(Im)patience is a Virtue: Time Preference and Incentives for Innovative Work*

Evan Sadler†

November 1, 2016

Abstract

Innovation benefits from low-powered incentives, yet many innovative firms provide very strong incentives to their employees. I offer an explanation based on a characteristic feature of creative work: dead ends. A principal hires an agent to work on a sequence of projects, each of which may succeed with some independent probability. The agent chooses whether to exert effort and when to abandon a project to start a new one. Switching projects incurs a delay. High rewards for success can slow innovation because the agent spends too much time on unpromising projects: the agent is reluctant to incur the cost of delay, particularly if he is impatient. If the principal is less patient than the agent, she offers high rewards anyway, and the agent's total compensation grows without bound as the value of innovations grows. If the principal is more patient than the agent, there is a bound on performance incentives and total compensation.

1 Introduction

Creative work is a central part of what many employees do in the modern economy. Whether researching new drugs, designing new products, or developing new trading strategies, a large share of workers engage in innovative activity on a routine basis. At the same time, various forms of profit-sharing are an increasingly common part of compensation packages. For instance, Lerner and Wulf (2007) note increasing use of restricted stock and long-term options to incentivize corporate research and development managers. Kruse et al. (2006) similarly document the ubiquity of "shared capitalism" across many industries. This paper asks what role such performance based compensation should play in incentivizing creative work.

*I thank Daron Acemoglu, Abhijit Banerjee, Ben Golub, Aditya Kuvalekar, Elliot Lipnowski, and Harry Pei for helpful conversations and suggestions.

†Harvard University – esadler@fas.harvard.edu

A central message in empirical studies of innovation and incentives is that innovation benefits from low-powered incentives. Extensive research, both in the field and the laboratory, finds that tolerance for early failure, long-run performance incentives, and ceding creative control are associated with higher levels of innovation. In a corporate context, long-term incentives for research and development heads (Lerner and Wulf, 2007), longer stock option vesting periods (Yanadori and Marler, 2006), golden parachutes (Francis et al., 2011), and a failure tolerant culture (Tian and Wang, 2014) are all associated with firms creating more patents and more heavily cited patents. Ederer and Manso (2013) present evidence from a laboratory experiment that long-term incentives and tolerance for failure help motivate innovation, and Azoulay et al. (2011) show how grants that are more failure tolerant, and give researchers more flexibility over what projects to pursue, increase the output of academic scientists.

Nevertheless, incentives in innovative organizations are highly varied in strength. At one extreme, the routine output of an academic researcher has little if any near-term effect on compensation—long-term effects come from improved job prospects and reputation, which are often difficult to tie to any single piece of work. At the other extreme, developers in a technology start up may receive the majority of their compensation through shared ownership of the firm, providing very strong incentives to perform. Incentives for most corporate researchers fall somewhere in between. If the evidence in favor of low-powered incentives is so clear, why do we sometimes observe high-powered incentives in practice?

I study incentives for innovation in a principal-agent setting with an explicit model of the innovative process. Dead ends are an important feature that distinguishes creative labor from other types of labor. A given project or idea might not work out, and workers must decide not just whether to exert effort, but also how long to keep trying before switching to a new idea. This leads to a tension between incentivizing effort and wasting time on projects with little chance of success. Time preferences determine how the principal and the agent balance the two.

An agent works on a sequence of projects, each of which is viable with some independent probability. A viable project produces a breakthrough at a Poisson rate, while an unviable project never produces. At any time, the agent may give up on the current project and incur a delay to come up with a new idea. The principal offers a combination of a fixed wage and shared ownership to compensate the agent. If the agent accepts the principal's offer, he is employed in perpetuity. I provide a full characterization of the agent's experimentation policy given the contract.

I first examine how the rate of innovation varies with the agent's ownership stake in the output. As long as the value of innovation is sufficiently high, there is a unique level of ownership that maximizes the rate of innovation. When the agent receives weaker incentives, he gives up too easily in favor of new ideas. When the agent receives stronger incentives, he spends too much time working on dead ends. Maximizing the rate of innovation entails capping the reward the agent receives for each success: as the value of each innovation increases, the rate maximizing ownership stake converges to zero. This result offers a novel explanation why low-powered incentives can encourage innovation, with a nuanced account of when this is the case. Low-powered incentives are most important when innovations are valuable and the agent is impatient.

In maximizing her expected discounted payoff, an impatient principal faces the same temptation as the agent to keep pursuing unviable projects because a breakthrough may come more quickly. If the principal is less patient than the agent, the optimal contract has no fixed wage. All of the agent's compensation comes through shared ownership, and the agent's total compensation grows without bound as the value of innovations increases. If the principal is more patient than the agent, there is an upper bound on the agent's total compensation. As the value of innovations increases, the agent's ownership stake converges to zero. Consequently, we should expect innovative firms with high discount rates—firms with a high cost of capital or high opportunity costs—to make heavier use of performance based compensation and to offer higher compensation overall.

Hitting dead ends is an inherent part of innovation that has important implications for the structure of incentives. This feature gives us a new rationale why low-powered incentives facilitate faster innovation while simultaneously explaining why some innovative organizations may optimally use high-powered incentives. Impatience can motivate an agent to work too hard on a given project because the delay to come up with a new idea is costly. On the other hand, this means that impatient workers are less expensive to incentivize, which suggests that innovative firms may favor hiring impatient employees. Accounting for the effects of dead ends can also inform policies to encourage innovation. For instance, a patient planner designing a patent system may choose to limit protections not just because of positive externalities but because doing so can lead to faster innovation. If firms are impatient relative to the planner, limiting the payoff to innovation can incentivize more efficient experimentation from the planner's perspective.

