

A Theory of Monitoring and Internal Labor Markets¹

Gautam Bose

School of Economics, University of New South Wales

Kevin Lang

Department of Economics, Boston University, and NBER/IZA

August 15, 2011

¹We are grateful to Sambuddha Ghosh, Eddie Lazear, Andy Newman, Mike Waldman and participants at seminars and workshops at Boston University, the NBER and the University of New South Wales for helpful comments and suggestions. The usual caveat applies.

Abstract

We analyze a firm's job-assignment and worker-monitoring decisions when workers face occasional crises. Firms prefer to assign good workers to a difficult task and to not employ bad workers. Firms observe failures but only observe successfully resolved crises if they monitor the worker. If monitoring costs are positive but sufficiently small, for a range of probabilities that the worker is good, the firm assigns the worker to a low task (less sensitive to crises) and monitors her. At probabilities below this range and not too much above it, she is assigned to the low task and not monitored. At high probabilities of being good, she is assigned to the difficult task. We analyze the implications for internal labor markets of the case where a worker has the same *ex ante* probability of being good at all firms and learning is about ability at this particular firm.

We develop a model of monitoring and internal labor markets based on two key elements. First, monitoring is primarily used to catch errors and to evaluate workers and not, as assumed in the agency literature, primarily to deter shirking, stealing or other forms of moral hazard. Second, in most jobs much of the work is routine in the sense that it can be performed similarly by any worker who is reasonably well-trained for the job. Worker quality is primarily about the ability to handle the nonroutine elements of the job, which we call crises. We show that a model that combines these elements can explain why in many jobs wages are largely determined by initial human capital and seniority and that the model predicts or is consistent with many of the regularities in the literature on internal labor markets.

In the base model, workers can be either good or bad. All workers can perform routine tasks, but only good workers can solve crises. Workers may be assigned to either of two tasks of which one is more productive under routine circumstances but is also subject to more costly crises. After hiring a worker, firms choose one of three strategies: 1) place the worker in the higher productivity but more crisis-sensitive task and hope for the best, 2) place him in the less sensitive task and monitor his performance, or 3) place him in the less sensitive task and not monitor his performance. Given the structure of the base model, monitoring workers in the higher task is never optimal.

Initially we abstract from wages and the presence of an external labor market. We consider only the optimal allocation of workers to task and optimal monitoring of their performance. We show that if the cost of monitoring is finite but sufficiently small, then there are four assignment tiers. Workers whom the firm believes are very likely to be good are placed in the high task (group A). Those with a somewhat lower probability of being good are not monitored. After some time, if they have not been observed to mishandle a crisis, they are placed in the high-tier task (group B). Below this tier, workers are monitored (group C). Once a crisis occurs, worker quality is known, and the worker is assigned to the appropriate task. Finally, if the probability that the worker is good is sufficiently low, the firm again does not monitor him (group D). However, unless its prior that the worker is good is 0, after some time if the worker has not faced a crisis and failed, the firm's updated estimate of the probability that the worker is good becomes sufficient for the firm to monitor him.

We do not model wage determination but assume that the wage is an increasing function of the probability that the worker is good. We also assume that the worker and firm separate whenever the probability that the worker is good falls below its initial value, consistent with an interpretation

of worker quality as match-specific.

Under these assumptions, provided they do not separate, a worker who is initially placed in the low no-monitoring range receives continuous (automatic) wage increases until the probability that he is good rises sufficiently so that the firm monitors him. At this point, he reaches the “top of the scale” for his “job” and does not receive a wage increase until he is promoted.

While the base model fits some regularities reported in the literature on internal labor markets, it is inconsistent with some results regarding promotion patterns. Allowing for partial monitoring and for false positives (apparent crises that can be solved by all workers) addresses many of these inconsistencies and provides a basis for exploring the sources of variation in internal labor markets.

The next section briefly discusses some of the underpinnings of the model. It is followed by an intuitive presentation of the argument. The formal model is described and its implications for monitoring are analyzed in section three. In section four we develop the implications for internal labor markets and discuss their relation to the empirical regularities. Section five presents some extensions. The last section has concluding remarks.

1 Internal Labor Markets, Monitoring and Crises

In the last two decades the Baker, Gibbs and Holmstrom (1994a&b), hereafter BGH, analysis of the career patterns of managerial workers within a single service sector firm has justifiably received a great deal of attention. In particular, Gibbons and Waldman (1999a, 2006, see also 1999b) develop a simple and elegant model designed to address many of the features of the internal labor market described in BGH (1994b). To some extent, we, too, will seek to explain some of the regularities in the internal labor market studied by BGH, but we also emphasize that the features of that internal labor market are by no means universal.

1.1 Nonmanagerial Internal Labor Markets

There is a great deal of wage-setting at the individual level among the managerial employees in the firm studies by BGH. It is difficult to get precise information on how common it is for firms to set pay at the individual level. We provide some evidence from the UK Workplace Employment Relations Survey (WERS).

The WERS asks managers to list the determinants of pay in the largest *nonmanagerial* occupation in their establishment. Occupations are defined

in very broad categories (e.g. skilled workers, administrative and secretarial, technical). Because collective bargaining agreements may require firms to set wage differentials on the basis of objective measures, we limit our analysis to firms with no workers whose wages are set by collective bargaining. Among these nonunion establishments, roughly 20 percent report *no* variation in pay within their largest class of nonmanagerial workers except for hours, overtime and shift differentials, and another 3 percent report using only these factors and some form of objective pay for performance such as piece rates or commission.¹ Of the remaining establishments, a further 14 percent differentiate pay on the basis of skills/core competencies and/or job grade but not factors such as seniority or performance evaluations. Since the occupations considered are quite broad, such differences may reflect only differences between jobs (e.g. carpenter and electrician). We restrict the remainder of the analysis to establishments that clearly have some form of individualized pay that is not strictly piece-rate or similar.

