

MPRA

Munich Personal RePEc Archive

Dynamic Evaluation Design

Smolin, Alex

University of Bonn

October 2017

Online at <https://mpra.ub.uni-muenchen.de/84133/>

MPRA Paper No. 84133, posted 23 Jan 2018 15:27 UTC

Dynamic Evaluation Design*

Alex Smolin

October 12, 2017

Abstract

A principal owns a firm, hires an agent of uncertain productivity, and designs a dynamic policy for evaluating his performance. The agent observes ongoing evaluations and decides when to quit. While not quitting, the agent is paid a wage proportional to his perceived productivity; the principal claims the residual performance. After quitting, the agent secures a fixed safe payoff. I show that equilibrium evaluation policies are Pareto efficient and leave no rents to the agent. In a minimally informative equilibrium, for a broad class of performance technologies, the agent's wage deterministically grows with tenure.

Keywords: evaluation, information design, career concerns, bandit experimentation, downward wage rigidity, up-or-out

JEL Codes: C72, D82, D83, M52

*Institute for Microeconomics, University of Bonn, alexey.v.smolin@gmail.com. This paper builds on a chapter of my dissertation written at Yale University that has been previously circulated as “Optimal Feedback Design.” I thank Dirk Bergemann, Johannes Hörner, and Larry Samuelson for support and guidance throughout this project. I am grateful to Florian Ederer, Sergei Izmalkov, Emir Kamenica, Anton Kolotilin, Chiara Margaria, Benny Moldovanu, Pauli Murto, Anne-Katrin Roesler, and Áron Tóbiás for insightful conversations. Finally, I am thankful to participants of research seminars at Berlin, Bocconi, Bonn, Cornell Johnson, HSE (Moscow), Indiana Kelley, MPI Bonn, and Yale, as well as at the 5th World Congress of the Game Theory Society and the 2016 North American Summer Meeting of the Econometric Society.

1 Introduction

Evaluation of individual performance is an important part of organizational life. Although much evaluation is informal, most organizations have formal evaluation policies designed to collect and distribute performance information to employees.¹ As communication and information technologies advance, many companies find it easier to provide more evaluation. As a recent example, in August 2015 General Electric (GE) announced an ongoing shift from its legacy system of annual performance reviews to more frequent conversations between managers and employees via an online application.² This way, GE joined other high-profile companies such as Microsoft, Accenture, and Adobe in a move towards more frequent, exhaustive, and real-time evaluation. However, whenever adopting new evaluation policies the companies should ask: What is their effect on the overall performance? Would other evaluation policies do better? Ultimately, which evaluation policy is the best for the company?

In this paper, I develop a framework to analyze the design of evaluation policies. I consider a principal who owns a firm and hires an agent to work over time. The agent's productivity, his *type*, is initially uncertain to both players but affects the agent's ongoing performance. The performance is not directly observed but can be revealed through evaluations. While at firm, the agent is paid a wage proportional to his expected productivity; the principal claims the residual performance. In every period, the agent evaluates his career prospects and decides whether to quit. If the agent quits, he stops performing and secures an exogenously fixed safe payoff. Both players are risk neutral and discount the future at the same rate.

The principal designs and adopts a dynamic evaluation policy. The policy is a sequence of statistical experiments informative about past performance. The experiments can vary in what and when performance is assessed. The evaluation is costless but its design should take into account the agent's incentives. On one hand, the promise of future evaluations motivates the agent to stay at the firm and learn whether he is able to perform well. On the other hand, any evaluation may turn out negative and persuade the agent to quit. As a result, a partially informative policy can outperform both no-evaluation and complete-evaluation policies. I discuss it in detail in the Section 3 example.

In Section 4, I investigate equilibrium evaluation policies by developing an efficiency argument. First, I show that the design problem can be viewed as a dynamic persuasion problem in a bandit experimentation setting. Doing so simplifies the problem and allows me

¹According to [Murphy and Cleveland \(1995\)](#), between 74% and 89% of business organizations had formal performance appraisal policies by 1995.

²“GE’s Real-Time Performance Development,” *Harvard Business Review*, August 12, 2015, <https://hbr.org/2015/08/ges-real-time-performance-development>.

to characterize the set of feasible payoffs that can be possibly achieved in the relationship. Second, I study the set of implementable payoffs, achievable by some evaluation policy and the agent’s best response to it. I show this set includes all Pareto efficient payoffs that deliver the agent at least his safe option. I conclude that any equilibrium evaluation policy is efficient and leaves the agent with no rents. Furthermore, I characterize a minimally informative equilibrium policy. This policy turns out to be simple—it informs the agent whether he would quit if he could fully observe past performance but valued his outside option less.

In Section 5, I study effects of optimal evaluation on the agent’s career and wage dynamics. I observe several qualitative properties that hold in the minimally informative equilibrium. First, the agent’s wage is a deterministic function of tenure. Second, for a broad class of performance technologies, the agent’s wage monotonically increases with tenure. The increase reflects the ongoing positive selection and the corresponding growth of expected productivity. Third, the shape of the wage profile depends on performance technology. If the technology is coarse, so that performance comes as a stream of infrequent successes, then the wage increases at the revision dates, spaced sparsely over the agent’s career. In contrast, if the technology is detailed, so that performance can always reveal the agent’s incompetency, then the performance is constantly monitored and the wage strictly grows with tenure.

I discuss concrete implications of my findings for organizational behavior in Section 6. First, my analysis suggests that objective evaluation policies may be important in explaining economic dynamics commonly observed within firms. These include lack of wage variation within same-tenure cohorts, downward wage rigidity, and up-or-out contracts. Second, I provide a framework to assess existing evaluation practices. My results speak in favor of retrospective evaluation and provide a rationale for rating compression and leniency bias as techniques to maintain workforce morale. Finally, I highlight the important commitment role human resource departments may play in implementing optimal evaluation policies.

Related Literature My paper contributes to the literature on dynamic persuasion and information design built from static models of [Rayo and Segal \(2010\)](#) and [Kamenica and Gentzkow \(2011\)](#). [Orlov \(2016\)](#) studies the joint design of performance evaluations and monetary contracts when the agent exerts private efforts. [Renault, Solan, and Vieille \(2017\)](#) and [Ely \(2017\)](#) study dynamic persuasion with exogenous information flow and a myopic agent. [Orlov, Skrzypacz, and Zryumov \(2017\)](#) investigate a setting in which the principal lacks commitment between different periods. In a closely related paper, [Ely and Szydlowski \(2017\)](#) study a setting in which the relevant state deterministically evolves over time, resulting in a predictable reversal of incentives. They apply an argument analogous to mine and

show an optimal persuasion policy is efficient and leaves the agent with no rents.

My framework highlights the interplay between career concerns and performance evaluations. In my model, the agent’s quitting decision is publicly observable. In contrast, several papers investigate evaluation effects on private efforts in a framework of [Holmström \(1999\)](#). [Hansen \(2013\)](#) and [Rodina \(2016\)](#) study static incentives with a respective focus on partition evaluations and comparative statics. [Hörner and Lambert \(2016\)](#) study dynamic incentives with a focus on Gaussian policies.

Similar incentive effects are present in multistage contests and tournaments. [Ederer \(2010\)](#) compares effectiveness of complete- and no-evaluation policies in two-stage tournaments. [Halac, Kartik, and Liu \(2016\)](#) study an optimal design of general multistage contests and similarly focus on the extreme evaluation policies within each period. Nevertheless, [Goltsman and Mukherjee \(2011\)](#) highlight optimal evaluation policies in tournaments are generally partially informative.

Finally, my paper contributes to the literature on dynamic contracts without transfers. [Guo \(2016\)](#) studies dynamic delegation when the agent is privately informed. [Hörner and Guo \(2015\)](#) study dynamic resource allocation when the agent’s private information evolves over time. My paper highlights that information control may complement delegation and action control as a powerful management tool.

2 Model

A principal owns a firm and hires an agent. The relationship takes place in consecutive periods $t = 0, 1, 2, \dots$. At time 0, the agent’s productivity $\theta \in \Theta \subseteq \mathbb{R}^+$ is drawn according to a cumulative distribution G_0 . The productivity is fixed throughout the relationship and is not directly observed by either principal or agent. The players are symmetrically informed about the productivity with the prior expectation of productivity, $\mathbb{E}\theta$, being equal to θ_0 .

