decreases. This is intuitive as then a smaller reward is already sufficient to incentivize the agent, which causes

[12] Note that assertion (ii) holds strictly only if the respective FOSD shift occurs (also) to the right of $\widehat{p}$.

[13] It should be noted that when conducting a comparative analysis for β—e.g., when letting $\pi \to 0.5$—this is feasible only as long as (10) is still satisfied.

a smaller bias. We also have that p^* is lower as y increases.[14] The intuition for why an increase in Y leads to a lower value of p^* and, thus, to a lower distortion $p^* - \widehat{p} > 0$ is as follows. For higher Y, the principal has more incentives to make an unbiased decision because, as the residual claimant, he has more at stake: Holding the compensation contract fixed, for given $p^* > \widehat{p}$, the loss from an inefficient decision to reject the innovation, $Y(p^* - \widehat{p})$, is strictly increasing in Y. Furthermore, in analogy to the comparative analysis conducted in Corollary 1, we have that p^* also decreases if it becomes *ex ante* more likely that the innovation is profitable and, thus, efficient to implement.

As with Corollary 1, these comparative results for the limit case with $\pi = 0.5$ also hold if incentive pay is feasible, as the performance measure is informative, $\pi > 0.5$. In the latter case, if[15]

$$c\widehat{p} \leq w_m \left[\int_{\widehat{p}}^{\overline{p}} [p - \widehat{p}] f(p)\, dp \right]$$

holds, then as π increases from 0.5 to 1, the resulting cutoff p^* decreases from the upper boundary $\widetilde{p}$ to the efficient cutoff $\widehat{p}$. Finally, again mirroring the results from Corollary 1, if incentive pay is feasible, then an increase in w_m also allows a more-efficient (second-best) cutoff to be implemented.

Corollary 2. *If the decision is not delegated and if the first best cannot be achieved under the optimal contract, i.e., if* (10) *does not hold, then the resulting distortion* $p^* - \widehat{p} > 0$ *becomes smaller if*:

(i) *it is less costly to elicit the agent's effort* (*lower* c);
(ii) *the innovation is ex ante relatively more profitable as either the distribution* $F(p)$ *improves in the sense of FOSD or as* $\widehat{p}$ *decreases*;
(iii) *the performance measure is more precise* (*higher* π);
(iv) *the agent's market wage is higher* (*higher* w_m) *or*
(v) *the outcome in case of success is higher* (*higher* Y).

Proof. Suppose $\pi > 0.5$. Recall from the proof of Proposition 1 that the equilibrium cutoff p^* is the lowest value that solves (13). This implies that the derivative of the left-hand side of (13) must be strictly positive at the respective value of p^*. As we now marginally reduce c, this allows us to find a new, strictly smaller value in the neighborhood of p^*, such that (13) is still satisfied with equality, implying that the new optimal value for p^* must indeed be strictly smaller. Likewise, if a comparative change increases the right-hand side of (13) for all $\widehat{p} < p^* < \overline{p}$, then we can again conclude that the equilibrium value of p^* is strictly smaller.

To apply these insights, note, first, that the left-hand side of (13) is strictly increasing in Y (assertion (v)) as well as in w_m (assertion (iv)). We now show that the same holds for a variation in π (assertion (iii)). Substituting for $\phi(\cdot)$, the derivative is

$$[w_m + Y(p^* - \widehat{p})] \frac{1}{\phi^2(p^*)} \int_{p^*}^{\overline{p}} (p - p^*) f(p)\, dp > 0.$$

For assertion (ii), the part for $\widehat{p}$ follows immediately. As $\phi(p)/\phi(p^*)$ is strictly increasing in p, we next have that, for fixed p^*, the values of the two integrals on the left-hand side of (13) are indeed strictly higher after a shift in the sense of (strict) FOSD that applies (also) to $p > p^*$.

Finally, to consider the limit $\pi \to 0.5$, it is convenient to substitute $\alpha = 0$ into the binding incentive constraint (6), which together with $\pi = 0.5$, yields $\beta = 2[w_m + c/[1 - F(p^*)]]$. (Note that, by construction, β is then paid with probability $\frac{1}{2}$ irrespective of the realized value of p.) The resulting cutoff p^* as characterized in (16) is finally obtained by substituting $\pi = 0.5$ into (13). Here, as noted in the main text, only assertions (i), (ii), and (v) apply. □

Feasibility of an incentive scheme: If we were to hold p^* fixed, then the agent's incentives to exert effort in $t = 0$ would clearly be strictly increasing in β. However, as a higher value of β will also push up p^*, thus making it less likely that an innovation is actually implemented, it is not guaranteed that this will still relax the agent's incentive constraint. In fact, if the incentive problem is sufficiently severe, then it may no longer be feasible at all to incentivize the generation of new ideas.[16] More precisely, as a higher β pushes up p^* and, thus, reduces $1 - F(p^*)$, to relax the incentive constraint, it may be necessary to further increase β and so on: a potentially vicious circle that may prevent the incentive constraint from being satisfied for any β.

To make this more formal, we first revisit the case with $\pi = 0.5$. Recall that the resulting equilibrium cutoff p^* is the lowest value at which (16) is satisfied, provided that such a value $p^* < \overline{p}$ indeed exists. Recall, also, that in this case, the

[14] The left-hand side of (16) is non-monotonic in p^*, taking on the value of zero at both $p^* = \widehat{p}$ and $\overline{p}$. However, if a change in the exogenous parameters leads to an increase in the left-hand side of (16) for all $\widehat{p} < p^* < \overline{p}$, then the equilibrium value of p^*—i.e., the lowest value satisfying (16)—is again lower.

[15] This condition is obtained from setting $\pi = 1$ in (10).