1.1 Related Work

Economic theory offers two main explanations why low-powered incentives help encourage innovation. First, strong incentives may crowd out an agent's intrinsic motivation. Work in psychology documents that extrinsic rewards can impair an individual's intrinsic motivation to perform a task (Deci et al., 1999), and intrinsic motivation seems particularly important in creative pursuits. Bénabou and Tirole (2003) develop a principal-agent model in which this crowding out effect can appear. They require two conditions for incentives to discourage the agent. First, the principal must have some private information about either the task or the agent's ability. Second, the principal must be inclined to offer incentives when the agent has low ability or the task is difficult. In this case, strong incentives lead the agent to infer that success is unlikely, and the agent is reluctant to undertake similar tasks in the future.

A second explanation stems from an agent's aversion to uncertainty. In many cases, an agent can choose whether to innovate or to rely on current knowledge. Innovating entails a chance of failure, so rewards that are contingent upon success make pursuing innovation more risky for the agent. Manso (2011) formalizes this idea using a two-period model in which an agent can shirk, exploit an action with known payoffs, or explore an action with unknown payoffs. The optimal contract that induces exploration exhibits tolerance for early failure and rewards long-term success—the author explicitly ties this structure to common features of managerial compensation contracts like stock options and golden parachutes. In a related model, Byun (2015) shows that ambiguity about the likelihood of successful innovation exacerbates this effect. When an agent can choose whether or not to innovate, or can choose between radical or incremental innovation, strong performance incentives discourage the agent from choosing the riskier option.

To my knowledge, this is the first paper to suggest that dead ends and the need to backtrack can also explain why low-powered incentives encourage innovation. Research success often requires trying several different approaches, and there is a tension between continuing on the current path and coming up with a new idea. Impatience makes delay costly, and an agent may be tempted to keep working on an unpromising idea, hoping for a more immediate success. Even when there is no crowding out of intrinsic motivation, and no trade-off between exploration and exploitation, this tension gives cause to limit incentives. Moreover, this mechanism can furnish a reason why some firms use high-powered incentives despite the negative impact on innovation.

2 The Innovation Process

A principal manages a single agent, who is tasked with creating and implementing innovative projects. Time is continuous, and I represent a project as an exponential bandit. A new project is viable with independent probability p_0. While working on a project, at each instant the agent exerts effort $x \in \{0, 1\}$, incurring a flow cost $c > 0$ for effort 1. If the project is viable, and the agent invests effort, a breakthrough arrives at a Poisson rate $\lambda > 0$. If the project is unviable, or the agent makes no effort, a breakthrough never arrives. At any point, the agent may choose to abandon the current project and start a new project. Starting a new project incurs a time delay that follows an exponential distribution with parameter γ. After a breakthrough, the agent similarly incurs a delay before starting on the next project. The principal observes neither effort nor the decision to change projects.

The principal values each breakthrough at $y > 0$ and compensates the agent through a combination of a fixed wage and shared ownership. The agent receives the payment $w \geq 0$ upon hiring and a fraction $\alpha \in [0, 1]$ of an innovation's value when a breakthrough occurs. The agent discounts the future at a rate r, and the principal discounts the future at a rate $\tilde{r}$. The agent has an outside option equivalent to an immediate payment $\underline{w} > 0$. The principal offers a pair (w, α) to the agent, who either accepts or rejects. If the agent accepts, he chooses an experimentation policy to maximize his expected discounted payoff. The principal offers a contract to maximize her expected discounted profit.

2.1 Remarks on the Model

There are two components to the agent's ability. A higher λ means the agent is more efficient at implementing an existing idea—the agent is more productive. In contrast, a higher γ means the agent is more adept of coming up with new ideas—the agent is more creative. Both attributes are valuable, but they have different implications for the structure of incentives.

The agent incurs the flow cost c only when making effort on an existing project, not when attempting to come up with a new idea. The qualitative results would remain unchanged if we instead had separate flow costs $c_p > c_c$ for productive work and creative work respectively. In essence, I assume that the agent enjoys creative labor more than working through an existing project. We can solve the problem using the same techniques under the opposite assumption, but a number of comparative statics would reverse.

The contract (w, α) is a simple way to capture the combination of wages and profit-sharing that is so prevalent in employee compensation. A higher value of α translates to higher contingent payments and thus higher powered incentives. The analysis centers on how α affects the rate of innovation and how the principal chooses what α to offer. A separate question, which I do not address, is whether the structure (w, α) is optimal or approximately optimal in any sense. While these contracts are appealingly simple, it seems likely that the principal could do better with a more flexible compensation scheme. However, having distinct discount rates for principal and agent makes the problem ill-defined for a fully general payment structure. For instance, if the principal is less patient, she can always do better by offering payments further in the future. We avoid such problems with the given contractual form.

In many settings, we can sell the output to the agent to obtain an efficient level of effort. The assumption $w \geq 0$ precludes this possibility. For large scale research and development projects, especially projects that depend on significant capital investments, it is clear that selling the project to an agent is typically infeasible. While this feature is absent in many smaller scale innovative projects, such projects often involve relationship and asset specific innovation that makes selling the project equally prohibitive. A lot of innovative work builds on existing products, or is particular to a specific client's needs, and cannot reasonably be separated and sold to an employee carrying out the work.

3 The Agent's Problem

This section studies the agent's optimal behavior after accepting a contract (w, α). A few observations simplify our analysis. First, the agent's effort and project switching policy depends only on contingent payments—the value αy is all that matters for the agent's decision. Furthermore, if α is high enough to ever induce effort, the agent always makes effort. When the agent refrains from effort, it delays any contingent payoff without yielding information about project viability. Hence, if the agent can obtain a positive payoff from making effort, any such delay is strictly suboptimal. Consequently, the agent's policy amounts to choosing when to cease working on the current project.

While the agent works without a breakthrough, his belief p about the project's viability steadily decreases. We can characterize the optimal strategy via a threshold belief $\underline{p} \leq p_0$ at which the agent abandons the project. If we write $u(p)$ for the agent's expected continuation payoff given current belief p, we recognize this as a modified version of the exponential bandit problem of Keller et al. (2005). The key difference is that the value of the "safe" arm is determined endogenously as the value of incuring delay to start a new project.