Performance The productivity affects the agent’s performance at firm $y_t \in Y \subseteq \mathbb{R}$. Conditional on productivity, performance is independently and identically distributed across periods according to a cumulative distribution F_θ . The collection of distributions $\{F_\theta\}_{\theta \in \Theta}$ defines production capabilities of the firm and is called (*performance*) *technology*. Without loss of generality, the productivity is defined to equal expected performance:

$$\mathbb{E}[y_t | \theta] = \theta. \tag{1}$$

It follows that performance is informative about the agent’s productivity: consistently

higher performance suggests higher productivity. The overall informativeness and details of the learning process are determined by technology. If distributions F_θ have distinct supports for all θ , then a single performance outcome fully reveals productivity. In contrast, if supports for all θ coincide, then, generally, the productivity cannot be learned with certainty in any finite time. I put no assumptions on technology for the characterization of equilibrium payoffs in Section 4. I will impose a regularity assumption in Section 5 to establish downward wage rigidity. Performance is not directly observed by either party but can be revealed through evaluations as discussed below.

Strategies The principal can publicly reveal past performance through an *evaluation policy* she designs. The policy is costless and governs when and what performance information is publicly available. The evaluations are objective; their outcomes cannot be manipulated by the principal. At the same time, I place no restrictions on which evaluations the principal can conduct. That is, she can conduct complete evaluation, no evaluation, periodic reviews, grade evaluations, and so forth.

Formally, the principal chooses an evaluation policy m among all stochastic processes measurable with respect to past performance and evaluations.³ A policy can be represented by a sequence of random messages $\{m_t\}_{t=0}^\infty$ that are sent to the agent,⁴

$$m_t : Y^{t-1} \times M^{t-1} \rightarrow \Delta(M). \quad (2)$$

The message space M can be freely chosen by the principal. For concreteness, I let it consist of all finite length messages composed of Latin letters and numbers and include a zero length message \emptyset . The exact message labels do not matter because their meaning is determined solely by the law of m . The evaluation policy determines the process of learning about the agent's productivity at the firm. Denote the set of all possible evaluation policies of the form (2) by \mathcal{M} .

The concept of an evaluation policy is an extension of [Kamenica and Gentzkow \(2011\)](#)'s [Kamenica and Gentzkow \(2011\)](#) static persuasion policy and admits two possible interpretations. First, it can be viewed as a disclosure policy. In this interpretation, the principal constantly monitors the performance but is bound to communicate according to the policy chosen at the beginning of the game. Second, it can be viewed as a sequence of public experiments. In this interpretation, the principal does not observe performance directly but

³Randomization over several evaluation policies can be represented by a single evaluation policy with a combined evaluation law.

⁴I adopt a convention that for any stochastic process x its time- t realization is denoted by subscript x_t and the history up to time t , $\{x_s\}_{s \leq t}$, is denoted by superscript x^t .

commits to a sequential policy of public tests to inform both players on past performance. These two interpretations are equivalent in my setting because the principal has no private use for information.

It might be helpful to think about the evaluations being conducted by the human resource department of a firm. In this case, a realization m_t corresponds to an outcome of a particular evaluation. The evaluation policy m corresponds to the operating rules of the department and specifies which and how past performance is evaluated in any given period. Complete evaluation policy corresponds to $m_t \equiv y_{t-1}$, no evaluation policy to $m_t \equiv \emptyset$, periodic reviews to $m_t = y^{t-1}$ for some selected dates, and grade evaluations to $m_t = i$ whenever y_t occurs in the i -th element of some partition of Y .

Faced with the evaluation policy, the agent chooses whether and when to quit the firm. He decides based on past evaluations, which he correctly interprets according to Bayes' rule. The quitting is irreversible and terminates any interaction between the agent and the firm. The agent can quit at time 0.

Formally, the agent chooses a *quitting time* τ , which is a stopping time measurable with respect to the evaluation policy m :

$$\tau \text{ is a stopping time w.r.t. } m_0, m_1, \dots \quad (3)$$

If $\tau \equiv 0$, then the agent quits at time 0 and does not generate any performance. If $\tau \equiv \infty$, then the agent stays at the firm forever, irrespectively of past evaluations. Denote the set of all possible quitting times by \mathcal{T} .

Payoffs As long as the agent stays at the firm, the principal appropriates the performance outcomes and pays the agent a wage w_t . I assume the wage is proportional to the agent's expected performance,

$$w_t(m^t) = \alpha \mathbb{E}[y_t | m^t], \quad (4)$$

with the proportionality coefficient $\alpha \in (0, 1)$. Effectively, α is the share of the expected output appropriated by the agent and $1 - \alpha$ is the share appropriated by the firm. The interiority assumption $\alpha \in (0, 1)$ ensures that both players extract some surplus from the relationship. The relationship captures, in the simplest form, the reputation effects of performance evaluations. The agent wants to receive positive evaluations to be perceived as more productive because, in this case, he will be offered a higher wage. The relationship may be viewed as a reduced form outcome of competition within the firm's industry. Alternatively, it might be set by a regulator and specified in a contract.

As soon as the agent quits the firm, he secures a total payoff of $V \geq 0$. This payoff is

exogenous, commonly known, and fixed throughout the relationship. This is the agent's safe option. It can be alternatively interpreted as his (opportunity) cost of staying at the firm.

Both players are risk neutral and discount the future with a common discount factor δ . For given evaluation policy $m \in \mathcal{M}$ and quitting strategy $\tau \in \mathcal{T}$, the normalized expected payoffs of the parties are

$$U^P(m, \tau) = \mathbb{E}_{m, \tau} \left[(1 - \delta) \sum_{t=0}^{\tau-1} \delta^t (y_t - w_t) \right], \quad (5)$$

$$U^A(m, \tau) = \mathbb{E}_{m, \tau} \left[(1 - \delta) \sum_{t=0}^{\tau-1} \delta^t w_t + \delta^\tau V \right]. \quad (6)$$

Note that evaluation policy plays two implicit roles in the payoffs. First, it shapes the evolution of the agent's wage. Second, it provides the agent with information that guides his quitting decision.

Equilibrium I study perfect Bayesian equilibria of this game. For a given evaluation policy, the agent chooses a quitting time to maximize his total expected payoff. The strategy choice fully captures the agent's sequential rationality because his problem is time consistent. The principal anticipates the agent's best response and designs the evaluation policy to maximize her expected payoffs.

Definition 1. An evaluation policy m^* and a quitting time τ^* constitute an equilibrium if they solve the problem:

$$\max_{m \in \mathcal{M}, \tau \in \mathcal{T}} U^P(m, \tau), \quad (7)$$

$$\text{s.t. } \tau \in \arg \max_{\tau \in \mathcal{T}} U^A(m, \tau). \quad (8)$$

My goal is to characterize an equilibrium evaluation policy and payoffs in this game. This characterization further allows me to study equilibrium wage dynamics. To this end, for given strategies m and τ , define an *observed wage* W_t as the agent's wage conditional on staying at the firm,

$$W_t = w_t \mid \tau > t. \quad (9)$$

The observed wage at time t is, generally, a random variable because the agent can possibly stay at the firm with a wide range of past evaluations. In addition, define a wage profile as a collection of observed wages at different times $W = \{W_t\}_{t=0}^\infty$. The wage profile captures the dynamics of the agent's wage throughout his career at the firm. By definition, if the agent always quits before time t , then $W_t = \emptyset$.

In the next section, I illustrate the setting and discuss the main economic forces in a simple example. In Section 4, I show that an equilibrium exists and characterize the equilibrium payoffs. In Section 5, I study equilibrium wage dynamics and establish conditions under which the equilibrium wage profile is non-decreasing.