[16] Intuitively, as shown in the following section, we do not encounter such a problem if the decision is delegated to the agent, in which case the incentive constraint will always be satisfied for sufficiently high values of β.

agent receives an additional lump-sum payment b if his innovation is implemented, from which we have that the agent exerts effort only in case $b \geq c/[1-F(p^*)]$. While an increase in b relaxes the agent's incentive constraint if we hold p^* constant, note, also, that the true value of p^* is, in this case, given by $(p^*-\widehat{p})Y = b$ and, thus, increases with b. (Condition (16) is then obtained from these two (binding) constraints.)

Corollary 2 also generates implications for when it is more likely that the agent can indeed be incentivized to generate an innovation. Again, taking the case with $\pi = 0.5$, condition (16) is more likely to be satisfied for some value $p^* < \overline{p}$ if for any given value $\widehat{p} < p^* < \overline{p}$ the left-hand side increases or, alternatively, if the right-hand side decreases. This mirrors our previous comparative analysis for the (second-best) equilibrium cutoff, where we presumed that such a value $p^* < \overline{p}$ exists. The same observation applies to the case with $\pi > 0.5$. From these observations, together with Corollary 2, we immediately have the following result.

Corollary 3. *Without delegation, it is not always feasible to incentivize the agent. Ceteris paribus, the range of parameters for which this is feasible increases if we conduct a comparative change, as in assertions* (i)–(v) *of Corollary* 2.

Illustrative example: Take, again, the case in which $p \in [0, \overline{p}]$ is uniformly distributed with $0.5 < \overline{p} \leq 1$ and $\widehat{p} = 0.5$. For the case with $\pi = 0.5$, we obtain from condition (16), which becomes

$$(p^*-0.5)(\overline{p}-p^*) = \frac{\overline{p}c}{Y}, \tag{17}$$

that an equilibrium $p^* < \overline{p}$ only exists if[17]

$$\left(\frac{\overline{p}-0.5}{2}\right)^2 \geq \frac{\overline{p}c}{Y}.$$

Solving this, we obtain the condition

$$\overline{p} \geq \frac{\left(1+4\frac{c}{Y}\right)+\sqrt{\left(1+4\frac{c}{Y}\right)^2-1}}{2}. \tag{18}$$

To see how tight condition (18) is, we choose the most "favorable" distribution with $\overline{p} = 1$. In this case, we find, after some transformations, that (18) holds only if $c/Y \leq \frac{1}{16}$. To put this threshold into perspective, we calculate the expected surplus if the decision rule was contractible and, thus, equal to $p^* = \widehat{p} = 0.5$. Under the efficient-decision rule, it is then also efficient to elicit effort if $c/Y \leq \frac{1}{8}$. Hence, for all $c/Y \in (\frac{1}{16}, \frac{1}{8})$, it would be first-best optimal to incentivize the agent, but it is not feasible if the decision is not delegated and the performance measure is perfectly noisy.

To compare this with the case in which incentive contracts are feasible, though the first best cannot be obtained, we make the performance measure as informative as possible: $\pi = 1$. To simplify expressions for this illustration, we further set $w_m = 0$, which—as we know—limits the scope of incentive contracts. From Proposition 1 (namely, condition (13)), we then obtain that the equilibrium cutoff p^* is the lowest value that solves

$$(p^*-0.5)(\overline{p}-p^*)\left(\frac{\overline{p}+p^*}{2p^*}\right) = \frac{\overline{p}c}{Y}. \tag{19}$$

We have already arranged this expression to allow for a direct comparison with condition (17), which pins down s^* if no incentive contract is feasible. The difference is the last term in brackets on the left-hand side of (19). As this is strictly larger than one for $p^* < \overline{p}$, we can already confirm the following. First, if a solution exists, then the inefficiency is lower with the incentive contract. Second, the range of parameters for which such a solution $p^* < \overline{p}$ indeed exists is strictly larger with the incentive contract.

For $\overline{p} = 1$, we finally obtain from (19) that there is an interior solution if and only if $c/Y \leq 0.0736$. Recall that without the incentive contract the condition was $c/Y \leq \frac{1}{16} = 0.0625$.

4. The case with delegation

Delegating the decision seems like a natural response to the principal's commitment problem. Then the principal would no longer have the opportunity to inefficiently decide against the innovation in order to keep the employee's (expected) compensation lower. However, though delegation can cure the "underinvestment" problem, it gives rise to a problem of "overinvestment."

To see this, it is again helpful to start with the benchmark case in which the performance measure is perfectly uninformative. This allows to incentivize the agent by making a lump-sum payment if the innovation is implemented (as in Rotemberg and Saloner, 1994). With delegation, this results in an extreme problem of "overinvestment": Any such lump-sum payment, however, induces the agent to implement the innovation for *any* value of p. The role of incentive pay when the performance measure is informative ($\pi > 0.5$) is now to make this "overinvestment" problem less severe.

[17] For this, we maximise the left-hand side of (17) and subsequently substitute the maximand.

To see this, recall, first, that for $\beta > 0$ and $\pi > 0.5$, the agent's compensation, $\alpha + \beta\phi(p)$, is strictly higher the more promising the innovation is. Consequently, unless the agent never or always chooses to implement the innovation, he will again apply a cutoff rule characterized by some p^* with $\underline{p} < p^* < \overline{p}$. At this cut-off his market wage w_m, which equals the compensation that he receives without the innovation, is just equal to his expected compensation with the innovation, $\alpha + \beta\phi(p^*)$. From this indifference condition, we have that

$$\phi(p^*) = \frac{w_m - \alpha}{\beta}. \tag{20}$$

On the other hand, the agent will always want to implement the innovation if

$$\alpha + \beta\phi(\underline{p}) \geq w_m, \tag{21}$$

while the innovation would not be implemented with positive probability if

$$\alpha + \beta\phi(\overline{p}) \leq w_m. \tag{22}$$

Lemma 2. *In the case with delegation, for*

$$\phi(\underline{p}) < \frac{w_m - \alpha}{\beta} < \phi(\overline{p}),$$

there is a unique interior cutoff $\underline{p} < p^* < \overline{p}$, *as given by* (20), *such that the innovation is undertaken if and only if* $p \geq p^*$. *Otherwise, the innovation is always undertaken if* (21) *holds and never if* (22) *holds.*

Clearly, in equilibrium for all p, the case in which the agent exerts effort to generate an innovation but subsequently decides against its implementation will not arise.