Of those for whom there is clear differentiation in pay within similar jobs, 63 percent determine pay at least in part by age, experience, seniority and/or formal qualifications such as education but not on the basis of subjective performance measures. In contrast, only 5 percent use subjective performance evaluations but not the objective factors in the previous sentence and 32 percent use both. These numbers should be treated with considerable caution. Of managers who reported that performance appraisal or assessment was one of the factors that explain the differences in the level of pay of full-time workers in the largest occupation in their establishment, only about one-quarter also reported that merit pay, defined as pay related to a subjective assessment of individual performance by a supervisor or manager, was used *anywhere* in the establishment.

In sum, it is unclear just how prevalent individualized pay based on individual performance is, at least at the nonmanagerial level. It is, however, very clear that much of pay is determined by objective factors that *may* be correlates of productivity but are certainly not measures of productivity. Almost half of establishments that have differentiated pay do so on the basis of seniority within the firm and about three-quarters do so on the basis of either seniority or experience.

One of the tasks before us is therefore to explain why firms might rely on such factors, particularly seniority, rather than on either actual productivity or subjective performance measures.

¹All calculations use the establishment weights.

1.2 Monitoring

Payment on the basis of either measured output or subjective performance requires some form of monitoring. Surprisingly, the literature on monitoring in the labor market almost universally assumes that its sole purpose is to allow firms to punish workers who shirk, cheat or steal. Yet, this assumption both generates empirical predictions that are, at the least, problematic and is inconsistent with our everyday experience.

From a theoretical perspective, modeling monitoring as designed to enforce good behavior is problematic. The standard result from the literature on crime is that since detection is costly and deterrence depends on the probability of detection multiplied by the cost of punishment, fines should be as large as possible and monitoring should be minimal. Dickens, Katz, Lang and Summers (1989) refer to this as the monitoring puzzle. Akerlof and Katz show that the only solution to models of this sort is the one derived by Becker and Stigler (1974), which is to have workers “buy” their jobs and to have the purchase price returned to them when they retire. If workers’ ability to purchase their job is limited, firms may require them to engage in rent dissipating behavior (Murphy and Topel, 1990). Neither purchase of jobs nor obviously rent-dissipating requirements are a common feature of job contracts. If bonding is costless, more general earnings profiles are possible, but the logic of the argument requires that wages be less than value of marginal product early in seniority and more than VMP later (Lazear, 1979, 1981). The efficient contract sets hours so that VMP equals the worker’s marginal value of leisure, while workers would wish to choose to set the marginal value of leisure equal to their wage. This implies that junior workers will want to work less than required by the optimal contract while senior workers will want to work more. In fact, the desire to work less increases with seniority (Kahn and Lang, 1992, 1995).

Moreover, the view that monitoring serves primarily as a discipline device violates everyday experience. We do not check our research assistant’s work mainly to deter him from shirking, nor do we fire him if we catch an error. We want to avoid the cost of a mistake. The frequency with which we find errors will affect our assessment of the research assistant and thus our decision to rehire him. But in this sense monitoring is part of our process of evaluating the research assistant. The threat that we may fire or not rehire the assistant may have the additional effect of encouraging the assistant to work harder in order to reduce the frequency of errors, but our monitoring decision is mostly affected by our need to catch errors and to know whether the assistant is competent. As we become more convinced that the assistant

can handle the tasks to which he has been assigned, we are likely to reduce our monitoring. In this respect we resemble private sector managers who often monitor workers as part of an evaluation process and review work to determine that it has been done correctly.

In this paper, we focus on monitoring for the purpose of evaluation although in the extensions section, we address monitoring to catch errors. Formal modeling of monitoring solely for the purpose of catching and correcting mistakes provides us with few nontrivial insights, while combining the possibility of productive monitoring with evaluative monitoring changes the mathematics but not the essential message.

Our model is closest in spirit to Bjerk (2008) who assumes that firms learn less about worker productivity in low-level jobs where productivity is less responsive to ability. However, in that model, presumably because firms cannot capture rents from their knowledge of worker productivity and they cannot commit to a job assignment, variation in learning is an exogenous response to a job assignment decision made on the basis of where the worker's expected productivity is highest.

1.3 Crises

The way we model monitoring reflects our view that the issue is the ability to do the difficult parts of a job. In many, perhaps most, jobs much of the work is routine for any trained worker. We do not wish to imply that there are no specialized skills. Many of us do not know how to change the oil in a car, but for anyone trained as a mechanic, this is usually straightforward. Physicians frequently see and recognize the same cluster of symptoms, making certain diagnoses straightforward for anyone who has been properly trained. Most real estate transactions are sufficiently simple that in some U.S. states, no lawyer need be involved.

However, sometimes something nonroutine arises: the drain plug will not loosen; the symptoms do not quite fall into the usual cluster; previous sale of the property was mishandled. If the individual faced with the nonroutine task is skilled, she may be the only one who is ever aware of it: she finds an appropriate torque wrench and loosens the plug without stripping it; he diagnoses the condition accurately and prescribes the proper treatment; she contacts the lawyer who handled the previous sale and has the problem corrected. However, if the worker is unskilled at his job, the outcome may be very noticeable: he strips the plug; oil leaks out and the engine is seriously damaged; she produces a diagnosis of asthma when the problem is heart disease; the patient soon dies of a heart attack; the legal error from the

last sale goes unnoticed; the new owner later incurs considerable expense to establish rightful ownership. The famous con-artist, Frank Abagnale, claims to have worked for eleven months as chief resident pediatrician in a Georgia hospital until he was faced with an oxygen-deprived baby and was almost exposed (Abagnale, 1980).