3 Binary Example

To gain more intuition about the setting, consider the following example. The performance is binary, low or high, $Y = \{y^L, y^H\}$. Let $y^L = 0$ and $y^H = 1$ so that only high performance is valuable, to which I refer as a “success.” There are two possible types, $\Theta = \{\theta^L, \theta^H\}$, and each type is equally likely. By the definition (1), a type is equal to expected productivity which in this example coincides with a probability of success, $\Pr(y_t = y^H | \theta) = \theta$. Further, assume that only the high type is productive, $\theta^L = 0$ and $\theta^H = 1/2$. The following table summarizes the firm’s technology:

$f_\theta(y)$	y^L	y^H	
θ^L	1	0	.
θ^H	1/2	1/2	

Let the players split the surplus equally, $\alpha = 1/2$, and the agent’s safe option be $V = 1/5$. Finally, let the discount factor be $\delta \simeq 0.97$.⁵

No evaluation First, consider the case in which the principal adopts a no-evaluation policy, $m_t \equiv \emptyset$. In this case, the same evaluation message is sent irrespective of past performance and thus is completely uninformative. The firm effectively provides no feedback. As a result, the wage is equal to half of a prior expected productivity:

$$w_t \equiv w_0 = \alpha \theta_0 = \frac{1}{8}.$$

The agent’s best-response is straightforward. Whenever staying at the firm, he receives a flow payoff of $w_0 = 1/8$ and foregoes the opportunity flow of $V = 1/5$. Because $w_0 < V$, the agent quits at time 0.

As a result, in the absence of informative evaluations, the wage profile is empty and the firm is effectively not operating. The players’ payoffs are $(U^A, U^P) = (0.2, 0)$.

⁵The exact value required for a clean demonstration is $2/1365^{1/10}$.

Complete evaluation Now, consider the case in which the principal adopts a complete-evaluation policy, $m_t \equiv y_{t-1}$. In this case, each evaluation fully reveals the performance of the last period. The firm effectively collects all relevant performance information and timely provides it to the agent.

Under complete evaluation, the agent's wage depends on past evaluations. The wage starts at half of the prior productivity, $w_0 = 1/8$. As long as no successes occur, the wage gradually decreases according to Bayes' rule, $w_t^L = \frac{1}{2(1+2^t)}$. If a success occurs, the wage jumps to half of the expected productivity of a high type, $w^H \equiv 1/4$, and stays there forever. The single success reveals that the agent's type is high and guarantees a steady stream of future successes. The transition to one of these two wages ensures that the wage is a martingale. Note that this is a general feature that holds under any evaluation policy because the wage is proportional to expectations which are martingales by Bayes' rule.

Faced with these career prospects, the agent optimally quits in the first period when his wage drops below a cutoff wage \hat{w} . The cutoff depends on the discount factor. Higher δ translates into lower \hat{w} , because career concerns are more important, and the agent is willing to try to achieve a success for a longer time. For the considered discount factor the cutoff can be calculated to be anywhere in the region $(1/20, 1/12)$. The wage drops below \hat{w} if no success occurred at period $\hat{T} = 2$. Hence, the agent optimally works until time \hat{T} and continues working if and only if evaluations revealed a success in the past.

It follows that under complete evaluation the wage profile is random. Viewing the agent as a representative employee, one out of many independent draws, would result in two distinct features. First, there would be a cross-sectional variation of employees' wage in periods before \hat{T} : some employees are proven to be high types and some still try to achieve a success. Second, the wage of a given employee is likely to decrease during his career in the firm. The wage profile is illustrated in Figure 1.

The players payoffs are $(U^A, U^P) \simeq (0.21, 0.1)$. Note that $U^A > V$, so the agent strictly benefits from working at the firm as compared to his outside option.

Equilibrium evaluation Now, consider a partial evaluation policy, namely, a single revision policy, $m_T = y^{T-1}$ and $m_t = \emptyset$ for $t \neq T$. Under this policy, the principle provides no informative evaluations before or after the single revision date T when she conducts a full evaluation of past performance.

Under the single revision policy, the agent's wage at the firm exhibits a single jump at the revision date. No informative evaluations are conducted before the revision time. During this period, the agent's wage stands still at the prior level, $w_0 = 1/8$. At the revision time, the full evaluation of past performance is conducted. If a success is revealed, the agent

is proven to be of a high type, and his wage jumps up to $w_T^H = 1/4$. If no successes are found, then the productivity expectation drops, as does the wage, to $w_T^L = \frac{1}{2(1+2^T)}$. After the revision time, no informative evaluations are conducted, so the wage stays at the revision level forever.

Given the career prospects, the agent's optimal quitting time depends on the revision date and can take only values 0 or T . Indeed, in periods after the revision, $t \geq T$, the agent is essentially facing a no-evaluation policy; his wage is fixed. Consequently, because $w_T^H > V > w_T^L$, the agent immediately quits at time T if no success is revealed; otherwise, he stays at the firm forever. Before the revision, $t < T$, the agent's continuation payoff weakly increases in t as the revision time gets closer. Hence, the agent's quitting problem reduces to a binary choice: to either quit at time 0 or stay until the revision time and quit only if no successes are revealed.

If the revision date equals the complete evaluation time $T = \hat{T}$, then the agent prefers to stay towards the revision because this strategy delivers him payoff greater than V . In fact, this strategy delivers the same payoff as the best response to a complete evaluation policy. This is not a coincidence. Both strategies induce the same joint distribution of productivity and quitting time. They differ only in their wage profiles: before T , the wage is fixed under a revision policy and is random under a complete evaluation policy. However, because wage is a martingale, the agent receives, on average, the same wage, and, hence, the same payoff. Similarly, these strategies deliver the same payoff to the principal.

The principal can induce the agent to generate more surplus by postponing the revision until time $T > \hat{T}$. However, there is a limit on how late the revision can be performed because the agent may prefer to quit at time 0. A maximal incentive compatible revision time can be calculated to be $T^* = 10$.

In fact, the revision policy with the revision time T^* is optimal for the principal and, hence, is an equilibrium evaluation policy. Indeed, it delivers payoffs $U^{A^*} = V = 0.2$, $U^{P^*} \simeq 0.12$. Moreover, these payoffs are Pareto efficient because the agent never quits when successful. Because the agent can guarantee his safe option by quitting at time 0, the principal cannot achieve payoffs above U^{P^*} , and the result follows. The equilibrium wage profile is illustrated in Figure 1. Notably, the resulting equilibrium wage profile is deterministic and weakly increasing. Effectively, it exhibits an up-or-out pattern with downward wage rigidity.

I show in the next two sections that the main features of this example are general. First, the evaluation policy affects payoffs only through its effect on quitting time and not on the wage. Second, equilibrium payoffs are Pareto efficient, and the agent is left with no rents at his safe payoff. Finally, I show the equilibrium wage profile is always weakly increasing if




Figure 1: Wage profile under complete evaluation policy (left) and equilibrium evaluation policy (right).

there are only two states or if the technology satisfies general regularity assumptions.

4 Equilibrium Analysis

Equilibrium characterization requires finding an evaluation policy optimal for the principal which is a dynamic information disclosure problem. Such problems are known to be difficult due to their dynamic structure and multidimensionality. The characterization is further complicated because the information that can be disclosed is gradually and endogenously generated, and the agent is forward looking. Because of these features, the existing techniques of [Kremer, Mansour, and Perry \(2014\)](#) and [Ely \(2017\)](#) cannot be applied.

Instead, I develop and use an efficiency approach. First, I characterize the set of Pareto-efficient payoffs (Sections 4.1, 4.2). Second, I provide an upper bound on the principal's equilibrium payoffs. Finally, I show the upper bound can be achieved with a particular information policy (Section 4.3).

4.1 Payoff Transformation

Before characterizing the set of feasible and Pareto-efficient payoffs, it is useful to investigate the players' payoffs. Intuitively, because both players are risk neutral, the mean-preserving spread of a wage should not affect their payoffs in any period. Hence, receiving a fixed wage proportional to the expected performance at the beginning of a period should be payoff equivalent to receiving a bonus proportional to performance at the end of a period .

This intuition can be formalized. Applying the law of iterated expectations and the optional stopping theorem, the player's payoffs can be written solely in terms of performance and a safe option.

Lemma 1. (Payoff Transformation) *The players' payoffs can be written as*

$$U^P(m, \tau) = \mathbb{E}_{m, \tau} \left[(1 - \delta) \sum_{t=0}^{\tau-1} \delta^t (1 - \alpha) y_t \right], \quad (10)$$

$$U^A(m, \tau) = \mathbb{E}_{m, \tau} \left[(1 - \delta) \sum_{t=0}^{\tau-1} \delta^t \alpha y_t + \delta^\tau V \right]. \quad (11)$$

Lemma 1 implies the current setting is strategically equivalent in a sense of [Thompson \(1952\)](#) to the setting of persuasion of bandit experimentation.⁶ They have the same set of strategies and each strategy pair delivers the same payoffs to both players. Consequently, the two settings have the same feasible payoffs and equilibria. In this alternative setting, the agent sequentially pulls the arm of a slot machine and decides when to stop. Pulling the arm generates a stochastic reward, which is proportionally split between the agent and the principal. The reward depends on the machine's type and is not observed by the agent. The principal designs what reward information the agent observes to maximize her own payoffs.