Next, to ensure that in $t = 0$, the agent exerts effort, the incentive compatibility constraint (6) must still hold, as in the case without delegation. The only difference from the previous case is that, now, the respective cutoff p^* must be substituted from Lemma 2.

The role of incentive pay: Can incentive pay again ensure that the efficient decision rule applies with $p^* = \widehat{p}$? The somewhat surprising answer is that this is the case if and only if this is also the case without delegation. To see this, note, first, that if $p^* = \widehat{p}$ is feasible, then both with and without delegation, the optimal compensation is pinned down by the same two requirements. The first requirement is the agent's incentive compatibility constraint (6) evaluated at $p^* = \widehat{p}$. The second requirement is that the first-best decision rule is indeed applied, which holds in either case if and only if at $p = \widehat{p}$, the agent's expected compensation is independent of whether or not the innovation is implemented:[18]

$$\alpha + \beta\phi(\widehat{p}) = w_m. \tag{23}$$

The argument for why incentive pay mitigates the agent's "overinvestment" bias is now perfectly symmetric to the previous argument without delegation. There, incentive pay reduced the principal's "underinvestment" bias through making the adoption of the innovation more attractive at relatively low values of p. By analogy, reducing the agent's expected compensation for low p incentive pay pushes up p^* in the case of delegation and, thereby, reduces the "overinvestment" problem.

Proposition 2. *The first best can be achieved with delegation if and only if this is also possible without delegation. The unique optimal compensation is characterized by* (7) *and* (8), *as in the case without delegation. Otherwise—i.e., if* (10) *does not hold—the innovation is now implemented too often with* $p^* < \widehat{p}$, *in which case the optimal contract specifies* $\alpha = 0$ *and*

$$\beta = \frac{1}{2\pi - 1} \frac{c}{\int_{p^*}^{\overline{p}} [p - p^*] f(p)\, dp}. \tag{24}$$

Proof. The contract-design problem for the principal is the same as in the proof of Proposition 1. Likewise, if $p^* = \widehat{p}$ is feasible as (10) holds, then the characterization of the unique optimal contract follows again from (20), together with the binding incentive constraint (6), where we substitute $p^* = \widehat{p}$.

Next, if (10) does not hold, then the argument for why optimally $\alpha = 0$ must be satisfied is analogous to that in the proof of Proposition 1. Substituting $\alpha = 0$ and $\beta = \phi(p^*)/w_m$ from (20), the binding incentive constraint transforms to

$$w_m \int_{p^*}^{\overline{p}} \left[\frac{\phi(p)}{\phi(p^*)} - 1 \right] f(p)\, dp = c. \tag{25}$$

Note that if (25) has a solution satisfying $p^* > \underline{p}$, which is what we presently assume, then it is unique. Finally, (24) is obtained from the binding incentive constraint (6) together with the agent's decision rule in (20). □

By analogy with Corollary 2, if the first best is not feasible, the resulting inefficiency is lower if it is less costly to elicit the agent's effort, given that a lower reward is already sufficient to incentivize the agent. The same logic also applies for a FOSD

[18] More formally, we obtain condition (23) when substituting $p^* = \widehat{p}$ into either (3) or (20).

shift in the distribution function $F(p)$, which for given p^* makes it more likely that the innovation is ultimately implemented. Furthermore, as for Corollary 2, an increase in w_m also results in a lower inefficiency.

What is also intuitive, however, is that, with delegation, neither Y nor $\widehat{p}$ affects the prevailing cutoff p^* if the first best cannot be obtained. This follows, as for a given incentive contract, the agent does not take into account the principal's "costs" when a wrong decision is taken: i.e., the (opportunity) cost of $Y(\widehat{p}-p^*)$ if for $p^*<\widehat{p}$ the innovation is still implemented.

Corollary 4. *If the decision is delegated and if the first best cannot be achieved under the optimal contract—i.e., if* (10) *does not hold—then the resulting distortion* $\widehat{p}-p^*>0$ *becomes smaller if*:

(i) *it is less costly to elicit the agent's effort* (*lower* c);
(ii) *the innovation is ex ante relatively more profitable as either the distribution* $F(p)$ *improves in the sense of FOSD or as* $\widehat{p}$ *decreases*;
(iii) *the performance measure is more precise* (*higher* π) *and*
(iv) *the agent's market wage is higher* (*higher* w_m).

Proof. Note, now, that (25) pins down a unique cut-off p^* as the left-hand side is strictly decreasing in p^*. From continuity, assertions (i)–(iv) follow, then, as the right-hand side is strictly decreasing in c, while the left-hand side strictly decreases in $\widehat{p}$ and strictly increases in π, w_m, and from a FOSD shift in $F(p)$. □

Illustrative example: Take, again, the case in which $p\in[0,\overline{p}]$ is uniformly distributed with $0.5<\overline{p}\le 1$ and where $\widehat{p}=0.5$. To simplify expressions, we stipulate also that $\pi=1$. If the first best is not feasible, we have for this case that p^* solves

$$\frac{(\overline{p}-p^*)^2}{2\overline{p}p^*}=\frac{c}{w_m}. \tag{26}$$

Incidentally, note that in the present case with $\underline{p}=0$, there always exists a solution value $p^*>\underline{p}$. Hence, as the performance measure is perfectly informative, under delegation, the "overinvestment" problem will never become too extreme. Solving (26), we obtain explicitly

$$p^*=\overline{p}\left[\left(1+\frac{c}{w_m}\right)-\sqrt{\left(1+\frac{c}{w_m}\right)-1}\right],$$

which is strictly increasing in $\overline{p}$ and strictly decreasing in c/w_m.[19]

5. Implications

5.1. Contract design and over- vs. underinvestment

We showed that the role of incentive pay is "robust" under the two considered regimes and should, thus, be pervasive in innovative firms. A number of recent papers (cf. Lerner and Wulf, 2007) have documented the spread of performance-related pay in such industries and firms, as well as a possibly causal link between incentive pay and innovative performance.