2 Intuition

It is very costly to assign a worker to a job for which he does not have the requisite skills. Only when the firm has reasonable confidence about the worker's ability to respond to crises will the worker be assigned to high-level jobs where crises can be costly. But it is also valuable to make full use of skilled workers who are capable of responding to crises, rather than to assign them to low-level jobs where their skill is less valuable. Therefore the firm wants to learn whether or not its workers are skilled.

Failure to solve a crisis is usually very noticeable. So one option for the firm is simply to wait. As time passes, it becomes less and less likely that the worker has never been faced with a crisis, and if the firm has not observed a crisis, more and more likely that the worker addressed one or more crises successfully.

Alternatively, the firm can monitor the worker to determine if a crisis is occurring. In this case, it knows whether or not there has been a crisis rather than simply inferring the likelihood of the crisis from the length of time the worker has been at risk of facing a crisis. In either case, the firm will know if the worker has failed to solve a crisis. The advantage of monitoring is that if the worker solves a crisis, the firm learns this immediately, and because it now knows that the worker can handle crises, can move him instantly to a job where his skill is more valuable.

When should the firm monitor the worker, and when should it simply wait? If the firm's prior that the worker is good is very low, then monitoring has very little benefit; the probability that the firm will observe a crisis and discover that the worker is actually good is, by definition, very low. Therefore the expected benefit is low and monitoring is unlikely to be profitable.

Consider now the case of an unmonitored worker who is just about to be promoted assuming that she does not fail to solve a crisis during the very brief period before promotion. By definition, *ex ante* the firm expects to make (almost) the same profit whether it places her in the high or low job. Thus, if the firm monitors the worker, it gives up an expected zero

flow-benefit but pays the flow cost of monitoring until a crisis arrives and thus has a finite expected cost.

So the firm will not monitor workers who are either very unlikely to be good or who are sufficiently likely to be good that they, with high probability, will be promoted shortly anyway. If any workers are monitored, it will be those with an intermediate probability of being good.

We make this argument formally in the next section.

3 The Formal Model

3.1 Workers and jobs

An employer hires a worker whom he can put in a high task (H) or a low task (L). The worker's productivity in a given task depends on his type, which may be good (G) or bad (B). Both types of worker produce a flow output normalized to zero per unit time in the low task and $g > 0$ per unit time in the high task. θ_0 will be the firm's prior belief that the worker is good. For the moment it is irrelevant whether type is general or firm-specific, but it may be helpful to think of it as firm-specific. In the next section, when we discuss internal labor markets, we will treat the value of θ_0 when the worker begins working for the firm as the probability that the worker is good at a randomly chosen firm so that a worker with an initial θ_0 equal to 0.7, will be good at 70 percent of possible firms at which he may work and bad at the remaining 30 percent.

Crises occur in both H and L tasks with a Poisson arrival rate λ . This assumption ensures that task assignment is unaffected by its impact on learning about productivity. Assumptions of this nature are common in the literature on internal labor markets (e.g. Gibbons and Waldman, 1999a). It might be more natural to assume that crises are more common in the high tasks, but it greatly complicates the math, with, as far as we can tell, little additional insight.

Bad workers *fail* when a crisis occurs. Failure generates negative output of $-c_l$ in the L -task and $-c_h$ in the H -task, with $c_h > c_l$. If a worker is bad and a crisis occurs, then the failure is immediately observed and the worker's type is revealed. Good workers resolve crises when they occur, with no impact on productivity. Thus if the worker is good, the occurrence of a crisis can be known only if the worker is actively monitored, in which case the worker's type is revealed. We assume that $g - \lambda c_h < -\lambda c_l$ so that the expected flow of output net of costs associated with crises is more negative

when a bad worker is placed in the H -task than when he is placed in the L -task.

Time is continuous and the future is discounted at a rate r .

Under complete information, it is clear that good workers will be put in the H -task and bad workers will leave the firm since their productivity is negative. We assume that if a worker is revealed to be bad, she separates from the firm immediately. This assumption is natural if type is firm-specific. Whether the assumption makes sense if type is general will depend on how the labor market is structured.

3.2 Monitoring

The employer can use one of two strategies to assess workers: monitor (M) or not-monitor (N). Under strategy N he assigns the worker to a task and does not monitor him. Thus he only gets confirmation of the worker's type if the worker is bad and a crisis occurs, in which case the worker fails. If the worker is good, or until a crisis arrives, the employer observes nothing, and can update his beliefs about the worker as time passes.

Under strategy M, the employer monitors the worker until a crisis occurs, at which time he learns the worker's type. There is a flow cost of b per unit time of monitoring, and the cost must be borne until a crisis occurs. Note that under the monitoring strategy the employer's beliefs remain unchanged until a crisis occurs, at which time the employer knows the worker's type. Therefore, if it was optimal to monitor the worker, it will continue to be optimal until the firm observes a crisis.

3.3 Optimal promotion without monitoring

First consider a worker who is not monitored by the firm. If θ_0 is not too high, the employer will place the worker in the L -task. If the worker fails at some time, his type is revealed to be bad, and he will leave the firm. If he has not failed until time t , the employer updates his belief about the worker's type to $\theta(t, \theta_0)$. As time passes and $\theta(t, \theta_0)$ becomes sufficiently high, the employer may promote the worker to the H -task. Similarly, a worker who comes in with a sufficiently high prior at time 0 will be placed immediately in the H -task.

Given Poisson arrival, the density function for the arrival of the first crisis is $\lambda e^{-\lambda t}$, hence the probability that the first crisis arrives by time τ is $p(\tau) = 1 - e^{-\lambda\tau}$. Thus the probability that a bad worker does not fail by time τ is $1 - p(\tau) = e^{-\lambda\tau}$.