The main result of this section characterizes the threshold belief $\underline{p}$ at which the agent quits the current project. For notational convenience, I define

$$\rho = \frac{\underline{p}(1-p_0)}{p_0(1-\underline{p})}, \tag{1}$$

and let T denote the amount of time the agent spends on a project before giving up. Note there is a one-to-one correspondence between ρ, T, and $\underline{p}$.

Theorem 1. *The agent works if and only if $\alpha y > \frac{c}{\lambda p_0}$. In this case, the equation*

$$\frac{\gamma}{r}\rho\left(1 - \frac{\lambda\rho^{\frac{r}{\lambda}}}{\lambda+r}\right) + \rho\frac{\lambda\alpha y - c}{\lambda\alpha y(1-p_0)} = \frac{\gamma}{\lambda+r} + \frac{c}{p_0\lambda\alpha y} \tag{2}$$

uniquely characterizes the belief at which the agent abandons a project.

Proof. The proof entails solving a Bellman equation for $u(p)$ and explicitly computing $u(p_0)$ to obtain the endpoints. See Appendix for details. □

The belief p is a convenient state variable to characterize the agent's decision, but the time an agent spends working on a project is typically easier to observe. Consequently, the most useful comparative statics results describe how the time T, rather than the belief $\underline{p}$, changes with model parameters.

Proposition 1. *The agent's strategy satisfies:*

(a) The time T is decreasing in c and increasing in α and y.

(b) The time T is increasing in p_0 and decreasing in γ.

(c) The time T is increasing in r.

Proof. We derive each of these results by implicitly differentiating (2) with respect to the variable in question. Begining with the cost c, we have

$$\frac{\partial\rho}{\partial c}\left(\frac{\gamma}{r} + \frac{\lambda\alpha y - c}{\lambda\alpha y(1-p_0)} - \frac{\gamma}{r}\rho^{\frac{r}{\lambda}}\right) - \frac{\rho}{\lambda\alpha y(1-p_0)} = \frac{1}{p_0\lambda\alpha y}.$$

Solving gives

$$\frac{\partial\rho}{\partial c} = \frac{r\left(p_0\rho + 1 - p_0\right)}{rp_0(\lambda\alpha y - c) + \gamma\lambda\alpha y p_0(1-p_0)(1-\rho^{\frac{r}{\lambda}})} > 0,$$

which implies ρ is increasing in c, so T is decreasing. For α we have

$$\frac{\partial\rho}{\partial\alpha}\left(\frac{\gamma}{r}(1-\rho^{\frac{r}{\lambda}}) + \frac{\lambda\alpha y - c}{\lambda\alpha y(1-p_0)}\right) + \rho\frac{c}{\lambda\alpha^2 y(1-p_0)} = -\frac{c}{p_0\lambda\alpha^2 y},$$

giving

$$\frac{\partial \rho}{\partial \alpha} = \frac{-cr\left(p_0\rho + 1 - p_0\right)}{p_0 r\alpha(\lambda\alpha y - c) + \gamma\lambda\alpha^2 y p_0(1-p_0)(1-\rho^{\frac{r}{\lambda}})} < 0. \tag{3}$$

The calculation for $\frac{\partial \rho}{\partial y}$ is essentially identical, proving part (a).

Proceeding similarly, we compute

$$\frac{\partial \rho}{\partial p_0} = \frac{-r\left((1-p_0)^2 + \rho p_0^2(\lambda\alpha y - c)\right)}{r p_0^2(1-p_0)(\lambda\alpha y - c) + \gamma\lambda\alpha y p_0^2(1-p_0^2)(1-\rho^{\frac{r}{\lambda}})},$$

$$\frac{\partial \rho}{\partial \gamma} = \frac{\lambda\alpha y(1-p_0)\left(r + \lambda\rho^{1+\frac{r}{\lambda}} - (\lambda+r)\rho\right)}{r(\lambda+r)(\lambda\alpha y - c) + \gamma\lambda\alpha y(\lambda+r)(1-p_0)(1-\rho^{\frac{r}{\lambda}})}, \text{ and}$$

$$\frac{\partial \rho}{\partial r} = -\gamma\frac{r^2 - \rho(\lambda+r)^2 + \lambda\rho^{1+\frac{r}{\lambda}}(\lambda+2r) - r(\lambda+r)\rho^{1+\frac{r}{\lambda}}\ln\rho}{\gamma r(\lambda+r)^2(1-\rho^{\frac{r}{\lambda}}) + r^2(\lambda+r)^2\frac{\lambda\alpha y - c}{\lambda\alpha y(1-p_0)}}.$$

Proving the remaining claims is equivalent to showing that $\frac{\partial \rho}{\partial p_0} < 0$, that $\frac{\partial \rho}{\partial \gamma} > 0$, and that $\frac{\partial \rho}{\partial r} < 0$. The derivative $\frac{\partial \rho}{\partial p_0}$ is clearly negative as $\lambda\alpha y > \frac{c}{p_0} \geq c$. To see that $\frac{\partial \rho}{\partial \gamma}$ is positive, observe that the expression

$$r + \lambda\rho^{1+\frac{r}{\lambda}} - (\lambda+r)\rho$$

is decreasing in ρ whenever $\rho < 1$, taking the value zero exactly when $\rho = 1$. Hence, the numerator is positive, and the derivative is positive. Finally, to see that $\frac{\partial \rho}{\partial r}$ is negative, we show that

$$r^2 - \rho(\lambda+r)^2 + \lambda\rho^{1+\frac{r}{\lambda}}(\lambda+2r) - r(\lambda+r)\rho^{1+\frac{r}{\lambda}}\ln\rho$$

is positive for $\rho \in (0,1)$. Differentiating this expression with respect to ρ gives

$$(\lambda+r)^2\left(\rho^{\frac{r}{\lambda}}\left(1 - \frac{r}{\lambda}\ln\rho\right) - 1\right),$$

and differentiating a second time yields

$$-\frac{r^2}{\lambda^2}(\lambda+r)^2\rho^{\frac{r}{\lambda}-1}\ln\rho > 0.$$

The expression and its first derivative are both zero at $\rho = 1$. Since the second derivative is positive, the first derivative is negative on $(0,1)$, and hence the function is positive on $(0,1)$ as desired.