In light of this literature, an at-first counterintuitive implication of our model is that the performance sensitivity of compensation for key staff—e.g., in a firm's research department—would not necessarily depend on their authority and decision rights. (In particular, if the first best is feasible, then contracts are identical with and without delegation.) Wulf (2007) shows that division managers that are also corporate officers (and, thus, may have more-far-reaching decision rights) do not have more-high-powered (i.e., more-performance-sensitive) incentive schemes.[20]

How the sensitivity of performance pay changes in our model follows immediately from a comparative analysis of the first-best contract (cf. Proposition 1), under which both $\alpha>0$ and $\beta>0$ hold strictly if (10) also holds strictly.

Corollary 5. *Under both delegation and non-delegation, if the uniquely optimal compensation contract specifies* $\alpha>0$ *and* $\beta>0$, *which also implies that the first best is feasible, then the following comparative results hold: The incentive scheme is steeper, both in terms of a higher* β *and a higher ratio* β/α, *if*

(i) *it is more costly to elicit the agent's effort* (*higher* c):
(ii) *the performance measure is less precise* (*lower* π) *and*
(iii) *the innovation is ex ante relatively less profitable as either the distribution* $F(p)$ *deteriorates in the sense of FOSD or as* $\widehat{p}$ *decreases.*

[19] Note that with $\overline{p}=1$, we obtain for $c/w_m=\frac{1}{4}$ that $p^*=\widehat{p}=\frac{1}{4}$, confirming our previous result that for $c/w_m\le\frac{1}{4}$ the first best is feasible.

[20] Albeit, given the different scope of their responsibilities, performance measures are wider or more narrowly defined for the respective managers.

As some of these predictions—most notably the negative correlation between the steepness of pay and the informativeness of the performance measure—may, at first, be counterintuitive, recall that our model abstracts from risk aversion and that performance pay serves to induce (more) efficient decision making (instead of higher effort).

If the first best is not feasible, then our model implies "overinvestment" under delegation: An innovation is undertaken too often from the perspective of overall efficiency. The agent's bias arises endogenously. Without delegation, there is "underinvestment," where now the principal's bias is endogenous. Thus, rewarding the agent for an implemented innovation is at the heart of the problem as it biases the agent towards and the principal against implementing an innovation—and, at the same time, it provides the cure: namely, to structure the agent's reward as performance-sensitive incentive pay, provided that this is feasible.

Our insights straddle those of Jensen (1993), who argues for high-powered incentives within corporate research facilities, and Holmstrom (1989), who warns that high-powered incentives in an R&D setting could lead to a "multitasking" problem. While with delegation, our model explicitly incorporates such a multi-tasking role for the agent, we also show that incentive pay may resolve the conflict between the agent's two tasks of generating innovations and deciding on their ultimate implementation.

If the first best cannot be achieved, then our results also suggest that delegation leads to more innovation but reduces the gross returns across all implemented innovations. (Formally, this follows immediately from the lower cutoff p^*.) If firms or divisions could be separated into those with more or less delegation of control rights, then this would provide a potentially testable implication of our model: more innovative activity with delegation, but at the cost of a lower gross return.

This exercise would, however, take the choice of organization to be exogenous. We show in the following section, in which we endogenize the choice between delegation and non-delegation, that, in this case, these implications are likely to be even stronger. That is, as less-profitable innovative tasks are more likely to be delegated, this further reduces the average return of innovations that were implemented under delegation.

5.2. *Optimal choice between delegation and non-delegation*

We now compare the performance of the two organizational forms. As noted in the Introduction, we must presume that both are indeed feasible. In the case of delegation, this requires a formal and credible (re-)allocation of authority. Likewise, without delegation, it must be ensured, in particular, that the principal has access to all relevant information as captured by p. For now, we take this as given.

When comparing the outcomes with and without delegation, it is now convenient to denote the equilibrium cutoffs by p_D^* and p_{ND}^*, respectively. Recall, first, from Propositions 1 and 2 that if (10) holds, then $p_D^* = p_{ND}^* = \widehat{p}$, while otherwise both decisions are biased and satisfy $p_D^* < \widehat{p} < p_{ND}^*$. Given that the firm is always the residual claimant, which formally follows from the fact that the agent's incentive constraint (6) is satisfied with equality under the respective optimal contract, the firm will optimally choose the organizational form under which the expected inefficiency is smaller. For a given cutoff p^*, the expected inefficiency is given by

$$L := Y\left|\int_{p^*}^{\widehat{p}} (\widehat{p} - p) f(p)\, dp\right|, \tag{27}$$

where we have to substitute either $p^* = p_D^*$ or $p^* = p_{ND}^*$.

For a comparison of the two regimes, we first consider the case where under the respective optimal contracts the deviation from the first-best decision is (arbitrarily) small. Formally, this allows us to compare the cases with and without delegation by differentiating the loss function L around the efficient decision rule $p_{ND}^* = p_D^* = \widehat{p}$. We conduct the comparative analysis in terms of c. For this, we can first define the threshold $c = c_{FB}$ at which condition (10), which ensures that the first-best decision rule is feasible, is just satisfied with equality. As we then marginally increase c above c_{FB}, the first best can no longer be obtained both with and without delegation.