Theorem 3.1. *If the firm does not monitor the worker, then it promotes the worker to the H -task when its assessment of the probability that the worker is a good worker reaches*

$$\theta^* = \frac{\lambda(c_h - c_l) - g}{\lambda(c_h - c_l)} \quad (1)$$

provided that

$$\theta_0 < \theta^*$$

and the value of this strategy is given by

$$U^*(\theta_0) = \frac{\lambda}{r(\lambda + r)} \theta_0 g \left[\frac{\theta_0}{1 - \theta_0} \frac{g}{\lambda(c_h - c_l) - g} \right]^{\frac{r}{\lambda}} - (1 - \theta_0) c_l \frac{\lambda}{\lambda + r} \quad \text{for } \theta_0 \leq \theta^*. \quad (2)$$

If $\theta_0 \geq \theta^*$, the firm places the worker in the H task immediately.

(All proofs are in the appendix.)

Note that (1) has a natural interpretation. The worker is promoted when the expected flow payoffs in the L and H tasks are equal, that is

$$-(1 - \theta^*)\lambda c_l = g - (1 - \theta^*)\lambda c_h. \quad (3)$$

This follows from the assumption that learning about productivity is independent of task assignment. Therefore the assignment decision is determined solely by the effect on expected output. It is plausible that the arrival rate of crises would be faster in the H -task. In this case, workers would be promoted earlier than implied by equation (1). However, we have not obtained any additional insights from allowing for different rates of arrival of crises and therefore have not pursued this path.

3.4 Payoff with the monitoring strategy

When the employer monitors the worker he knows when the first crisis occurs, and immediately identifies the worker's type. Before the arrival of the first crisis no new information is generated, so there is no continuous updating of beliefs.

Let θ_0 be the prior that the worker is good. When the first crisis arrives, with probability θ_0 the (good) worker resolves the crisis and is promoted to the H -task, with the complementary probability he fails and leaves the firm. In either case the employer ceases to monitor him. Recall that monitoring has a flow cost of b per unit time. We prove in the appendix that

Theorem 3.2. *The value of the monitoring strategy is given by*

$$\tilde{U}(\theta_0) = \frac{1}{r(\lambda + r)} [\lambda\theta_0g - rb] - (1 - \theta_0)c_l \frac{\lambda}{\lambda + r}. \quad (4)$$

We assume that

$$\frac{g}{r} - \frac{b}{\lambda} > 0. \quad (5)$$

If not, even if the firm knew that the worker was good and even if the only way it could assign the worker to the H -task was by monitoring and observing him solve a crisis, it would prefer not to do so.

3.5 Optimal monitoring

Next we compare the two strategies to determine which one yields the greater expected payoff.

Theorem 3.3. *There is always a range $[0, \theta_a)$ and a range $(\theta_b, \theta^*]$ in which it is efficient not to monitor the worker.*

Theorem 3.3 establishes that workers who are very unlikely to be good and workers who are close to promotion will not be monitored, but it does not establish that firms will ever monitor workers in order to determine their quality.

Under what circumstances will the firm engage in monitoring? The following theorem addresses this question.

Theorem 3.4. *There is a range in which monitoring is preferred to non-monitoring if*

$$\left(\frac{g(\lambda + r)}{g\lambda - br} \right)^{\frac{\lambda+r}{r}} < \frac{\lambda(c_h - c_l) - g}{b}. \quad (6)$$

Condition (6) is not particularly informative. We can derive somewhat more informative conditions. Recall that, by assumption, both sides of the inequality are positive.

As $rb \rightarrow g\lambda$, the left-hand side goes to infinity while the right-hand side remains finite. Thus when monitoring costs are high, not surprisingly the firms never finds it efficient to monitor. On the other hand, when monitoring costs are sufficiently low, there is a range in which monitoring is efficient.

The gain from monitoring is that the firm is assured that it never places a bad worker in the H -task. The cost of doing so is given by $\lambda(c_h - c_l)$. Again not surprisingly, as this term gets large, there is always a range in




Figure 1: Monitoring and Task Assignment

which monitoring is efficient. When it gets small, or equivalently when the benefit from placing a good worker in the H -task gets small, monitoring is never efficient.

An increase in the frequency of crises, λ , lowers the left-hand-side and increases the right-hand side of inequality (6). Thus more frequent arrival of crises is associated with a larger range of other parameters consistent with monitoring. Conversely, a higher rate of time discounting is in many ways similar to a lower rate of arrival of crises and thus is associated with a more restricted set of parameters consistent with some monitoring.

The relation among θ , task assignment and monitoring is summarized in figure ??.

3.6 Example

Suppose that $r = .05$, $g = 1$, $\lambda = .2$, and $c_h - c_l = 10$, which implies that workers are promoted to the high task when there is a 50% probability that they are good. Then provided b is less than about .24, there will be a range in which there is monitoring. When $b = .2$, the firm does not monitor any worker for whom its estimate of θ is less than about .13, monitors those for whom its estimate of θ falls between about .13 and .35 and does not monitor those above this range.

4 Internal Labor Markets and Wage Profiles

To address the implications of the model for internal labor markets and wage profiles, we must model turnover and wage determination.

Turnover: Our model of turnover is simple. We interpret $\tilde{\theta}_0$, the value of θ when the worker first arrives at the firm, as the *ex ante* probability that the worker will be good at a randomly selected job (possibly from within a class of jobs). Any information about the ability of the worker to resolve crises is firm-specific. Therefore if the updated assessment falls to $\theta < \tilde{\theta}_0$, it

is efficient for the worker and firm to separate, and they do so. Otherwise it is efficient for the worker and firm to continue their relationship, and there is no separation. Thus in the version of the model developed so far, the worker remains with the firm unless he fails to resolve a crisis in which case the worker quits or is fired. Note that we do not really require that success or failure at the current firm provide no information about productivity elsewhere, only that it is more informative about productivity at this firm.