Further, the payoff representation (10) and (11) highlights the nature of the conflict between the players. The principal does not have a safe option. She achieves maximal payoffs when the agent generates maximal expected surplus in the relationship. Because all types are positive, $\Theta \subseteq \mathbb{R}^+$, it happens when the agent never quits. In contrast, the agent has a safe option. He would like to quit when his continuation payoff falls below it. Clearly, then, his payoffs are maximized when the principal provides complete evaluation because it allows him to make the most informed quitting decision.

Proposition 1. (Preferred Strategies) *The principal's payoff is maximized if the agent never quits, $\tau \equiv \infty$, irrespective of the evaluation policy. The agent's payoff is maximized if the principal provides complete evaluation $m_t \equiv y_{t-1}$ and the agent quits optimally.*

4.2 Feasible Payoffs

I proceed with characterizing the set of feasible payoffs. Recall from standard game-theoretic terminology that a pair of strategies $m \in \mathcal{M}$, $\tau \in \mathcal{T}$ delivers payoffs (u^A, u^P) if given the strategies the payoff of the agent equals to u^A and a payoff of the principal equals to u^P . In turn, payoffs (u^A, u^P) are *feasible* if they can be delivered by some players' strategies. The payoffs are (weakly Pareto) *efficient* if there are no strategies that deliver strictly greater payoffs to both players. Denote the set of all feasible payoffs by W .

By standard arguments, the set of feasible payoffs is compact and convex. Its boundary ∂W can be characterized by the supporting hyperplane theorem. In particular, for any

⁶[Bergemann and Välimäki \(2008\)](#) provide a recent overview of the bandit experimentation literature.

payoffs $(u^A, u^P) \in \partial W$, there are payoff coefficients $\lambda^A, \lambda^P \in \mathbb{R}$ with at least one coefficient different from 0, such that:

$$(u^A, u^P) \in \arg \max_{(u^A, u^P) \in W} \lambda^A u^A + \lambda^P u^P. \quad (12)$$

Conversely, for any payoff coefficients $\lambda^A, \lambda^P \in \mathbb{R}$, a solution to the problem (12) belongs to the boundary ∂W . Positive agent's payoff coefficients, $\lambda^A > 0$, correspond to the east boundary of W . Positive principal's coefficients $\lambda^P > 0$ correspond to the north boundary of W . Positive coefficients of both players $\lambda^A, \lambda^P \geq 0$ span all efficient payoffs.

In other words, boundary payoffs can be delivered by some strategies that maximize different linear combination of players' payoffs. This maximization is equivalent to an optimal stopping problem with a modified safe option and performance value. In fact, a straightforward application of Lemma 1 shows the problem (12) is equivalent to maximizing or minimizing payoffs of a fictitious agent with a virtual safe option V' :

$$U^F(m, \tau, V') = \mathbb{E}_{m, \tau} \left[(1 - \delta) \sum_{t=0}^{\tau-1} \delta^t \alpha y_t + \delta^\tau V' \right]. \quad (13)$$

Lemma 2. (Feasible Payoffs) *The set of feasible payoffs W is compact and convex. Moreover:*

1. Its west boundary is delivered by $(m', \tau') \in \arg \min_{m, \tau} U^F(m, \tau, V')$ for $V' \geq 0$,
2. Its east boundary is delivered by $(m', \tau') \in \arg \max_{m, \tau} U^F(m, \tau, V')$ for $V' \geq 0$, and
3. Its efficient payoffs are delivered by $(m', \tau') \in \arg \max_{m, \tau} U^F(m, \tau, V')$ for $V' \in [0, V]$.

Figure 2 depicts a set of feasible payoffs of the binary example in Section 3.⁷ Points A and D , as well as any payoffs on the segment AD , can be delivered by a quitting time that does not depend on evaluations. Point A corresponds to payoffs $(\alpha \theta_0, (1 - \alpha) \theta_0)$ and can be delivered by the agent never quitting, $\tau \equiv \infty$. This is the best outcome for the principal. Point D corresponds to payoffs $(V, 0)$ and is delivered by agent quitting at time zero, $\tau \equiv 0$. This is the worst outcome for the principal. The payoffs on the segment AD can be delivered by a randomization between these two quitting times.

To obtain payoffs outside the segment AD , the quitting time must depend on productivity through an informative evaluation policy. Roughly, to deliver payoffs to the west of AD , the strategies should use performance information in a Pareto-destructive way. In particular, the west boundary is delivered by strategies that *minimize* the payoff of an agent with a

⁷The shape of the feasible payoff set depends on performance technology. In general, its boundary does not have to be piecewise linear.




Figure 2: A set of feasible payoffs W . Drawn for the binary example of Section 3.

virtual safe option. In contrast, to deliver payoffs to the east of AD , the strategies should use the information in a Pareto-improving way. In particular, the west boundary can be delivered by strategies that *maximize* the payoff of an agent with a virtual safe option. An arc CD corresponds to virtual options $V' \geq V$. An arc AC corresponds to virtual options $0 \leq V' \leq V$. Point C corresponds to a virtual safe option $V' = V$. This is the best outcome for the agent. Point B maximizes the principal's payoff among all payoffs that deliver at least a safe option payoff to the agent.

4.3 Equilibrium Payoffs

The notion of feasibility ignores players' incentives. In particular, all feasible payoffs can be delivered by a complete evaluation policy because the quitting time can ignore any additional information. However, under complete evaluation, the agent would act in his own interests. He would choose a quitting time to maximize his payoff, which would correspond to point C in Figure 2.

To incorporate the players' incentives I refer to mechanism-design terminology and say an evaluation policy $m \in \mathcal{M}$ *implements* payoffs (u^A, u^P) if the payoffs are delivered by the policy and some agent's best response to it. In turn, payoffs (u^A, u^P) are *implementable* if they can be implemented by some evaluation policy.

In these terms, a complete evaluation policy implements the payoffs of point C . However, it is not the only evaluation policy that implements them. Consider a policy that sends only two messages: "stay" and "quit." Let the policy mimic the agent's best response under a

complete-evaluation policy; that is, to send a “quit” message only after those performance histories at which the agent himself would quit. If the agent follows the recommendations, then the joint distribution of quitting time and performance will be the same. By Lemma 1, his payoff then equals the payoffs of point C . Because it is his maximal feasible payoff, he cannot do better than follow the recommendations, and so this recommendation policy would implement payoffs C .

Definition 2. An evaluation policy is a *recommendation policy* if it places a positive probability on at most two messages: “stay” and “quit.” A recommendation policy is *incentive compatible* if following the recommendations is an agent’s best response.

In fact, the agent cannot do better than follow the recommendations of an arbitrary recommendation policy that mimics his best response. This follows from the standard argument of Myerson (1986). Consider an arbitrary evaluation policy m and a recommendation policy m' that mimics an agent’s best response to m . Because the agent always knows his actions, the policy m' provides weakly less information than m . Hence, his payoff cannot be higher than that under m . Following the recommendations delivers the agent the same payoff as under m and hence is a best response to m' .

In other words, the recommendation policies provide minimal information for the agent to make his quitting decision. The principal does not need to provide any information besides that. Note that Lemma 1 is crucial for this observation because it establishes that the net payoff effect on evaluation policy comes only through its effect on a quitting time and not on a wage.

Lemma 3. (Recommendation principle) *All implementable payoffs can be implemented by incentive compatible recommendation policies.*

It is instructive to compare the recommendation principle with the informativeness principle of ?. The recommendation principle states that the principal may provide minimal information to the agent. Under a recommendation policy, the agent observes only “stay” or “quit” messages without any additional performance details. In contrast, the informativeness principle states that the agent’s wage should depend on the finest details of agent’s performance even when he is risk averse. This highlights the difference in focuses of the two papers. Specifically, I study a design of information for a fixed payoff structure (4), and ? studies a design of a reward contract for a fixed information structure.