□

Parts (a) and (b) of Proposition 1 are fairly intuitive. Think of T as a measure of how hard the agent works. If you lower the cost of working, or increase the benefit, the agent works harder. If a project is viable with higher probability, the agent finds it worthwhile to work a little longer. If the agent is more creative and comes up with new ideas faster, changing projects is less costly, and he spends less time on any one project. A similar exercise for λ shows that ρ is decreasing in λ—if the agent is more efficient, he works to a lower belief threshold $\underline{p}$—but the time this takes may be lower or higher depending on other parameters.

Part (c) is more subtle. Switching projects is relatively costly for an impatient agent because the delay entails a larger discount to his reward. As a result, the impatient agent will work harder for lower α than a more patient agent. While larger contingent payments induce the agent to work more, the agent's patience limits how far this can go. From (2), we see that as αy approaches infinity, or as c approaches 0, the value of ρ approaches the solution to

$$\frac{\gamma}{r}\rho\left(1-\frac{\lambda\rho^{\frac{r}{\lambda}}}{\lambda+r}\right)+\frac{\rho}{1-p_0}=\frac{\gamma}{\lambda+r}. \tag{4}$$

Write $\overline{\rho}$ for this limiting value. The limit $\overline{\rho}$ follows the same comparative statics as in Proposition 1, implying there are levels of effort that only very impatient agents will ever exert.

Equation (4) plays an important role in our later analysis of the principal's problem. Since the principal does not incur the cost of effort, the analogous equation replacing r with $\tilde{r}$ describes the experimentation policy that is optimal from the principal's perspective. An implication is that whether the principal can induce the agent to adopt her optimal experimentation policy depends on the relative time preferences of the two players.

4 The Rate of Innovation

This section demonstrates how low-powered incentives can lead to faster innovation. We first find the experimentation policy that maximizes the rate of innovation. Building from the comparative statics in Proposition 1, we show that for αy sufficiently high, the agent works harder than this policy prescribes. When contingent payments are large, and the agent is impatient, the agent spends too much time working on dead ends because he finds more value in the small chance of a quick success. Limiting the agent's reward can therefore increase the rate of innovation, particularly when r is high and innovations are very valuable.

Write μ for the expected time that passes between two breakthroughs. If T is the time spent on an idea before giving up, then we can compute μ recursively as

$$\begin{aligned}\mu &= \frac{1}{\gamma}+p_0\int_0^T \lambda t e^{-\lambda t}\mathrm{d}t+\left(1-p_0(1-e^{-\lambda T})\right)(T+\mu)\\ &= \frac{1}{\gamma}+\frac{p_0}{\lambda}(1-e^{-\lambda T})+(1-p_0)T+\left(1-p_0(1-e^{-\lambda T})\right)\mu.\end{aligned}$$

Solving yields

$$\mu=\frac{\lambda+p_0\gamma(1-e^{-\lambda T})+\gamma\lambda(1-p_0)T}{p_0\gamma\lambda(1-e^{-\lambda T})}=1+\frac{\lambda-\gamma(1-p_0)\ln\rho}{p_0\gamma\lambda(1-\rho)}.$$

This is a convex function of ρ on $(0,1)$. Taking first order conditions, we find a unique minumum characterized by

$$\gamma(1-p_0)\left(1-\rho(1-\ln\rho)\right)-\lambda\rho=0. \tag{5}$$

Write ρ^* for the solution that minimizes μ and hence maximizes the rate of innovation.

Theorem 2. *The value $\overline{\rho}$ is strictly decreasing in r with*

$$\lim_{r\to 0}\overline{\rho} = \rho^*.$$

Proof. Taking $\alpha y \to \infty$ in part (c) of Proposition 1 implies $\overline{\rho}$ is decreasing in r. We can rewrite (4) as

$$\gamma(1-p_0)\left[1-\rho\left(1+\frac{\lambda}{r}(1-\rho^{\frac{r}{\lambda}})\right)\right]-(\lambda+r)\rho = 0.$$

Taking the limit as $r \to 0$ yields (5). □

Together with Proposition 1, Theorem 2 offers a detailed account of how the rate of innovation varies with the agent's incentives. Fixing α and r, the value of μ is a U-shaped function of y. For low-value innovations, the contingent payment induces little effort, and the agent spends most of his time trying to come up with fresh ideas. For high-value innovations, the agent works too hard on each individual project, wasting effort on dead ends. Fixing the reward αy, the value μ is a similarly U-shaped function of r. When the agent is patient, he makes less effort and too readily gives up on each project. When the agent is impatient, the cost of delay motivatives him to work too hard.

Maximizing the rate of innovation often requires low-powered incentives, particularly when innovations are valuable or the agent is impatient. Here, we capture the strength of incentives through the agent's ownership stake α, but the same logic applies to other performance dependent consequences. For instance, we could obtain similar comparative statics if the agent earned a fixed wage but was at risk of getting fired for failure to perform.[1] Low-powered incentives in this case correspond to increased tolerance for failure, which studies show encourages innovation.

The impact of dead ends on the rate of innovation suggests broader policy implications. Patent protections are one way that governments calibrate the gains that firms derive from investments in innovation. The positive externalities from free dissemination and use of new discoveries are one reason why governments limit protections. Our results offer a novel rationale for the same: by encouraging a firm to try new lines of research, limiting the reward may actually spur faster innovation. Based on these findings, a patent system designed to maximize the rate of innovation would provide weaker protections for more valuable discoveries.