Wages: We do not fully model wage determination. The precise wage will depend on the institutional and informational assumptions we make. However, we assume that the wage is an increasing function of θ . We find this assumption plausible but recognize that it is not trivial. For example, in our baseline model, Nash bargaining would generate a reduction in the wage as workers move from the no-monitoring regime to the monitoring regime.

There are at least two reasons for maintaining the assumption that the wage is increasing in θ . The first is behavioral. Workers show a strong preference for upward-sloping wage profiles even when a flat profile would have a higher present value (Loewenstein and Sicherman, 1991). It is hardly plausible that they would respond well to being told, “Congratulations! Our opinion of you has improved. Therefore we are cutting your pay.” For reasons that are not well understood, within firms, nominal wage cuts are rare (Baker, Gibbs and Holmstrom, 1994a). Dohmen (2004) reports that workers wages are not cut even when they are demoted. A nontrivial minority even received cost-of-living increases.

The second reason is that monitoring is non-contractible. Therefore workers may be unwilling to accept lower wages in return for being monitored. In this case, monitoring would not be chosen efficiently as assumed above, but the broad characteristics of the model would not change.

We note also that since wages are adjusted discretely, implicit declines in wages due to the onset of monitoring might not be detectable in the data.

The Internal Wage Profile: Recall that there are four ranges in the data. For $\theta < \theta_a$, the worker is assigned to the L -task and is not monitored. For $\theta_a < \theta < \theta_b$, the worker is assigned to the L -task and monitored. For $\theta_b < \theta < \theta^*$, the worker is assigned to the L -task and not monitored. For $\theta > \theta^*$, the worker is assigned to the H -task and not monitored. The wage profile for workers who remain with the firm therefore depends on $\tilde{\theta}_0$, the value of θ when the worker is hired:

1. Workers with a low $\tilde{\theta}_0$ will be placed into the low no-monitoring range. If they remain with the firm, the firm gradually increases its assessment of θ , and therefore the wage, until $\theta = \theta_a$. At this point, the firm

begins monitoring the worker, and there is no updating of θ until a crisis comes along. Wages remain fixed until the worker either leaves the firm or is promoted to the high task.

2. Workers with a somewhat higher initial θ will be placed immediately into the monitoring range. Although their wage will generally be higher than $w(\theta_a)$ since most will have $\tilde{\theta}_0 > \theta_a$, in other respects they are similar to workers who started at a lower θ_0 and rose to $\theta = \theta_a$.
3. Workers with a yet higher $\theta_b \leq \tilde{\theta}_0 < \theta^*$ remain in the L -task and receive continuous wage increases until $\theta = \theta^*$, at which point they are promoted to the H -task and continue to receive wage increases that are asymptotic to the wage associated with $\theta = 1$.
4. Finally, any worker hired with a high initial θ is placed directly into the H -task and receives continuous wage increases in a manner analogous to those promoted from the upper no-monitoring range.

The relation between θ and seniority (t) at two different entry values of θ is summarized in 2. The lower line represents workers entering with low values of θ . The upper line represents those entering with relatively high values of θ but not sufficiently high to be placed in the H -task immediately. *The Hierarchical Structure:* Heretofore we have referred to tasks rather than to jobs. Yet in many organizations, collections of tasks that appear quite similar have different job titles (secretary I and II, tenured associate and full professor). In the empirical literature, hierarchies are sometimes determined by transition patterns across occupational titles as in Baker, Gibbs and Holmstrom (1994a). In our model, it is natural to define three occupation titles: $LT1$, consisting of workers in the low no-monitoring and monitoring zones, $LT2$, consisting of the high no-monitoring zone, and HT . Although $LT2$ is higher paid than $LT1$, neither feeds into the other. Instead both feed into HT . So $LT1$ and $LT2$ appear to share a location at the bottom of the hierarchy below HT .

5 Extensions

The biggest empirical weakness of the base model is it predicts counterfactually that promotions are concentrated at the top of the scale of each of the lower level jobs in the hierarchy. Yet, the evidence strongly suggests that promotions come from most parts of the wage distribution within a level of the hierarchy. The strong (and false) prediction is a direct consequence of assuming that monitoring is either complete or nonexistent. Moreover, all promotions end up at either the bottom of the upper level of the hierarchy




Figure 2: Seniority, Assessed Ability and Task Assignment

or at its top. This reflects the assumption that when monitoring occurs, it is fully informative.²

In the remainder of the paper we relax these two assumptions. We first consider the case where different monitoring intensities are possible. We then explore the consequences of allowing false negatives and false positives.

5.1 Partial Monitoring

We maintain the other assumptions of the model but assume that the firm can vary the effort with which it monitors the worker. The firm can choose the effort and the corresponding flow cost b of monitoring. If the worker resolves a crisis, the firm observes the success with probability $p = p(b)$. We write the inverse function $b = b(p)$ and assume that $b' \geq 0$ and $b'' \geq 0$.

In the appendix we derive a result (theorem A.1) that parallels theorem (3.3), which is restated below:

Theorem 5.1. *If $b'(0) > 0$ and $b''(p) > 0 \forall p \in [0, 1]$, there is always a range $[0, \theta_a]$ and a range $(\theta_b, \theta^*]$ in which it is efficient not to monitor the worker.*

Needless to say, if $b'(0)$ is sufficiently small, there will also be a range in which in which the firm does at least some monitoring, and if $b'(1)$ is sufficiently small, there will also be a range with complete monitoring. Depending on the shape of the $b(p)$ function, the solution can be bang-bang as in our base model.