I proceed with a characterization of implementable efficient payoffs. The agent can secure the payoff V by quitting at time 0. Hence, payoffs that deliver to the agent payoffs less than V cannot be implemented. In Figure 2, this means no efficient payoffs to the west of point B

can be implemented. At the same time, point C is implementable by a complete-evaluation policy. It turns out that all efficient payoffs on the segment BC are implementable as well.

Lemma 4. (Implementable Payoffs) *No payoffs with $u^A < V$ are implementable. All efficient payoffs with $u^A \geq V$ are implementable.*

This lemma builds on the efficient payoff characterization of Lemma 2 and the recommendation principle of Lemma 3. By the payoff characterization, the efficient payoffs can be delivered by strategies (m, τ) that maximize a payoff of an agent with a virtual safe option $V' \in [0, V]$. It turns out that the agent, if faced with a recommendation policy that mimics τ , is willing to follow recommendations. Recommendations to quit are incentive compatible because even the agent with a lower safe option is willing to follow them. Recommendations to stay are incentive compatible because the agent's continuation payoff weakly increases with tenure.

In equilibrium, the principal chooses an evaluation policy to maximize her payoffs. Equivalently, the principal maximizes her payoff among all implementable payoffs. It is clear from Lemma 4 that she optimally chooses an efficient payoff that delivers V to the agent (point B in Figure 2). The following theorem states the result formally.

Theorem 1. (Equilibrium Payoffs) *Equilibrium payoffs exist and are unique and strictly efficient:*

$$\begin{aligned} u^{A*} &= V, \\ u^{P*} &= \max_{(u^P, V) \in W} u^P. \end{aligned}$$

Proof. Because the set of feasible payoffs W is compact, such payoffs (u^{A*}, u^{P*}) exist. By Lemma 4, u^{P*} is an upper bound on the principal's payoffs and can be implemented, hence (u^{A*}, u^{P*}) are equilibrium payoffs. Because the efficient frontier is strictly decreasing in u^A , there are no other equilibrium payoffs and (u^{A*}, u^{P*}) are strictly efficient. \square

Theorem 1 establishes that despite the conflict of interests within the firm the equilibrium outcome is Pareto efficient. However, control of performance information allows the principal to extract all rents from the agent—his payoff is equal to his safe option V . Note that these results hold even though the principal can control only information and not a monetary incentive scheme.

5 Wage Profile

Theorem 1 establishes uniqueness of equilibrium payoffs. However, several equilibrium evaluation policies could possibly deliver these payoffs but result in different wage profiles. In what follows, I concentrate on a particular equilibrium in which the principal uses an incentive-compatible recommendation policy. By Lemma 3, this equilibrium always exists and there are at least two reasons to concentrate on it. First, the recommendation policies are minimally informative; any other policy can be Blackwell garbled into a recommendation policy without affecting the players' payoffs. Providing minimal information may be desirable to avoid its misuse by the agent in ways not conceivable by the principal. Second, the recommendation policies minimize wage volatility; any additional information results in a mean-preserving wage split. Minimizing volatility may be desirable if the agent is (marginally) risk averse.

Definition 3. An equilibrium (m^*, τ^*) is *minimally informative* if m^* is a recommendation policy and τ^* follows its recommendations.

In a minimally informative equilibrium, the wage profile is deterministic. Indeed, in any period, there is a unique history of past evaluations that results in the agent staying at the firm; namely, a sequence of recommendations to stay. Consequently, the equilibrium agent's career takes a particularly simple form. There is a deterministic wage profile W and commonly known performance requirements to stay at the firm. If the agent's performance satisfies these requirements, he stays at the firm; otherwise, he quits and secures a safe option.

The equilibrium wage profile depends on the equilibrium recommendation policy, which, in turn, depends on technology details. However, by Theorem 1, the equilibrium policy is efficient. It allows me to establish wage profile properties that hold for a broad class of technologies.

5.1 Downward Wage Rigidity

A naive intuition suggests that an average productivity and, hence, the wage should increase with tenure because equilibrium is efficient. Indeed, efficiency is commonly associated with ongoing positive selection that eliminates bad performers and retains good performers. Such positive selection can be implemented through a sequence of history-dependent cutoffs so that the agent is recommended to quit whenever his productivity expectation drops below the corresponding cutoff. Under such a policy, the expected productivity would increase after every history and, hence, the average productivity would increase with tenure.

Definition 4. A recommendation policy is a *cutoff policy* (in expectations) if there exists a cutoff function $q_t : Y^{t-1} \rightarrow \mathbb{R}$ such that

$$m_t = \begin{cases} \text{"stay,"} & \text{if } \mathbb{E}[\theta \mid y^{t-1}] > q_{t-1}(y^{t-2}), \\ \text{"quit,"} & \text{if } \mathbb{E}[\theta \mid y^{t-1}] < q_{t-1}(y^{t-2}). \end{cases}$$

The naive intuition overlooks the fact that, in general, efficient selection should account for the whole profile of productivity beliefs, not just productivity expectation. Roughly, a greater productivity variance increases the chances of being highly productive and, hence, increases the option value of quitting and experimentation. An agent with a lower expected productivity but higher chances of being very productive can be worth keeping, whereas an agent with higher but certain productivity can be worth terminating. As a result, an efficient policy may not be cutoff, and the resulting equilibrium wage may decrease.

Nevertheless, I show that efficient policies are cutoff in many cases. By Lemmas 2 and 3, any efficient policy maximizes a payoff of an agent with some virtual safe option $V' \in [0, V]$. Such policy is Markov in the productivity beliefs, so the space of beliefs can be split in two sets: the set at which the agent stays and the set at which the agent quits. The exact characterization depends on parameters of the problem and can be intractable. However, it can be obtained in the following cases.

If there are only two productivity types, $|\Theta| = 2$, then there is a threshold expectation q^c so that under any efficient policy, the agent stays if the productivity expectation is above the threshold and quits otherwise.⁸ That is, an optimal policy is cutoff with the cutoff function being constant.

If there are more than two types, $|\Theta| > 2$, then belief is a multidimensional object, and the efficient policy can be characterized only under additional assumptions.

Definition 5. A technology is *regular* if it admits densities and $\forall \theta, \theta' \in \Theta, y, y' \in \text{supp} f_\theta \cap \text{supp} f_{\theta'}, \theta' > \theta, y' > y$

$$f_{\theta'}(y') f_\theta(y) \geq f_{\theta'}(y) f_\theta(y') \quad (\text{MLRP}).$$

Under the regular technology, the conditional performance distributions satisfy the monotone likelihood property. Most technologies used in the literature satisfy this condition. The regularity assumption adds structure necessary to analyze the multiple-type case. Banks and Sundaram (1992) use the regularity condition to establish that an optimal strategy is cutoff.

⁸This standard result is also proven in the Appendix.

In both cases, in a minimally informative equilibrium, the evaluation policy is cutoff and wage profile exhibits downward rigidity. The following theorem summarizes the results.

Theorem 2. (Downward Wage Rigidity) *If there are only two types $|\Theta| = 2$ or the technology is regular, then in any minimally informative equilibrium:*

1. An evaluation policy m is a cutoff policy, and
2. The wage profile W is deterministic and weakly increasing.

The theorem establishes sufficient conditions under which the equilibrium wage exhibits downward rigidity. The conditions are plausible in that they are satisfied in most existing models of career concerns and experimentation. Interestingly, the proof shows even a stronger statement: not only the expected productivity increases, but also the whole profile of productivity beliefs shifts upwards in an MLRP sense.

5.2 Wage Profile Shape

Theorem 2 establishes that under general conditions the minimally informative wage profile is deterministic and weakly increasing. Nevertheless, the exact shape of the wage profile depends on technology details. To understand the role of technology on equilibrium, it is important to be able to calculate the wage profile.

The following algorithm builds on the results of Section 4 and calculates wage profile for a given technology, discount factor, and safe option. In the first step, it solves the maximization problems from Lemma 2 for all virtual safe options $V' \in [0, V]$. These problems can be without loss reduced to standard optimal stopping problems by setting evaluation policy to complete evaluation. Denote the corresponding optimal stopping times $\hat{\tau}(V')$. In the second step, the algorithm calculates the payoffs of the agent with a safe option V if he follows $\hat{\tau}(V')$. Denote the payoffs by $\hat{u}^A(V')$. In the last step, the algorithm finds V' such that $\hat{u}^A(V') = V$ so that the agent is left with no rents. Such V' exists because $\hat{u}^A(V')$ is continuous, $\hat{u}^A(V) \geq V$, and $\hat{u}^A(0) = w_0 < V$. Denote the corresponding safe option by V'^* . The evaluation policy that recommends $\hat{\tau}(V'^*)$, and the agent strategy that follows the recommendations, constitute an equilibrium by Theorem 1 and Lemma 3. The resulting wage profile and quitting rate can then be calculated from their definitions either in closed form or by Monte Carlo simulations.