5 The Principal's Problem

The principal's impatience affects the extent to which her goals align with maximizing the rate of innovation. Even when low-powered incentives clearly encourage faster innovation, the principal may prefer high-powered incentives if $\tilde{r}$ is high. This provides a rationale for

[1] Deadline effects would make the analysis far more complicated, but we can use the same intuition.

some firms to use strong performance incentives despite slower innovation. As we see later, this also implies disparities in the total compensation offered to comparably skilled workers at different firms.

The principal offers the contract (w, α) to maximize her expected profit, given the agent's optimal response. The principal pays w upfront, and at each breakthrough, she earns $(1-\alpha)y$. The agent's outside option constrains the principal to offer an expected payoff of at least $\underline{w}$. Our first result characterizes the principal's profit as a function of the contract.

Proposition 2. *If the agent accepts the contract (w, α), the principal earns*

$$(1-\alpha)y\frac{\lambda p_0(\gamma+\tilde{r})(1-\rho^{1+\frac{\tilde{r}}{\lambda}})}{(\lambda+\tilde{r})\left(\tilde{r}+\gamma(1-p_0)(1-\rho^{\frac{\tilde{r}}{\lambda}})\right)+\gamma p_0\tilde{r}(1-\rho^{1+\frac{\tilde{r}}{\lambda}})} - w, \tag{6}$$

where ρ is given by (2).

Proof. Let $\pi(w, \alpha)$ denote the principal's profit, and let $\tilde{\pi}(\alpha) = \pi(w, \alpha) + w$ denote the profit when we ignore the fixed payment. Following a calculation similar to that for the agent's utility, we have

$$\begin{aligned}
\tilde{\pi}(\alpha) &= p_0\left((1-\alpha)y + \frac{\gamma}{\gamma+\tilde{r}}\tilde{\pi}(\alpha)\right)\int_0^T \lambda e^{-(\tilde{r}+\lambda)t}\mathrm{d}t \\
&\quad + (1-p_0+p_0e^{-\lambda T})e^{-\tilde{r}T}\frac{\gamma}{\gamma+\tilde{r}}\tilde{\pi}(\alpha) \\
&= \left(p_0\frac{\lambda+re^{-(\tilde{r}+\lambda)T}}{\lambda+\tilde{r}} + (1-p_0)e^{-\tilde{r}T}\right)\frac{\gamma}{\gamma+\tilde{r}}\tilde{\pi}(\alpha) \\
&\quad + p_0(1-\alpha)y\frac{\lambda}{\lambda+\tilde{r}}(1-e^{-(\tilde{r}+\lambda)T})
\end{aligned}$$

Solving this yields

$$\tilde{\pi}(\alpha) = \frac{\lambda p_0(1-\alpha)y(\gamma+\tilde{r})(1-\rho^{1+\frac{\tilde{r}}{\lambda}})}{(\lambda+\tilde{r})\left(\tilde{r}+\gamma(1-p_0)(1-\rho^{\frac{\tilde{r}}{\lambda}})\right)+\gamma p_0\tilde{r}(1-\rho^{1+\frac{\tilde{r}}{\lambda}})},$$

and (6) follows. □

Using this characterization of the principal's profit, we can derive the main result on the optimal level of profit sharing. The contract structure is highly sensitive to the discount rates of the principal and the agent. If the principal is less patient than the agent, then the agent's entire compensation always comes from shared ownership, and the agent's total compensation grows without bound as the value of innovations increases. In this case, the agent can earn far more than his outside option when y is large. If the principal is more patient than the agent, then the ownership stake α converges to zero as y grows, and there is a finite upper bound on the agent's total compensation. In many cases, the agent is never paid more than the outside option, no matter how valuable innovations are.

Theorem 3. *Suppose $\tilde{r} \geq r$. Then $w = 0$, and the performance reward αy increases without bound in y.*

Suppose $\tilde{r} < r$. Then there exists an upper bound on the performance reward αy, and there is an upper bound on the agent's total compensation.

Proof. We compute the total derivative of the function $\tilde{\pi}(\alpha)$ from the proof of Proposition 2:

$$\tilde{\pi}'(\alpha) = -\frac{\tilde{\pi}(\alpha)}{1-\alpha} + \frac{\partial \tilde{\pi}}{\partial \rho}\frac{\partial \rho}{\partial \alpha}.$$

The first term is the cost of paying higher rewards for each breakthrough, and the second term is the effect of giving the agent stronger incentives. Looking at the second term in more detail, the derivative $\frac{\partial \rho}{\partial \alpha} < 0$ is given in (3). The other derivative takes the form

$$\frac{\partial \pi}{\partial \rho} = -A(1-\alpha)y\left[\frac{\gamma}{\tilde{r}}\rho\left(1 - \frac{\lambda \rho^{\frac{\tilde{r}}{\lambda}}}{\lambda + \tilde{r}}\right) + \frac{\rho}{1-p_0} - \frac{\gamma}{\lambda + \tilde{r}}\right],$$

where A is a positive function of the parameters, excluding y and α. Notice that the bracketed term mirrors equation (4). There is a unique $\tilde{\rho}$ such that the derivative is zero, and the derivative is positive if and only if $\rho > \tilde{\rho}$. Recall the notation $\overline{\rho}$ for the solution to (4). The comparative statics in part (c) of Proposition 1 apply, implying that $\overline{\rho} \geq \tilde{\rho}$ if $\tilde{r} \geq r$, and $\overline{\rho} < \tilde{\rho}$ if $\tilde{r} < r$.

Suppose $\tilde{r} \geq r$. This means that $\frac{\partial \pi}{\partial \rho}\frac{\partial \rho}{\partial \alpha}$ is positive at any value of ρ the principal can induce the agent to choose. Since the principal is less patient than the agent, it is cheaper to provide incentives by offering payment in the future, and there is a benefit from providing stronger performance incentives. Therefore, there is no reason to offer a fixed payment w.