The more interesting case is when monitoring increases smoothly between θ_a and some θ_A at which p equals 1 (full monitoring). It remains at 1 for $[\theta_A, \theta_B]$ and then decreases smoothly between θ_B and θ_b . Then if workers are hired with $\tilde{\theta}_0 < \theta_a$, as in the baseline case, they will not be monitored, but, unless the worker fails to resolve a crisis, the firm's assessment of θ will rise continuously until it reaches θ_a . Thereafter the firm continues to update θ . If no crisis is observed, θ rises towards θ_A .³

But the firm may observe the worker resolving a crisis, in which case she is immediately promoted. In the region between θ_a and θ_A , the probability

²We draw on a growing literature, which we do not attempt to summarize, that examines the internal labor markets of individual firms. See Baker, Gibbs and Holmstrom (1994a&b), Dohmen (2004), Dohmen, Kriechel and Pfann (2004), Flabbi and Ichino (2001), Gibbs and Hendricks (2004), Grund (2005), Kwon (2006), Medoff and Abraham (1980, 1981) and Treble et al (2001). Not all of the results are found in each of these papers, but we have tried to be careful not to include results with contradictory evidence.

³It appears to us that θ will reach θ_A only asymptotically, but we have not proved this.

of promotion is strictly increasing in θ since both the probability of being good and the probability of being observed solving a crisis rise with θ .

If the worker is hired with $\theta_a < \tilde{\theta}_0 < \theta_A$, the situation is similar except that she never experiences the no monitoring regime.

If the worker is hired with $\theta_A < \tilde{\theta}_0 < \theta_B$, the firm does not update θ except simultaneously with a separation or promotion.

If $\theta_B < \tilde{\theta}_0 < \theta_b$, the firm continuously updates θ and gradually reduces monitoring. It is clear that for θ close to θ_b , the probability of promotion must be lower than for θ close to θ_B , but we have not been able to establish whether the relation between the probability of promotion and θ is monotonic in this range and expect that it need not be. Finally we note that if $\theta_b < \tilde{\theta}_0 < \theta^*$, the worker is not monitored. In the absence of a failed crisis, θ is updated continuously until it reaches θ^* and the worker is promoted to the high job.

As in the base model, there are no promotions from the L -task to the H -task originating at $\theta < \theta_a$ or $\theta_b < \theta < \theta^*$, and all promotions are from one task to the other. Thus, we continue to have two separate “jobs” at the lower level of the hierarchy leading to the upper level.

5.2 False negatives

Because only bad workers fail to resolve crises in the base model, failures always cause separations. Workers never remain with the firm after a negative shock to θ . Therefore, there are never real wage decreases except possibly for the effects of macroeconomic shocks outside the model.

In this section, we show that if we allow good workers to solve crises only with some probability, $\gamma < 1$, then there is an initial probationary period during which any failure causes a separation. After the probationary period, the firm may respond to a failure by reducing its assessment of θ depending on the worker’s history.

To analyze this case, we return to the assumption that monitoring is a binary decision. We maintain the assumption that the firm and worker separate whenever θ falls below $\tilde{\theta}_0$, the value of θ when the worker joined the firm.

Theorem 5.2. *A sufficient condition for a failure to result in a separation is that the worker has not been observed previously to solve a crisis and*

$$t < -\lambda^{-1} \ln \frac{1 - \gamma}{2 - \gamma}. \quad (7)$$

It should be noted that condition (7) is only relevant for workers who spent their “probationary period” in the no-monitoring range. Intuitively, during this period, the firm’s beliefs about θ improve sufficiently that the failure is not sufficient to lower this assessment below its initial level. However, if the worker did not spend the requisite time in the no-monitoring range, then the firm’s assessment of θ will not have risen sufficiently to offset the reduction resulting from failure, except in the special case of a worker who was monitored and observed to solve a crisis.

Note also that the probation period will always result in a substantial fraction of bad workers separating from the firm. Condition (7) can be rewritten as

$$1 - e^{-\lambda t} > \frac{1}{2 - \gamma}. \quad (8)$$

The left-hand side of (44) is the probability that the worker will have faced a crisis before time t . The right-hand side is always greater than .5. So if good workers are very bad at solving crises, fifty percent of workers must have failed (and been fired) before workers reach the point that a failure does not cause separation. If good workers can solve half of the crises they face, then two-thirds of workers will have faced a crisis before reaching the level of seniority at which a failure does not lead to separation. Of these, all the bad workers and half the good workers would have been fired for failing to solve the crisis. And, of course, if good workers always solve crises, the “probationary period” lasts forever.

Finally, we should emphasize that surviving the probationary period does not mean having tenure, only that the first failure does not induce a separation. Multiple failures may still result in a separation.

Without specific assumptions about parameter values, it is impossible to determine what fraction of workers will ever experience a reduction in θ . It seems to us that in many settings the proportion is likely to be modest but consistent with the description “not rare,” but others may have different priors about γ , λ , the distribution of $\tilde{\theta}_0$ and the distribution of the other fundamental parameters of the model. What the model does predict strongly is that real wage reductions not associated with macroeconomic phenomena will be rare early in a worker’s tenure with the firm. We are not aware of any results on this point.