In the rest of the section, I use this algorithm to calculate and contrast wage profiles and quitting rates under coarse and detailed performance technologies.

Coarse Performance In many industries, performance measures are coarse. A lawyer performance is captured by the number of successful trials, a consultant performance is measured by outcomes of his past projects, and a drug laboratory performance is captured by the number of new drugs developed. In all these cases, a performance outcome within any period is limited to few options that cannot fully reveal productivity. As a result, productivity learning is limited, and beliefs take a limited set of values. It turns out that this can translate into discontinuity of the equilibrium wage profile and a promotion-like career consisting of infrequent performance revisions followed by either quitting or staying at a higher wage.

I illustrate with the following example. The performance is binary, low or high, $Y = \{y^L, y^H\}$. Without loss of generality, $y^L = 0$, $y^H = 1$. There are two equally likely types, $\Theta = \{\theta^L, \theta^H\}$. The type corresponds to the probability of generating high performance,

$$\begin{array}{c|cc} f_{\theta}(y) & y^L & y^H \\ \hline \theta^L & 1 - \theta^L & \theta^L \\ \theta^H & 1 - \theta^H & \theta^H \end{array} .$$

with $0 < \theta^L < \theta^H < 1$ so that no performance realization is conclusive. Observing high performance y^H increases productivity expectation; observing low performance y^L decreases it.

Because the type is binary, as discussed in Section 5.1, an equilibrium evaluation policy is cutoff with a constant cutoff q . The cutoff is chosen so that the agent obtains a payoff V by following the recommendations. The cutoff and the corresponding wage profile and quitting rate can be calculated by Monte Carlo simulations and presented in Figure 3.

Note how the agent’s equilibrium career can be read off these plots. The agent starts working at period 0. In the first several periods, no evaluation is conducted—performance is not sufficiently informative to bring the expectations below q . During this time, the agent’s expected productivity and, hence, the wage, stay the same. In the beginning of fourth period, the first evaluation of past performance is conducted. If the performance was sufficiently bad, in this case corresponding to few high realizations, the wage drops below q , and the agent quits. Otherwise, the wage experiences a jump, and the agent stays at the firm for at least the next few periods. Such infrequent revisions are repeated throughout the agent’s career.

Detailed Performance Now, consider the case with detailed performance. For example, a manager can reveal his incompetency in day-to-day interactions with his clients and em-




Figure 3: Wage profile (left) and quitting rate (right) in a minimally informative equilibrium. Coarse performance. $\Theta = \{\theta^L, \theta^H\}$, $Y = \{0, 1\}$, $\Pr(y = 1 | \theta) = \theta$, $\alpha = 1/2$, $\theta^L = 0.2$, $\theta^H = 0.4$, $V = 0.6$. Computed numerically by Monte Carlo simulation.

ployees. In this case, it is possible to swing productivity expectations below the equilibrium cutoff in every period. As a result, the quitting may happen in any period, and the wage strictly increases with tenure.

I illustrate this case with the following example. Performance outcomes are rich, $Y = \mathbb{R}$. There are two possible types, $\Theta = \{\theta^L, \theta^H\}$, and each type is equally likely. Performance is distributed according to a Gaussian distribution with fixed variance and the mean being equal to the productivity:

$$y_t = \theta + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2).$$

In this way, lower performance decreases expected productivity. Importantly, the performance can be very conclusive—for any prior expectation below θ^H , a probability of interim expectation being arbitrarily close to θ^L is positive. As a result, there is a positive selection in every period, and the wage is strictly increasing with tenure. The equilibrium cutoff, wage profile, and the quitting rate can be computed numerically by the algorithm and are presented in Figure 4.

The equilibrium career with detailed performance differs from the career with coarse performance. The evaluations are conducted in every period, which results in a strictly increasing wage profile and strictly positive quitting rate in every period. However, in both cases the wage profile exhibits a (roughly) S shape that reflects the learning pattern and the rate of selection. At the beginning of the career, the quitting rate is low because there is little time to accumulate sufficiently negative performance information. Late in the career, the quitting rate is also low because, if retained, the agent was proven to have performed

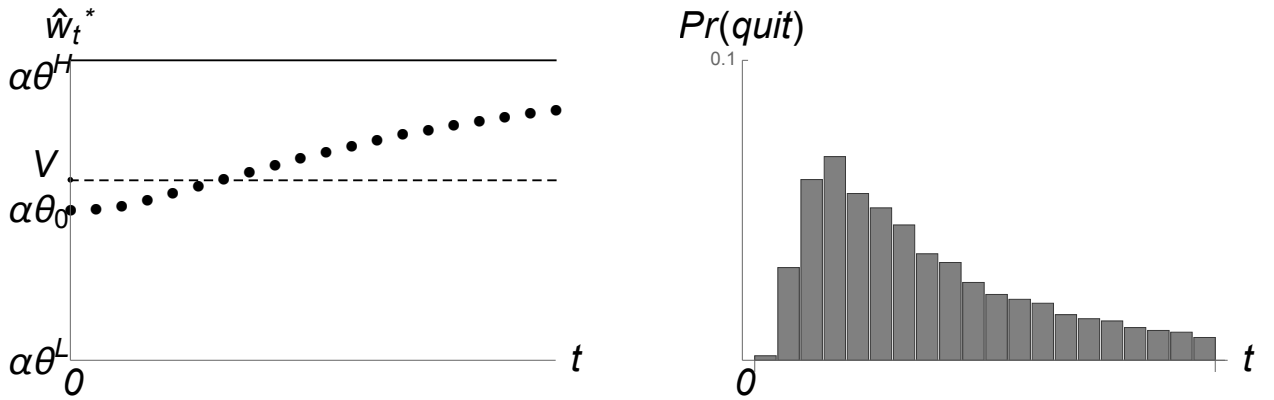


Figure 4: Wage profile (left) and quitting rate (right) in a minimally informative equilibrium. Rich performance. $\Theta = \{\theta^L, \theta^H\}$, $y_t = \theta + \varepsilon_t$, $\varepsilon_t \sim N(0, \sigma^2)$, $\sigma = 0.6$, $\alpha = 1/2$, $\theta^L = 0$, $\theta^H = 1$, $V = 0.6$. Computed numerically by Monte Carlo simulation.

well and is likely to be a high type. Most selection happens in the middle of the career, when information sufficient to swing expectations below a cutoff has been accumulated but much productivity uncertainty remains.

6 Discussion

6.1 Wages and Career

The analysis draws attention to a role evaluations may play in explaining internal labor market dynamics. I show organizations may systematically structure their evaluations in a way that results in wage patterns observed in practice. First, in a minimally informative equilibrium, the wage is a deterministic function of tenure. That is, all workers from the same cohort receive the same salary. Second, I show the equilibrium evaluation results in downward wage rigidity. Under complete evaluation, the wage may decrease, reflecting lower expectation following bad performance. However, under equilibrium evaluation, the wage never decreases—bad evaluations are infrequent and result in immediate resignation of the employee. As bad performers quit over time, the productivity expectations and the wage of retained employees increase. Notably, the wage grows even in the absence of human capital accumulation and elaborate incentive schemes. These qualitative properties are robust in a broad class of performance technologies and not driven by specific assumptions about the production process.

Interestingly, the analysis also provides an alternative view on career ladders. Charac-

terizing features of a career ladder are infrequent evaluations that result in an employee quitting the job or getting a promotion associated with a permanent wage increase. I show in Section 5.2 that this up-or-out structure can be driven by equilibrium evaluation when performance technology is coarse. The evaluation explanation is based solely on positive employee selection and does not require any increase of responsibilities or authority to come with promotion. This explanation may be particularly relevant for professional firms in which employee talents have high impact on performance and are learned gradually over time.