Proposition 3. *If the incentive problem is not severe, such that* (10) *holds, then the firm is indifferent between delegating the decision or not. If the first best cannot be obtained, but if the inefficiency remains* (*arbitrarily*) *small, namely as* $c - c_{FB} > 0$ *remains* (*arbitrarily*) *small, then delegation is strictly preferred if*

$$\frac{c_{FB}}{Y \int_{\underline{p}}^{\overline{p}} [p - \widehat{p}] f(p)\, dp} > \frac{1}{2}. \tag{28}$$

Proof. See Appendix.

As noted above, the explicit condition (28) is derived by differentiating at $c = c_{FB}$. As such, the general validity of the derived comparative statics result is, of course, limited. Still, condition (28) provides some useful intuition. Delegation is more likely to be the best choice (for low inefficiencies) if the innovative task requires higher effort (the numerator on the left-hand side of (28)) relative to the value that it adds (the denominator).[21] In this case, the required share β of total proceeds Y will also have to be higher.

For small inefficiencies, the analysis in the proof of Proposition 3 delineates parameter ranges for which either delegation or non-delegation becomes optimal. Next we can further generalize the result that non-delegation becomes optimal for high Y, even as the inefficiencies are substantial. The key to this conclusion is the following observation. If the first-best decision rule is not feasible, then under delegation, the respective cutoff p_D^* is *independent* of Y. This follows immediately from the fact that the agent cares only about his incentive component β and not about the firm's overall payoff Y in the case of success. On the other hand, the higher is Y the smaller becomes the principal's own opportunism (or commitment) problem. Intuitively, this follows, as with a higher Y more is at stake for the firm if an inefficient decision is made.

Proposition 4. *As Y becomes sufficiently large, delegation is no longer optimal.*

Proof. We denote the respective inefficiencies by L_D and L_{ND}. We argue, first, that if for given $Y = Y'$ we have that $L_D > L_{ND}$, then this holds also for all higher values $Y > Y'$. To see this, note that from $L_D > L_{ND}$ we have that

$$\int_{p_D^*}^{\widehat{p}} (\widehat{p} - p) f(p)\, dp > \int_{\widehat{p}}^{p_{ND}^*} (\widehat{p} - p) f(p)\, dp.$$

Consequently, even if p_{ND}^* stayed constant (just as p_D^*), the change from Y' to Y would have a larger effect on L_D than on L_{ND}. The result is then further strenghtened by the observation that p_{ND}^* becomes, in fact, strictly lower, which formally follows from the comparative statics argument in Corollary 2, together with the fact that the left-hand side of (13) is strictly increasing in Y.

We show next that for sufficiently high Y, it must indeed hold that $L_D > L_{ND}$. We do this by deriving a somewhat stronger result, showing that $L_{ND} \to 0$ as $Y \to \infty$. For this note, first from (13) that $p_{ND}^* \to \widehat{p}$ as $Y \to \infty$. Consequently, we have from (6) that

$$\beta_{ND} \to \beta_\infty := \frac{c - w_m[1 - F(\widehat{p})]}{\int_{\underline{p}}^{\overline{p}} \phi(p) f(p)\, dp}.$$

From the decision rule in (3), we have next that

$$Y(p_{ND}^* - \widehat{p}) \to \beta_\infty \phi(\widehat{p}) - w_m = \frac{c\phi(\widehat{p}) - w_m \int_{\underline{p}}^{\overline{p}} [\phi(p) - \phi(\widehat{p})] f(p)\, dp}{\int_{\underline{p}}^{\overline{p}} \phi(p) f(p)\, dp}. \tag{29}$$

Thus, we have, finally, from (29) together with $p_{ND}^* \to \widehat{p}$ that indeed $L_{ND} \to 0$. □

In general, the optimal choice between the two organizational forms relies on a comparison of the respective profit levels (or inefficiencies) at the equilibrium cutoffs p_D^* and p_{ND}^*. As these also depend on local conditions of the respective functions, most notably $f(p)$, we can not obtain further general comparisons. Nevertheless, by considering some extreme cases, we can obtain further insights. For this, note that if either $w_m = 0$ or $\pi = 0.5$ applies, then the "overinvestment" problem under delegation is always extreme: $p^* = \underline{p}$. If the prior value of the innovation is now negative—i.e., if $E[p] < \widehat{p}$—then we can immediately conclude that delegation can never be optimal if either the performance measure is very noisy or if the agent's reservation (or market) wage is very low.[22]

Illustrative example: Take, again, the case in which $p \in [0, \overline{p}]$ is uniformly distributed. We set $\overline{p} = 1$ but leave $\widehat{p}$ unspecified. The crux is now that we can only solve explicitly for p_{ND}^* in the special case in which $\pi = 0.5$. Recall that, in this case, the "overinvestment" problem with delegation is extreme and results in $p_D^* = 0$. Without delegation, p_{ND}^* is the lowest value satisfying

$$Y(p_{ND}^* - \widehat{p})(1 - p_{ND}^*) = \frac{c}{Y}, \tag{30}$$

provided such a value $p_{ND}^* < 1$ exists. As p is uniformly distributed and as $\widehat{p} - p_D^* = \widehat{p}$, we have that the expected inefficiency is strictly smaller under delegation only if $p_{ND}^* - \widehat{p} > \widehat{p}$ and, thus, if

$$p_{ND}^* > 2\widehat{p}. \tag{31}$$

Solving (30), we obtain that

$$p_{ND}^* = \frac{1}{2}\left[(1 + \widehat{p}) - \sqrt{(1 + \widehat{p})^2 - 4\left(\widehat{p} + \frac{c}{Y}\right)}\right].$$

[21] Note that this is compatible with condition (14), which ensures that under the first-best decision rule, eliciting effort is indeed efficient.

[22] Of course, it is then still necessary that the agent can be incentivized without delegation.