Moreover, suppose we let y increase to ∞ while holding the reward αy constant. The values $\tilde{\pi}(\alpha)$ and $\frac{\partial \tilde{\pi}}{\partial \rho}$ both grow proportionally to y, while the term $\frac{\partial \rho}{\partial \alpha}$ grows as $\frac{1}{\alpha}$. This implies that for sufficiently high y, the derivative $\tilde{\pi}'(\alpha)$ is positive, and this is true for any value of αy. Consequently, the agent's total compensation will grow without bound in y.

Suppose $\tilde{r} < r$. This means that $\overline{\rho} < \tilde{\rho}$, and there exists $\overline{w} < \infty$ such that $\alpha y = \overline{w}$ induces the agent to choose $\tilde{\rho}$. The principal never sets $\alpha > \frac{\overline{w}}{y}$ because π' is negative past this point. Suppose the contract $\left(0, \frac{\overline{w}}{y}\right)$ amounts to less than the agent's reservation wage. It should be clear that the principal will never offer more than the contract $\left(w^*, \frac{\overline{w}}{y}\right)$ with w^* chosen to make the contract equivalent to the reservation wage: this contract induces the best possible effort from the agent with the lowest cost to achieve that effort. If the contract $\left(0, \frac{\overline{w}}{y}\right)$ offers more than the agent's reservation wage, then the principal will never incur more than the cost of this contract—again, it obtains the best possible effort for the lowest cost. As an exercise, one can show that in the optimal profit sharing arrangement, the principal always offers somewhat less than $\alpha = \frac{\overline{w}}{y}$, making up the difference with a fixed payment if necessary, because the difference in discount rates makes it relatively cheap to compensate the agent upfront. $\square$

Theorem 3 establishes a connection between time preferences and the use of performance incentives for innovative work. A patient principal offers bounded performance rewards and may just match the agent's outside option. In contrast, an impatient principal offers stronger performance rewards, and potentially much higher compensation. We should expect the highest paid individuals to work for companies with high discount rates and to derive much of their compensation from profit sharing rather than a fixed salary. This prediction largely agrees with compensation patterns we observe in financial services and in technology start ups. Similarly skilled workers in more patient organizations—like universities—should expect lower overall wages that are far less sensitive to performance.

Similarly, patient agents can command higher performance rewards and higher compensation than their less patient peers. This suggests that innovative organzations may prefer employees with high discount rates: Proposition 1 implies we can obtain any given level of effort more cheaply from a less patient agent. Acemoglu et al. (2016) present evidence that firms with young managers are more innovative, stressing that a culture open to disruption plays a key role both in promoting innovation and allowing young managers to rise in the ranks. To the extent that young people are less patient, as suggested in some studies (Harrison et al., 2002; Bishai, 2004), the lower cost of providing incentives offers an alternative account of this pattern.

6 Remarks

Dead ends offer a new rationale for why low-powered incentives encourage creative work. The choice between continuing on a current project and incuring delay to start a new one is fundamentally distinct from the more classical exploration-exploitation tradeoff. In the latter case, strong incentives discourage risk taking for fear of failure. Here, strong incentives lead the agent to spend more time on projects that are more likely to fail: the agent is reluctant to switch projects due to the cost of delay. Limiting contingent payments encourages the agent to give up on unpromising ideas and try something new.

Since the cost of delay plays a prominent role in both the agent's and the principal's tradeoff, incentive schemes depend crucially on time preferences. Impatient agents are more responsive to incentives: they work harder for less. This suggests that time preferences may impact the matching between firms and workers, with innovative firms hiring less patient workers on average. A similar logic implies that impatient firms should offer employees more generous compensation that is more closely tied to performance. Dead ends can explain why some firms offer strong incentives even though it leads to less innovation in the long run: an impatient firm values the small chance of an immediate payoff more than a sustained high rate of innovation.

The findings suggest a number of broader patterns for innovative output. Holding all else equal, patient organizations should work on a larger number of distinct research projects and should reach breakthroughs more quickly than less patient organizations. We should also expect the most valuable breakthroughs to occur relatively slowly because researchers spend more time on dead ends. These patterns in turn suggest policy prescriptions to encourage innovation. A patent policy that treats discoveries of different values differently—offering less protection for more valuable inventions—may lead to faster innovation.

One open question is whether the combination of fixed wages and profit sharing constitutes an optimal contract for innovative work. I focus on this class of contracts as an approximation to commonly used incentive schemes, but it seems plausible that a more general contract could improve the principal's profit. Of course, real employment contracts are also more intricate than those studied here. In particular, the use of promotions and "up or out" policies in many firms complicates an analysis of incentives for innovation. It seems likely that such policies are related to screening employees based on ability and compensating investments in human capital, but they may well affect effort provision and project selection in innovative work.

References

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik (2016), "Young, Restless, and Creative: Openness to Disruption and Creative Innovations." Working Paper.

Azoulay, Pierre, Joshua Graff Zivin, and Gustavo Manso (2011), "Incentives and Creativity: Evidence from the Academic Life Sciences." *RAND Journal of Economics*, 42, 527–554.

Bénabou, Roland and Jean Tirole (2003), "Intrinsic and Extrinsic Motivation." *Review of Economic Studies*, 70, 489–520.

Bishai, David (2004), "Does Time Preference Change with Age?" *Journal of Population Economics*, 17, 583–602.

Byun, Seong (2015), "Incentivizing Innovation Under Ambiguity." Working Paper.

Deci, Edward, Richard Koestner, and Richard Ryan (1999), "A Meta-Analytic Review of Experiments Examining the Effects of Extrinsic Rewards on Intrinsic Motivation." *Psychological Bulletin*, 125, 627–668.