6.2 Evaluation Practices

My analysis provides a novel perspective on evaluation practices in organizations. First, it is commonly observed that evaluations feature rating compression and lenience bias—most employees are bunched at the highest grades with lower grades reserved for exceptionally bad performance. This feature is often attributed to behavioral biases of evaluators and criticized for lowering evaluation informativeness. However, my analysis shows these practices can in fact benefit a firm. In a minimally informative equilibrium, while staying at the firm, all employees receive the same “good” grade. A “bad grade” is received only for exceptionally bad performance, which lowers the wage dramatically and leads to immediate resignation. Second, there is no agreement on whether evaluations should be retrospective, that is, account for past performance. My analysis suggests that such an accounting is generally necessary. The equilibrium evaluation policy is Markov in belief; that is, it based on the whole performance history. In general, this policy cannot be implemented through a sequence of grades based solely on current performance.

The equilibrium evaluation policy can also have a psychological interpretation for maintaining workforce morale. The organizational literature has long recognized that optimal evaluation provision must balance between two opposing forces—the employee’s need for evaluation and the damage to his self-esteem. On one hand, evaluation is a valuable input for employees’ future decisions. On the other hand, any evaluation may turn out negative and discourage an employee. My model can be viewed as a formalization of these effects with the agent’s morale captured by his productivity belief.⁹ The equilibrium policy can then be interpreted as a policy that optimally trades-off these two competing effects.

⁹Fang and Moscarini (2005) study a similar trade-off in a static setting.

6.3 Commitment

My analysis relies on the ability of the firm to commit to its evaluation policy. In fact, several layers of commitment matter. First, I assume evaluations are objective, so that the information is verifiable. If the evaluation would be “cheap talk,” then the equilibrium recommendation policy would not be credible—the principal would always recommend the agent to stay. Second, I assume the principal commits to the evaluation rules at the beginning of the relationship and cannot change them down the road. This assumption matters because the equilibrium evaluation policy is not time consistent. Indeed, under the equilibrium policy, the only binding constraint is the participation constraint at time 0. As long as the agent stays at a firm, his continuation value increases, and his incentive constraints become slack. Hence, even if the players are symmetrically informed, the principal would like to change the evaluation policy after a recommendation to stay to extract more continuation surplus. The incentives to change the policy are even stronger if the principal can observe additional information besides evaluations. For these reasons, the equilibrium policy is not implementable without commitment. Because any equilibrium policy without commitment can be replicated with commitment, the commitment power is strictly valuable for the firm.

One way for a firm to sustain evaluation commitment is to have an independent human resource department. The evaluation policy then corresponds to rules it uses in evaluating their employees, such as timing and criteria of conducting evaluations. As long as the rules are commonly known and the department is independent, the department has no incentives to not follow them. This would make the evaluations credible and allow the firm to implement the equilibrium policy.

7 Conclusion

Evaluation of individual performance is an important part of organizational life. Although much evaluation is informal, most firms have formal evaluation policies designed to collect and distribute performance information to employees. In this work, I showed an equilibrium design of these policies can explain many features of internal labor market observed in practice, including downward wage rigidity. Information control allows the organization to extract all rents from its employees and achieve a Pareto-efficient outcome despite conflict of interests.

For clarity and conciseness, I abstracted away from many realistic features of internal labor markets. There may be richer allocation decisions—the employee may decide how much effort to exert, or the firm may choose what tasks to assign the employee. Performance

information may not be under full control of the firm—some performance outcomes may be publicly observable. Richer monetary contracts may be available. Incorporating these features would make the model more applicable in practice and constitute a plausible venue for future work. The efficiency approach developed in this paper may facilitate analysis of these richer settings as well. In the meantime, the proposed evaluation theory could be viewed as complimenting existing theories of wage dynamics and internal labor markets.

References

- BANKS, J. S. AND R. K. SUNDARAM (1992): “Denumerable-Armed Bandits,” *Econometrica*, 1071–1096.
- BERGEMANN, D. AND J. VÄLIMÄKI (2008): “Bandit Problems,” in *New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Basingstoke, UK: Palgrave Macmillan.
- EDERER, F. (2010): “Feedback and Motivation in Dynamic Tournaments,” *Journal of Economics & Management Strategy*, 19, 733–769.
- ELY, J. C. (2017): “Beeps,” *American Economic Review*, 107, 31–53.
- ELY, J. C. AND M. SZYDLOWSKI (2017): “Moving the Goalposts,” *Working paper*.
- FANG, H. AND G. MOSCARINI (2005): “Morale Hazard,” *Journal of Monetary Economics*, 52, 749–777.
- GITTINS, J. C. AND D. M. JONES (1974): “A Dynamic Allocation Index for the Sequential Design of Experiments,” in *Progress in Statistics*, ed. by I. Vincze, J. Gani, and K. Sarkadi, Amsterdam: North-Holland Pub. Co., 241–266.
- GOLTSMAN, M. AND A. MUKHERJEE (2011): “Interim Performance Feedback in Multistage Tournaments: The Optimality of Partial Disclosure,” *Journal of Labor Economics*, 29, 229–265.

- GUO, Y. (2016): “Dynamic Delegation of Experimentation,” *American Economic Review*, 106, 1969–2008.
- HALAC, M., N. KARTIK, AND Q. LIU (2016): “Optimal Contracts for Experimentation,” *Review of Economic Studies*, 83, 1040–1091.
- HANSEN, S. E. (2013): “Performance Feedback with Career Concerns,” *Journal of Law, Economics, and Organization*, 29, 1279–1316.
- HOLMSTRÖM, B. (1999): “Managerial Incentive Problems: A Dynamic Perspective,” *The Review of Economic Studies*, 66, 169–182.
- HÖRNER, J. AND Y. GUO (2015): “Dynamic Mechanisms without Money,” *Working Paper*.
- HÖRNER, J. AND N. LAMBERT (2016): “Motivational Ratings,” *Working Paper*.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KREMER, I., Y. MANSOUR, AND M. PERRY (2014): “Implementing the ‘Wisdom of the Crowd’,” *Journal of Political Economy*, 122, 988–1012.
- MURPHY, K. R. AND J. CLEVELAND (1995): *Understanding Performance Appraisal: Social, Organizational, and Goal-Based Perspectives*, Sage.
- MYERSON, R. B. (1986): “Multistage Games with Communication,” *Econometrica*, 54, 323–358.
- ORLOV, D. (2016): “Optimal Design of Internal Disclosure,” *Working Paper*.
- ORLOV, D., A. SKRZYPACZ, AND P. ZRYUMOV (2017): “Persuading the Principal to Wait,” *Working Paper*.
- RAYO, L. AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118, 949–987.

RENAULT, J., E. SOLAN, AND N. VIEILLE (2017): “Optimal Dynamic Information Provision,” *Games and Economic Behavior*, 104, 329–349.

RODINA, D. (2016): “Information Design and Career Concerns,” *Working paper*.

THOMPSON, F. B. (1952): “Equivalence of Games in Extensive Form,” Technical report rm-759, RAND Corporation, Washington, D.C.

8 Appendix

Proof of Lemma 1 Take any evaluation policy m and consider a stochastic process $X = \{X_s\}_{s \geq 1}$ with

$$X_s = (1 - \delta) \sum_{t=0}^{s-1} \delta^t \left(y_t - \mathbb{E}[\theta \mid m^t] \right).$$

The process X is a bounded martingale with respect to a filtration generated by m because

$$\begin{aligned} \mathbb{E}[X_{s+1} \mid m^s] &= X_s + (1 - \delta) \delta^s \mathbb{E}[y_s - \mathbb{E}[\theta \mid m^s] \mid m^s] \\ &= X_s + (1 - \delta) \delta^s (\mathbb{E}[y_s \mid m^s] - \mathbb{E}[\theta \mid m^s]) \\ &= X_s + (1 - \delta) \delta^s (\mathbb{E}[\theta \mid m^s] - \mathbb{E}[\theta \mid m^s]) \\ &= X_s. \end{aligned}$$

The second line follows from the law of iterated expectations and the third line follows from the productivity definition (1). It follows from the optional stopping theorem that for any stopping time τ measurable with respect to m :

$$\mathbb{E}[X_\tau] = \mathbb{E}[X_0] = 0.$$

Equivalently, by wage definition (4), for any evaluation policy m and a stopping time τ ,

$$(1 - \delta) \mathbb{E}_{m,\tau} \sum_{t=0}^{\tau} \delta^t w_t = (1 - \delta) \sum_{t=0}^{\tau} \delta^t \alpha y_t.$$

The result follows. □

Proof of Lemma 2. As shown in the main body of the paper all boundary payoffs can be obtained by strategies that maximize a linear combination of players' payoffs. By Lemma 1, this combination can be written as

$$\lambda^A U^A(m, \tau) + \lambda^P U^P(m, \tau) = \mathbb{E}_{m,\tau} \left[(1 - \delta) \sum_{t=0}^{\tau} \delta^t \left(\lambda^A \alpha + \lambda^P (1 - \alpha) \right) y_t + \delta^\tau \lambda^A V \right].$$

If $\lambda^A \alpha + \lambda^P (1 - \alpha) > 0$, then maximizing the payoff combination is equivalent to maximizing a payoff of an agent with a virtual safe option $V' = \frac{\lambda^A \alpha}{\lambda^A \alpha + \lambda^P (1 - \alpha)} V$. If $\lambda^A \leq 0$, then clearly $\tau \equiv \infty$ so all such coefficients correspond to payoffs $(\alpha \theta_0, (1 - \alpha) \theta_0)$. If $\lambda^A > 0$, then the payoffs belong to the east boundary and $V' > 0$. In this case, efficient payoffs correspond to $\lambda^P > 0$ and $V' \in (0, V]$.