Plugging this back into (31), we have the requirement that

$$\frac{c}{Y} < \widehat{p}(1 - 2\widehat{p}). \tag{32}$$

This confirms Proposition 3, where we generally showed that for sufficiently small inefficiencies, non-delegation is optimal only if c/Y remains sufficiently small. However, condition (32) also illustrates that the comparative statics in other parameters of the model may generally not be monotonic. According to (32), the threshold for c/Y for which non-delegation remains optimal is non-monotonic in $\widehat{p}$.[23]

6. Discussion

So far, our analysis has been restricted to a game in which, first, the firm offered a simple contract (α, β) and, second, the respective decision maker chose whether or not to implement the innovation. If the first best is not feasible, one may ask whether the firm could not devise a different mechanism. To discuss this question, it is convenient to address the two organizational regimes separately.

The case without delegation: If there is inefficient decision making in the case without delegation, then the innovation is not implemented for all $p \in [\widehat{p}, p^*)$, even though this would be efficient (and strictly so for all $p > \widehat{p}$). This bias is due to the fact that the agent's expected compensation, $\alpha + \beta\phi(p)$, exceeds his market wage w_m. If the agent accepted a small cut in his incentive component β, given that α is already equal to zero, then both sides could be made better off: As the principal implemented the innovation, the agent's expected compensation would still lie above w_m, while the principal also could pocket a share of the resulting efficiency gain $Y(p - \widehat{p})$.

To study the possibility of renegotiations, we stipulate that after the principal has observed p and announced his planned decision, the agent can preempt a possible rejection of the innovation by proposing a modified contract $(\widetilde{\alpha}, \widetilde{\beta})$. After accepting or rejecting this offer, the principal will, finally, either implement the innovation or not. Our key assumption is now that p is only privately observed by the decision maker, i.e., the principal.[24]

Clearly, if the principal truthfully announced whether or not given the initial contract he would want to implement the innovation, then the agent could make a mutually beneficial new proposal $(\widetilde{\alpha}, \widetilde{\beta})$. For instance, with $\widehat{\beta} := w_m/\phi(\widehat{p})$, any offer specifying $\widetilde{\alpha} = 0$ and $\widehat{\beta} < \widetilde{\beta} < \beta$ would be accepted with positive probability—i.e., for some set $p \in [\widetilde{p}^*, p^*)$—and would, in the case of acceptance, make both parties strictly better off.[25] The crux is, however, that when expecting such a new and more attractive offer, all "types" of the principal would want to pretend that they will subsequently reject the innovation—i.e., also "types" $p \geq p^*$. Hence, while the agent could indeed push down p^* by offering a wage cut, it is likely that this would only lead to a windfall gain for the principal without affecting his decision. If the agent still preferred such a revised offer $(\widetilde{\alpha}, \widetilde{\beta})$, however, then it is immediate that the initial contract (α, β) could not have been the principal's optimal commitment offer.

For the previous argument, we have considered only a renegotiation offer that adjusted the bonus β—i.e., from β down to some $\widetilde{\beta}$—while leaving the base wage constant at $\alpha = \widetilde{\alpha} = 0$. In order to "target" only "types" in $p \in [\widehat{p}, p^*)$ but not those in $p \geq p^*$, it may be asked whether the agent cannot make a different offer by adjusting both the base wage and the bonus. Indeed, as different "types" p have different relative preferences for an increase in the base wage or the bonus, it is indeed possible to sort between different types. This raises, more generally, the issue of whether a *menu* that is offered initially would not be more efficient. A menu of contracts would be given by $\{(\alpha_i, \beta_i)\}_{i \in I}$, where I is some index set. From this menu, the principal would then choose a contract $i \in I$ after observing p. The use of a non-degenerate menu is, however, strictly suboptimal. This follows from the following argument.

Any menu from which the principal would choose a different contract for different realizations of $p \geq p^*$ would, by revealed preferences, yield a higher expected payoff for the principal, compared to the case in which the offer is restricted to "marginal contract" only—i.e., the contract that the principal would pick at $p = p^*$. Exchanging the menu for this "marginal contract" would relax the agent's incentive constraint (6), while leaving the cutoff p^* unchanged. This makes it possible to adjust the contract by reducing either α (if this is feasible) or β until the agent's incentive constraint becomes again binding, given the newly obtained cutoff, which, as the principal's payoff from the innovation increases for all p, is strictly lower. The outcome is, thus, more efficient, such that, given that the agent's incentive constraint still holds with equality, the principal is strictly better off.

Proposition 5. *Consider the case without delegation. Then, the optimal contract from Proposition* 1 *will not be renegotiated in equilibrium. In addition, the offer from Proposition* 1 *also remains uniquely optimal when allowing the principal to offer menus.*

Proof. See Appendix.

[23] Note here that we clearly need $\widehat{p} < 0.5$ to make delegation feasible, given that from $E[p] = 0.5$, net profits equal $0.5 - \widehat{p} - c$.

[24] If the principal could make another offer, we would, thus, encounter a game of signaling. Though this would complicate the following argument, the main result would not be affected. (More formally, this follows also from our subsequent discussion of menus of offers.)

[25] To be precise, the offer would reverse the principal's decision for all $p \in [\widetilde{p}^*, p^*)$, where $\widetilde{p}^*$ solves (3) for the newly offered contract.

The case with delegation: Matters are more straightforward in the case with delegation. As long as the agent receives the (market) wage w_m when his innovation is not implemented, he will want to implement the innovation also for all $p \in [p^*, \widehat{p})$. Even if p were common knowledge at this stage, there cannot be any successful renegotiations of the (on-the-job) compensation (α, β). (The possibility to pay the agent more than w_m if the innovation is not implemented is discussed in Section 7.)

Furthermore, note that also in the case with delegation, it is still not possible to improve efficiency by initially specifying a menu of contracts, from which the agent now picks a contract after privately observing p. This could improve efficiency only if more of the agent's expected reward were thereby shifted into high p-states. As for high p, the agent prefers a steeper incentive scheme, for a non-degenerate menu incentive compatibility would require that some contract with $\alpha_i > 0$ was chosen for lower values of p. Including such a flatter contract in the menu would, however, not be optimal as it would push p^* further down instead of reducing the inefficiency $\widehat{p} - p^* > 0$.