5.3 False positives

The assumption that only good workers solve crises produces the strong and empirically false result that monitored workers who resolve crises are

always promoted to the top of the next level of the hierarchy. In this subsection we consider what happens if monitoring can produce false positives, that is the worker can appear to have resolved a crisis when none existed. Let δ be the arrival rate of false positives and let $\mu = \delta / (\delta + \lambda)$ be the proportion of apparent crises that are not really crises. Then if the firm observes that a worker with $\theta = \theta_0$ has “resolved a crisis,” the firm’s assessment of θ will be updated to

$$\theta = \frac{\theta_0}{(1 - \mu)\theta_0 + \mu}. \quad (9)$$

It is both intuitive and straightforward to show that if, when δ is 0, there is an interval of θ_0 for which the firm monitors the worker, then for δ sufficiently small, there is still an interval for which the firm monitors the worker and, if the worker appears to resolve a crisis is promoted to the H -task although the updated θ is less than one.

False positives leading to promotion to the H job We begin by assuming that false positives are sufficiently rare that within the monitoring range, if the firm believes the worker has solved a crisis, it promotes him to the H job. Having derived the monitoring range, we will then have to verify that this assumption holds.

Suppose that after observing a “resolved crisis” $\theta > \theta^*$ and normalize the starting time after promotion to 0. Then after promotion, we have

$$\begin{aligned} U(\theta) &= \theta \int_0^{\infty} e^{-rt} g dt + (1 - \theta) \left[\int_0^{\infty} g e^{-rt} e^{-\lambda t} dt - c_h \int_0^{\infty} e^{-rt} \lambda e^{-\lambda t} dt \right] \\ &= \frac{\theta g}{r} + \frac{(1 - \theta)(g - \lambda c_h)}{\lambda + r}. \end{aligned} \quad (10)$$

$$(11)$$

Suppose next that the next apparent crisis occurs at time t , then the value of the monitoring strategy is

$$\begin{aligned} U(\theta_0, t) &= -b \int_0^t e^{-rt} dt - c_l (1 - \theta_0) e^{-rt} (1 - \mu) + [(1 - \mu)\theta_0 + \mu] U(\theta) \\ &= -\frac{b}{r} + \left(\frac{b}{r} - c_l (1 - \theta_0) (1 - \mu) + [(1 - \mu)\theta_0 + \mu] U(\theta) \right) e^{-rt}. \end{aligned} \quad (13)$$

Integrating over t

$$U(\theta_0) = \quad (14)$$

$$\begin{aligned} & -\frac{b}{r} + \left(\frac{b}{r} - c_l(1-\theta_0)(1-\mu) + [(1-\mu)\theta_0 + \mu]U(\theta) \right) \int_0^\infty e^{-rt}(\lambda + \delta) e^{-(\lambda+\delta)t} dt \\ & = -\frac{b}{r} + \frac{(b - c_l(1-\theta_0)(1-\mu)r + [(1-\mu)\theta_0 + \mu]rU(\theta))(\lambda + \delta)}{r(\lambda + \delta + r)} \end{aligned} \quad (16)$$

gives the value of the monitoring strategy.

Substituting for μ , $U(\theta)$ and θ

$$U(\theta_0) = \frac{(\lambda + \delta)\theta_0 g}{r(\lambda + \delta + r)} - \frac{b}{(\lambda + \delta + r)} - \frac{(1 - \theta_0)\lambda c_l}{(\lambda + \delta + r)} + \frac{\delta(g - \lambda c_h)(1 - \theta_0)}{(\lambda + r)(\lambda + \delta + r)}. \quad (17)$$

When $\delta = 0$, this reduces to the earlier expression.

The existence of false positives does not affect the value of the waiting strategy. Therefore, it will be optimal to monitor if

$$\begin{aligned} & \frac{(\lambda + \delta)\theta_0 g}{r(\lambda + \delta + r)} - \frac{b}{(\lambda + \delta + r)} - \frac{(1 - \theta_0)\lambda c_l}{(\lambda + \delta + r)} + \frac{\delta(g - \lambda c_h)(1 - \theta_0)}{(\lambda + r)(\lambda + \delta + r)} \\ & > \frac{\lambda}{r(\lambda + r)} \theta_0 g \left(\frac{\theta_0}{1 - \theta_0} \frac{g}{\lambda c - g} \right)^{\frac{r}{\lambda}} - (1 - \theta_0) c_l \frac{\lambda}{\lambda + r} \end{aligned} \quad (18)$$

or

$$rb < \frac{\delta r(g - \lambda c(1 - \theta_0)) + \lambda \theta_0 g(\delta + \lambda + r) \left(1 - \left(\frac{\theta_0 g}{(1 - \theta_0)(\lambda c - g)} \right)^{\frac{r}{\lambda}} \right)}{r + \lambda} \quad (19)$$

As we would expect, the rhs is declining in δ . Higher arrival rates of false positives shrink the range of θ_0 for which monitoring is optimal. The rest goes through. The rhs is zero at $\theta_0 = \theta^*$ and negative at $\theta = 0$ although this last result is misleading because $\theta = 0$ is inconsistent with switching to the high task after a positive result.

If all positives result in promotion to the high task, then we require that the inequality be reversed at the lowest θ_0 leading to promotion to the H job. This lowest θ_0 is given by

$$\theta_m = \delta \frac{\lambda c - g}{\lambda(c\delta + g)}.$$

Substituting into condition (19) gives the following condition

$$rb > \frac{-\delta rg \frac{\lambda c - g}{c\delta + g} + \delta \frac{\lambda c - g}{c\delta + g} g (\lambda + \delta + r) \left(1 - \left(\frac{\delta}{\lambda + \delta}\right)^{\frac{r}{\lambda}}\right)}{r + \lambda}. \quad (20)$$

If δ is sufficiently close to 0, then despite false positives, all monitored workers who appear to resolve a crisis will be promoted to the high job. Note that for δ close to 0, the conditions under which monitoring is optimal will be similar to those with no false positives.