If $\lambda^A \alpha + \lambda^P (1 - \alpha) < 0$, then maximizing the payoff combination is equivalent to minimizing a payoff of an agent with a virtual safe option $V' = \frac{\lambda^A \alpha}{\lambda^A \alpha + \lambda^P (1 - \alpha)} V$. If $\lambda^A \geq 0$, then the optimal quitting time is $\tau \equiv 0$, corresponding to payoffs $(V, 0)$. If $\lambda^A < 0$, then the payoffs belong to the west boundary and $V' > 0$. The result follows. \square

Proof of Lemma 4. Consider any efficient payoffs (u^A, u^P) . If $u^A < V$, then the payoffs cannot be implemented because the agent can secure a payoff V by quitting at time 0. If $u^A \geq V$, then by Lemma 2 the payoffs can be delivered through strategies (m, τ) that maximize a payoff of an agent with a safe option $V' \in [0, V]$. Consider the recommendation policy m' that mimics τ . I show that m' is incentive compatible.

Recommendations to quit are incentive compatible. If the agent is recommended to quit at time t , his wage at the firm stays constant at some level w_t^Q hereafter. Because the strategy maximizes a virtual payoff, $w_t^Q \leq V' \leq V$. Hence, the agent with a safe option V agrees to quit as well.

Recommendations to stay are incentive compatible. Recommendation at time 0 is incentive compatible because it gives the agent a payoff of at least u^A and $u^A \geq V$. Incentive compatibility at time $t+1$ follows from the incentive compatibility at time t . Indeed, denote by \hat{U}_t^A the agent's optimal continuation payoff conditional on being recommended to stay at time t . Because the recommendation to stay is incentive compatible at time t and the recommendations to quit are incentive compatible everywhere, \hat{U}_t^A is a convex combination of w_t , \hat{U}_{t+1}^A , and V ,

$$\hat{U}_t^A = (1 - \delta) w_t + \delta \Pr(m_{t+1} = \text{"stay"} | m^t \equiv \text{"stay"}) \hat{U}_{t+1}^A + \delta \Pr(m_{t+1} = \text{"quit"} | m^t \equiv \text{"stay"}) V.$$

The agent obtains a continuation payoff V if he quits immediately so $\hat{U}_t^A \geq V$. The agent obtains a continuation payoff w_t if he never quits so $\hat{U}_t^A \geq w_t$. Thus, for the equality to hold, it must be that $\hat{U}_{t+1}^A \geq \hat{U}_t^A \geq V$. By forward induction, all recommendations to stay are incentive compatible. This completes the proof. \square

Lemma 5. Fix $V' \geq 0$, $\Theta = \{\theta^L, \theta^H\}$, and an arbitrary technology. Consider the following maximization problem:

$$\max_{m \in \mathcal{M}, \tau \in \mathcal{T}} \mathbb{E}_{m, \tau} \left[(1 - \delta) \sum_{t=0}^{\tau-1} \delta^t \mathbb{E}[\theta | m^t] + \delta^\tau V' \right].$$

Then without loss of generality $m \equiv y$ and the following statements are true.

1. If $V' > \theta^H$, then $\tau \equiv 0$ is uniquely optimal.

2. If $V' < \theta^L$, then $\tau \equiv \infty$ is uniquely optimal.

3. If $V' \in [\theta^L, \theta^H]$, then any optimal strategy τ is equivalent (in terms of quitting distributions) to a cutoff quitting strategy with some fixed cutoff $q \in [\theta^L, \theta^H]$.

Proof. The problem is a standard bandit problem of an agent choosing between a risky arm with payoffs $\mathbb{E}[\theta | m^t]$ and a safe arm with a payoff V' . This is a Markov decision problem with an expected productivity $\hat{\theta}_t \triangleq \mathbb{E}[\theta | y^t]$ being a sufficient statistic for beliefs. An optimal quitting strategy is characterized by Gittins and Jones (1974) to be an index policy. The risky arm is assigned an index $\xi(\hat{\theta}_t)$. The agent quits the arm as soon as the index drops below V' . The decision maker can arbitrarily randomize at $\xi(\hat{\theta}_t) = V'$. The index is

$$\begin{aligned}
\xi(\hat{\theta}_t) &\triangleq \sup_{\tau} \frac{\mathbb{E} \left[\sum_{t=0}^{\tau} \delta^t \mathbb{E}[\theta | m^t] \mid \hat{\theta}_t \right]}{\mathbb{E} \left[\sum_{t=0}^{\tau} \delta^t \mid \hat{\theta}_t \right]} \\
&= \sup_{\tau} \frac{\frac{\hat{\theta}_t - \theta^L}{\theta^H - \theta^L} \mathbb{E} \left[\sum_{t=0}^{\tau} \delta^t y_t \mid \theta = \theta^H \right] + \frac{\theta^H - \hat{\theta}_t}{\theta^H - \theta^L} \mathbb{E} \left[\sum_{t=0}^{\tau} \delta^t y_t \mid \theta = \theta^L \right]}{\frac{\hat{\theta}_t - \theta^L}{\theta^H - \theta^L} \mathbb{E} \left[\sum_{t=0}^{\tau} \delta^t \mid \theta = \theta^H \right] + \frac{\theta^H - \hat{\theta}_t}{\theta^H - \theta^L} \mathbb{E} \left[\sum_{t=0}^{\tau} \delta^t \mid \theta = \theta^L \right]} \\
&= \sup_{\tau} \frac{(\hat{\theta}_t - \theta^L) \theta^H \mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^H \right] + (\theta^H - \hat{\theta}_t) \mu^L \mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^L \right]}{(\hat{\theta}_t - \theta^L) \mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^H \right] + (\theta^H - \hat{\theta}_t) \mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^L \right]} \\
&= \sup_{\tau} \theta^L + \frac{(\hat{\theta}_t - \theta^L) (\theta^H - \theta^L) \mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^H \right]}{(\hat{\theta}_t - \theta^L) \mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^H \right] + (\theta^H - \hat{\theta}_t) \mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^L \right]} \\
&= \theta^L + \frac{(\hat{\theta}_t - \theta^L) (\theta^H - \theta^L)}{(\hat{\theta}_t - \theta^L) + (\theta^H - \hat{\theta}_t) \kappa}
\end{aligned}$$

where $\kappa \triangleq \inf_{\tau} \frac{\mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^L \right]}{\mathbb{E} \left[1 - \delta^{\tau+1} \mid \theta = \theta^H \right]} > 0$ and the second and third lines follow from the optional stopping theorem and the law of iterated expectations. Note that κ depends only on the performance technology F_{θ} and not on $\theta^H, \theta^L, \theta_0$, or V' . Thus, $\xi(\hat{\theta}_t)$ strictly increases with $\hat{\theta}_t$ and the result follows. \square

Proof of Theorem 2. If the type is binary, $|\Theta| = 2$, then by Lemma 5 for any virtual safe option $V' \in [0, V]$, an optimal policy is cutoff. Hence, the wage profile is weakly increasing.

If the type is not binary, $|\Theta| > 2$, but the performance technology is regular then by Banks and Sundaram (1992), an agent-optimal policy is cutoff. It follows that the wage profile is weakly increasing. \square