Proposition 6. *Consider the case with delegation. Then, the optimal contract from Proposition* 2 *will not be renegotiated in equilibrium. In addition, the offer from Proposition* 2 *also remains uniquely optimal when allowing the principal to offer menus.*

Proof. See Appendix.

Private information: The preceding discussion of the case with delegation does not indicate that if p were observed by both sides, then there would not exist other mechanisms that would lead to the efficient outcome through either explicitly or implicitly making the agent's compensation depend on p.[26] The required flexibility may, however, add additional costs, such as transactions costs from negotiations. In addition, if the agent's superior were not the owner of the firm, additional agency costs could arise if the two sides conspire to increase the agent's compensation. These additional considerations are, however, clearly outside the considered model. Instead, we showed that if p is the decision maker's private information, then our previous contractual restrictions are without loss of generality.

As noted in the Introduction, p may be privately observed by the agent in the course of generating the innovation. Regardless of the firm's "formal organization," the firm would, thus, have to actually delegate the decision to the agent. In other cases, however, it may be the agent's superior or some other manager who possesses key information about the innovation's (market) potential. Furthermore, even if it were possible, in principle, to make p common knowledge, this may not always be practical—e.g., as observing or sharing knowledge of p may involve additional costs. For instance, managers further up the line would have to continuously spend time and effort to keep track of ongoing research so as to ultimately have the same information as the respective innovators.

7. Conclusion

This paper poses the question of which organizational and contractual choices can best ensure, first, that new ideas are generated and, second, that these ideas are implemented if and only if they are profitable. We found a key role for incentive pay. Both with and without delegation, incentive pay, which links the agent's reward for generating an innovation to the ultimate performance of the innovation, reduces the bias in decision making. This is achieved by rewarding the agent more when the innovation is highly profitable and less when the innovation is only moderately profitable. This, in turn, reduces the bias of the respective decision maker: The agent has fewer incentives to implement unprofitable innovations, which leads to less " overinvestment;" and the principal has more incentives to implement innovations, which leads to less "underinvestment."

By pointing out this novel role of incentive pay, we also qualify the results in Rotemberg and Saloner (1994). We show that rewarding an employee for a new innovation only through a lump-sum payment, as in their model, results in the worst bias for the respective decision maker—i.e., the employee with delegation or the firm without delegation. Steeper incentives that are linked to the success of the innovation mitigate this bias.

Throughout our analysis, we assumed that the agent is employed under an at-will contract. This implied that in order to sustain a compensation above his market wage, the agent had to make himself less indispensable, which was the case if his innovation was implemented. In the working paper version of this paper, we contrasted this with the case in which the firm could promise job security through contractually committing to "severance pay." This increased the agent's payoff without an innovation. Intuitively, while this can reduce the inefficiency in decision making, it reduces the agent's incentives to generate an innovation in the first place. In addition, compensation that pays agents an above-market wage even if they remain fully dispensable risks attracting employees who are not even able to perform the required innovative activity.

[26] A particularly simple mechanism that still allocates all "residual" decision rights to the agent would be one under which the agent's pay remains initially unspecified. Without an innovation, this still leads to a wage of w_m. If an innovation were generated, which the agent could then still unilaterally decide to implement, the respective wage must be negotiated. For all $p > \widehat{p}$, it remains to share the resulting efficiency gains $Y(p - \widehat{p})$. The agent's expected reward could be "fine-tuned" through the choice of the ensuing game of negotiations (e.g., the timing of the different moves as in Rubinstein and Wolinsky, 1992). (Cf. also, more generally, Maskin and Tirole, 1999 on the implementation with "unforeseen" but commonly observed contingencies.)

Acknowledgment

I am thankful to two referees, as well as to the associate editor for helpful comments.

Appendix A. Omitted proofs

Proof of Proposition 3. Take, first, a marginal change to $c > c_{FB}$. We want to derive conditions for when the resulting change in L_{ND} exceeds that in L_D. Next, note that p_D^* is always continuously differentiable. Though this does not hold for p_{ND}^*, it will become clear that it holds for $p_{ND}^* = \widehat{p}$.

With a continuously differentiable cutoff p^*, we have that

$$\frac{dL}{dc} = f(p^*)Y\left|\frac{dp^*}{dc}(p^* - \widehat{p})\right|.$$

Clearly, at $p^* = \widehat{p}$, this is zero by definition of $\widehat{p}$. To compare L_N and L_{ND} in an arbitrarily small neighborhood $c > c_{FB}$, we consider the second derivative, which evaluated at $p^* = \widehat{p}$ becomes

$$\left.\frac{d^2L}{dc^2}\right|_{p^*=\widehat{p}} = f(p^*)Y\left(\frac{dp^*}{dc}\right)^2.$$

Consequently, delegation is better in the neighborhood $c > c_{FB}$ if

$$-\left.\frac{dp_D^*}{dc}\right|_{p_D^*=\widehat{p}} < \left.\frac{dp_{ND}^*}{dc}\right|_{p_{ND}^*=\widehat{p}}. \tag{33}$$

Note now that we have from implicit differentiation of (25) that

$$\left.\frac{dp_D^*}{dc}\right|_{p_D^*=\widehat{p}} = -\frac{1}{w_m\left[\frac{\phi'(\widehat{p})}{[\phi(\widehat{p})]^2}\int_{\underline{p}}^{\overline{p}}\phi(p)f(p)\,dp\right]}.$$