Ederer, Florian and Gustavo Manso (2013), "Is Pay-for-Performance Detrimental to Innovation?" *Management Science*, 59, 1496–1513.

Francis, Bill, Iftekhar Hasan, and Zenu Sharma (2011), "Incentives and Innovation: Evidence from CEO Compensation Contracts." Working Paper.

Harrison, Glenn, Morten Lau, and Melonie Williams (2002), "Estimating Individual Discount Rates in Denmark: A Field Experiment." *American Economic Review*, 92, 1606–1617.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005), "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73, 39–68.

Kruse, Douglas, Joseph Blasi, and Rhokeun Park (2006), "Shared Capitalism in the U.S. Economy: Prevalence, Characteristics, and Employee Views of Financial Participation in Enterprises." In *Shared Capitalism at Work: Employee Ownership, Profit and Gain Sharing, and Broad-based Stock Options* (Douglas Kruse, Richard Freeman, and Joseph Blasi, eds.), University of Chicago Press.

Lerner, Josh and Julie Wulf (2007), "Innovation and Incentives: Evidence from Corporate R&D." *Review of Economics and Statistics*, 89, 634–644.

Manso, Gustavo (2011), "Motivating Innovation." *Journal of Finance*, 66, 1823–1860.

Tian, Xuan and Tracy Wang (2014), "Tolerance for Failure and Corporate Innovation." *Review of Financial Studies*, 27, 211–255.

Yanadori, Yoshio and Janet Marler (2006), "Compensation Strategy: Does Business Strategy Influence Compensation in High-Technology Firms?" *Strategic Management Journal*, 27, 559–570.

A Appendix

Proof of Theorem 1

Taking the ownership stake α as given, we can derive a Bellman equation to characterize the agent's behvaior. Write $u(p)$ for the agent's continuation payoff when currently working on a project that is viable with probability p and τ for the exponentially distributed delay. Since the agent always has the option to incur delay and start a new project, we have for all p that

$$u(p) \geq \mathbb{E}\left[e^{-r\tau}\right] u(p_0) = \frac{\gamma}{\gamma + r} u(p_0).$$

This will hold with equality at the abandonment belief $\underline{p}$.

During a "period" $\mathrm{d}t$, the agent makes a breakthrough with probability $\lambda \mathrm{d}t$ if the project is viable, and with probability zero otherwise. In the absence of a breakthrough, the agent updates her belief to

$$p + \mathrm{d}p = \frac{p(1 - \lambda \mathrm{d}t)}{1 - p + p(1 - \lambda \mathrm{d}t)}.$$

Equivalently, the belief change is $\mathrm{d}p = -\lambda p(1-p)\mathrm{d}t$.

At the same time, the agent incurs a flow cost $c\mathrm{d}t$ of effort. The value function therefore satisfies

$$u(p) = \max\left\{\frac{\gamma}{\gamma + r} u(p_0),\ -c\mathrm{d}t + e^{-r\mathrm{d}t}\mathbb{E}[u(p + \mathrm{d}p) \,|\, p]\right\}.$$

The agent makes a breakthrough with probability $\lambda p \mathrm{d}t$, yielding a continuation payoff of $\alpha y + \frac{\gamma}{\gamma+r} u(p_0)$. Otherwise, the continuation payoff is $u(p) + u'(p)\mathrm{d}p$. Taking $e^{-r\mathrm{d}t} \approx 1 - r\mathrm{d}t$, we arrive at the Bellman equation

$$u(p) = \max\left\{\frac{\gamma}{\gamma + r} u(p_0),\ \frac{1}{\lambda p + r}\left(\lambda p\left[\alpha y + \frac{\gamma}{\gamma + r} u(p_0) - (1-p)u'(p)\right] - c\right)\right\}.$$

Observe that the flow payoff can be positive if and only if $\lambda p_0 \alpha y > c$, implying the first claim.

As long as the agent continues on the current project, the value function satisfies the differential equation

$$(1-p)u'(p) + \frac{r + \lambda p}{\lambda p} u(p) = \alpha y + \frac{\gamma}{\gamma + r} u(p_0) - \frac{c}{\lambda p}.$$

The solutions of this differential equation are of the form

$$u(p) = \frac{\lambda p}{\lambda + r}\left(\alpha y + \frac{\gamma}{\gamma + r} u(p_0) + \frac{c}{r}\right) - \frac{c}{r} + C(1-p)\left(\frac{1-p}{p}\right)^{\frac{r}{\lambda}},$$

where C is an arbitrary constant.

To find C and the abandonment belief $\underline{p}$, we impose the necessary boundary conditions $u(\underline{p}) = \frac{\gamma}{\gamma+r} u(p_0)$ and $u'(\underline{p}) = 0$. We compute the derivative

$$u'(p) = \frac{\lambda}{\lambda + r}\left(\alpha y + \frac{\gamma}{\gamma + r} u(p_0) + \frac{c}{r}\right) - C\frac{r + \lambda p}{\lambda p}\left(\frac{1-p}{p}\right)^{\frac{r}{\lambda}}.$$

The condition $u'(\underline{p}) = 0$ gives

$$C = \left(\frac{\lambda}{\lambda + r}\right)\left(\frac{\lambda \underline{p}}{\lambda \underline{p} + r}\right)\left(\alpha y + \frac{\gamma}{\gamma + r} u(p_0) + \frac{c}{r}\right)\left(\frac{\underline{p}}{1-\underline{p}}\right)^{\frac{r}{\lambda}}.$$

Substituting this into $u(\underline{p}) = \frac{\gamma}{\gamma+r} u(p_0)$ gives

$$\begin{aligned}\frac{\gamma}{\gamma + r} u(p_0) = {} & \frac{\lambda \underline{p}}{\lambda + r}\left(\alpha y + \frac{\gamma}{\gamma + r} u(p_0) + \frac{c}{r}\right) - \frac{c}{r} \\ & + \left(\frac{\lambda}{\lambda + r}\right)\left(\frac{\lambda \underline{p}(1-\underline{p})}{\lambda \underline{p} + r}\right)\left(\alpha y + \frac{\gamma}{\gamma + r} u(p_0) + \frac{c}{r}\right).\end{aligned}$$