Proceeding likewise for p_{ND}^* we have next that

$$\left.\frac{dp_{ND}^*}{dc}\right|_{p_{ND}^*=\widehat{p}} = \frac{1}{Y\int_{\widehat{p}}^{\overline{p}}\frac{\phi(p)}{\phi(\widehat{p})}f(p)\,dp - w_m\left[\frac{\phi'(\widehat{p})}{[\phi(\widehat{p})]^2}\int_{\underline{p}}^{\overline{p}}\phi(p)f(p)\,dp\right]}.$$

Consequently, (33) holds if $Y < 2w_m\phi'(\widehat{p})/\phi(\widehat{p})$. As we can substitute $\beta\phi(\widehat{p}) = w_m$ for both delegation and non-delegation, given that (10) just holds with equality, we have the condition $Y < 2\beta\phi'(\widehat{p})$. That is, we have that

$$\frac{\beta}{Y} > \frac{1}{2(2\pi - 1)},$$

which from (7) becomes (28). □

Proof of Proposition 5. It remains to prove that the use of (non-degenerate) menus is strictly suboptimal if a simple contract cannot achieve the first best. It is convenient to now denote an incentive-compatible menu by $\{(\alpha(p), \beta(p))\}$, with p being the "type" of the principal.[27] The respective utility of the principal is then

$$U(p) := \max\{\widehat{p}Y - w_m, Yp - [\alpha + \beta\phi(p)]\},$$

from which it follows again immediately that there exists a cutoff p^* such that the innovation is implemented only if $p \geq p^*$. (Recall that we restrict consideration to the case where $\underline{p} < \widehat{p}$.) Note next that the agent's incentive-compatibility constraint (6) can now be rewritten as

$$\int_{p^*}^{\overline{p}} [Yp - U(p)]f(p)\,dp \geq c. \tag{34}$$

We argue to a contradiction and assume that a degenerate menu is offered (while (10) does not hold). From incentive compatibility for the principal (or, likewise, his "revealed preferences"), this implies that the left-hand side of the agent's incentive-compatibility constraint (34) would *strictly* increase if instead of offering the menu, the principal offered only the "cutoff contract" $(\alpha(p^*), \beta(p^*))$, which would also implement the same cutoff p^*.

[27] It is straightforward to extend the argument to the case of menus in which the principal would (at least for some p) be indifferent between various contracts in the menu and, thus, possibly randomize at the interim stage.

We now adjust $(\alpha(p^*), \beta(p^*))$ by reducing the agent's compensation. For the following argument, it is irrelevant whether this is done by reducing $\alpha(p^*)$, whenever this is still feasible as $\alpha(p^*) > 0$, or by reducing $\beta(p^*)$. This continuously changes the left-hand side of the (at first slack) incentive-compatibility constraint (6), given that p^* also changes continuously.[28] Clearly, by making this adjustment sufficiently large, we arrive at a point where the incentive constraint (6) becomes binding again. Together with a more efficient cutoff, which is closer to but still higher than $\widehat{p}$ given that (10) does not hold, the firm's payoff must be strictly higher. □

Proof of Proposition 6. By the arguments in the main text, it remains to rule out the optimality of menus. We denote again by $(\alpha(p^*), \beta(p^*))$ the "cutoff contract," which is chosen at $p = p^*$. Consider next the following implications of incentive compatibility for the agent in case of a non-degenerate menu. We must have that $\alpha(p^*) > 0$, as for $\alpha(p^*) = 0$ it would not be incentive-compatible that $p = p^*$ prefers this contract while some higher type $p > p^*$ prefers another contract with $\alpha(p) > \alpha(p^*)$ (together with a lower value of $\beta(p) < \beta(p^*)$).

Consider now an alternative contract $(\widetilde{\alpha}, \widetilde{\beta})$ with $\widetilde{\alpha} = 0$ that implements the same cutoff p^*: $\widetilde{\beta} = \alpha(p^*)/\phi(p^*) + \beta(p^*)$. Importantly, from $\widetilde{\beta} > \beta(p^*)$ this implies that for all $p > p^*$, the agent's expected compensation is strictly higher with the new contract, which relaxes the incentive constraint (6). From there on, the argument is analogous to that in the proof of Proposition 5. That is, we can now decrease $\widetilde{\beta}$, thereby pushing up $p^* < \widehat{p}$, until the constraint becomes binding again. The resulting contract is then strictly more profitable for the firm than the original menu. □

References

Aghion, P., Tirole, J., 1997. Formal and real authority in organizations. Journal of Political Economy 105, 1–29.
Alonso, R., Matouschek, N., 2008. Optimal delegation. Review of Economic Studies 75, 259–293.
Dessein, W., 2002. Authority and communications in organizations. Review of Economic Studies 69, 811–838.
Holmstrom, B., 1989. Agency costs and innovation. Journal of Economic Behavior and Organization 12, 305–327.
Inderst, R., Klein, M., 2007. Innovation, endogenous overinvestment, and incentive pay. Rand Journal of Economics 38 (4), 881–904.
Jensen, M.C., 1993. The modern industrial revolution, exit, and the failure of internal control systems. Journal of Finance 48, 831–880.
Lerner, J., Wulf, J., 2007. Innovation and incentives: evidence from corporate R&D. Review of Economics and Statistics 89, 634–644.
Levitt, S.D., Snyder, C.M., 1997. Is no news bad news? Information transmission and the role of early warning in the principal-agent model. Rand Journal of Economics 28, 641–661.
Maskin, E., Tirole, J., 1999. Unforseen contingencies and incomplete contracts. Review of Economic Studies 66, 83–114.
Rotemberg, J., Saloner, G., 1994. Benefits of narrow business strategies. American Economic Review 84, 1330–1349.
Rubinstein, A., Wolinsky, A., 1992. Renegotiation-proof implementation and time preferences. American Economic Review 82, 600–614.
Wulf, J., 2007. Authority, risk, and performance incentives: evidence from division manager positions inside firms. Journal of Industrial Economics 55, 169–196.

[28] Note that as p^* decreases, the adjustment may, at points, even relax (6), which is, however, irrelevant for the argument.