Multiplying through by $\lambda \underline{p} + r$, we find the quadratic term drops out, and we can solve for

$$\underline{p} = \frac{1}{\lambda \alpha y}\left(\frac{\gamma r}{\gamma + r} u(p_0) + c\right).$$

To complete the proof, we must compute $u(p_0)$. The expected payoff is the probability of finding a breakthrough in the current project times the discounted payoff from a breakthrough, plus the probability of not finding a breakthrough times the continuation payoff from starting over, less the accumulated flow cost of effort. Let T denote the amount of time the agent spends on any research project. From Bayes' rule, we have $e^{-\lambda T} = \frac{\underline{p}(1-p_0)}{p_0(1-\underline{p})} = \rho$. We calculate $u(p_0)$ as

$$\begin{aligned}u(p_0) &= p_0\left(\alpha y + \frac{\gamma}{\gamma + r} u(p_0)\right)\int_0^T \lambda e^{-(r+\lambda)t} \mathrm{d}t + \left(1 - p_0 + p_0 e^{-\lambda T}\right) e^{-rT}\frac{\gamma}{\gamma + r} u(p_0) \\ &\quad - p_0 \int_0^T \lambda e^{-\lambda t}\left(\frac{c}{r}(1 - e^{-rt})\right)\mathrm{d}t - \left(1 - p_0 + p_0 e^{-\lambda T}\right)\frac{c}{r}(1 - e^{-rT}) \\ &= \left(p_0 \frac{\lambda + r e^{-(r+\lambda)T}}{\lambda + r} + (1-p_0)e^{-rT}\right)\frac{\gamma}{\gamma + r} u(p_0) + p_0 \frac{\lambda \alpha y}{\lambda + r}\left(1 - e^{-(r+\lambda)T}\right) \\ &\quad - \frac{c}{r}\left(\frac{p_0 r}{\lambda + r}(1 - e^{-(r+\lambda)T}) + (1-p_0)(1 - e^{-rT})\right)\end{aligned}$$

Solving and substituting ρ for $e^{-\lambda T}$ gives

$$\begin{aligned}u(p_0) &= \frac{\lambda \alpha y p_0(\gamma + r)(1 - \rho^{1+\frac{r}{\lambda}}) - c p_0(\gamma + r)(1 - \rho^{1+\frac{r}{\lambda}}) - \frac{c}{r}(1-p_0)(\lambda + r)(\gamma + r)(1 - \rho^{\frac{r}{\lambda}})}{(\lambda + r)(\gamma + r) - p_0\gamma(\lambda + r\rho^{1+\frac{r}{\lambda}}) - \gamma(1-p_0)(\lambda + r)\rho^{\frac{r}{\lambda}}} \\ &= \frac{\lambda \alpha y p_0(\gamma + r)(1 - \rho^{1+\frac{r}{\lambda}}) - c(\gamma + r)\left(p_0(1 - \rho^{1+\frac{r}{\lambda}}) + \frac{\lambda + r}{r}(1-p_0)(1 - \rho^{\frac{r}{\lambda}})\right)}{(\lambda + r)\left(r + \gamma(1-p_0)(1 - \rho^{\frac{r}{\lambda}})\right) + \gamma p_0 r(1 - \rho^{1+\frac{r}{\lambda}})}\end{aligned}$$

Substituting this into $\underline{p} = \frac{1}{\lambda\alpha y}\left(\frac{\gamma r}{\gamma+r}u(p_0)+c\right)$ and simplifying gives

$$\underline{p} = \frac{\gamma p_0 r(1-\rho^{1+\frac{r}{\lambda}}) + \frac{c}{\lambda\alpha y}r(\lambda+r)}{(\lambda+r)\left(r+\gamma(1-p_0)(1-\rho^{\frac{r}{\lambda}})\right)+\gamma p_0 r(1-\rho^{1+\frac{r}{\lambda}})}.$$

From the definition of ρ we have

$$\frac{\underline{p}}{1-\underline{p}} = \frac{p_0}{1-p_0}\rho = \frac{\gamma p_0 r(1-\rho^{1+\frac{r}{\lambda}}) + \frac{c}{\lambda\alpha y}r(\lambda+r)}{\gamma(\lambda+r)(1-p_0)(1-\rho^{\frac{r}{\lambda}})+r(\lambda+r)\left(1-\frac{c}{\lambda\alpha y}\right)}.$$

Multiplying through by the denominator and simplifying gives our result:

$$\frac{\gamma}{r}\rho\left(1-\frac{\lambda\rho^{\frac{r}{\lambda}}}{\lambda+r}\right)+\rho\frac{\lambda\alpha y-c}{\lambda\alpha y(1-p_0)} = \frac{\gamma}{\lambda+r}+\frac{c}{p_0\lambda\alpha y}.$$

If we differentiate the left hand side with respect to ρ, we obtain

$$\frac{\gamma}{r}\left(1-\rho^{\frac{r}{\lambda}}\right)+\frac{\lambda\alpha y-c}{(1-p_0)\lambda\alpha y} > 0.$$

Therefore, the left hand side is strictly increasing in ρ. For $\rho \in (0,1)$, the left hand side ranges from zero to

$$\frac{\gamma}{\lambda+r}+\frac{\lambda\alpha y-c}{\lambda\alpha y(1-p_0)} > \frac{\gamma}{\lambda+r}+\frac{\frac{c}{p_0}-c}{\lambda\alpha y(1-p_0)} = \frac{\gamma}{\lambda+r}+\frac{c}{p_0\lambda\alpha y},$$

implying there is a unique solution. $\square$