If δ is sufficiently large, workers at the lower end of the monitoring range who appear to have solved a crisis may be promoted to the upper no-monitoring zone in the L -task. However, we can show that except in a knife-edge case, successful workers at the top end of the monitoring range will be promoted to the H -task.

It seems likely that, for δ sufficiently large and b sufficiently small, the updating of θ could leave the worker in the monitoring range. Since as b goes to zero, the entire range of $\theta < \theta^*$ is monitored and since as δ gets large the updated value of θ remains close to the value prior to the apparent crisis, it seems that this possibility must exist, but we have not explored it.

Finally, we have not explored formally the case in which crises are differentially informative. In this subsection we have assumed that crises are either real or false. Plausibly, crises differ in the likelihood that a bad worker can resolve them. Thus some crises would be more informative than others about worker ability. Solving a more informative crisis leads to a larger upward revision of θ .

Example In the example below, we set the value of output in the H -task to 1, the arrival rate of crises to 1 and the cost of a failed crisis to be 2.1 higher in the H -task than in the L -task. The arrival rate of false positives is also 1, the discount rate is .1 and the flow cost of monitoring is .125. The results are the following:

1. For θ less than about .34, the worker is not monitored. If he is not observed to have failed to resolve the crisis, the assessment of θ is continually increased until it reaches .34
2. For θ between roughly .34 and .355, the worker is monitored. If he appears to resolve a crisis he is promoted to the no-monitoring region; otherwise the assessment of θ is unchanged.
3. For θ between roughly .355 and .43, the worker is monitored. If he appears to resolve a crisis is promoted to the H -task; otherwise the assessment of θ is unchanged.

4. For θ between roughly .43 and .524 the worker is not monitored. If he is not observed to have failed to resolve a crisis, the assessment of θ is continually increased until it reaches .524 and the worker is assigned to the H -task.

Note that the internal labor market path will depend on the level of θ at which workers are typically hired. If most workers are hired when the probability of the worker being good is less than one-third, they will begin in the no-monitoring zone and remain for a while unless they are shown to be bad at the job. Eventually, they will hit the bottom of the monitoring zone where they will remain until they appear to face a crisis. After the apparent crisis, they will either be shown to be bad at this job and separate from it or will be promoted to the higher no-monitoring zone where they will remain and rise within the zone until they are promoted to the H -task or separate. Workers who are hired into the monitoring zone may be promoted into the monitoring zone or the H -task depending on the value of θ when they entered.

5.4 Monitoring to Correct Mistakes

So far we have assumed that the sole purpose of monitoring workers is to learn about their type. Monitoring the worker may also permit the firm to mitigate or eliminate the cost of any mistakes.

Allowing monitoring to eliminate costly mistakes in the L -task has no important implications provided that $b > \lambda c_l$ so that it is not efficient to monitor simply to catch mistakes even if the worker is known to be bad. The proof that there cannot be monitoring when θ is close to 0 and when θ is close to but less than θ^* goes through *mutatis mutandum*. Of course, by offsetting some of the cost of monitoring, the ability to correct mistakes will increase the range of parameters for which monitoring is optimal.

Monitoring to catch mistakes in the H -task is somewhat more interesting. We limit ourselves to a few remarks and do not analyze this case fully. We note that if monitoring fully eliminates mistakes, then a monitored worker has net output $-b_l$ in the L -task and $g - b_h$ in the H -task. In this case, monitoring will be used at most in one task. The case where monitoring is only used in the L -task was discussed in the previous paragraph. Suppose monitoring occurs only in the H -task, and that $b > g + \lambda c_l$. Then for low values of θ , workers are assigned to the L -task and are not monitored. When θ becomes sufficiently high, the worker is assigned to the H -task and monitored until a crisis arises.⁴ Workers with sufficiently high

⁴For $\theta = 0$, it is efficient to assign the worker to the L -task and not monitor him.

θ are assigned to the H -task and never monitored.

If monitoring in the H -task lowers but does not eliminate the cost of unresolved crises or, equivalently, eliminates the cost some, but not all, of the time, monitoring in the H -task is efficient if

$$b_h < (1 - \theta) \lambda (c_h - c_m) \tag{21}$$

where c_m is the cost of an unresolved crisis when the worker is monitored in the H -task. If $c_h - c_m$ is sufficiently large relative to b_h , there will be a range of θ for which, if workers are assigned to the H -task, they will be monitored.

The theorem below addresses the case where $b_h = b$ and monitoring does not affect the cost of unresolved crises in the L -task. The assumption about g ensures that a firm that monitored all workers regardless of θ would assign workers with low values of θ to the L -task and assign those with high values of θ to the H -task.

Theorem 5.3. *If $b > 0$, $g - \lambda c_m < -\lambda c_l$, and $b/(\lambda(c_h - c_m))$ is sufficiently small, there is always a range $[0, \theta_a)$ and a range $(\theta_b, \theta^*]$ in which it is efficient not to monitor the worker and assign him to the L -task and a range (θ^*, θ_c) in which it is efficient to monitor the worker and assign him to the H -task.*

We do not explore formally the conditions under which a monitoring range in the L -task exists. For sufficiently low b , monitoring will be desirable. We have established that whether monitoring and assignment to the L -task is more profitable than monitoring and assignment to the H -task depends on θ but not on b . Therefore for sufficiently low b , there will be values of θ for which workers will be monitored in the L -task.

Thus it is possible to have monitoring at both the intermediate range of the L -task and the bottom of the H -task, albeit with different goals. Perhaps strikingly, workers will never be promoted to an unmonitored range of the H -task unless they are known to be good ($\theta = 1$) although workers may be hired into this range.

We note also that only workers who are hired into the upper no-monitoring range of the L -task will ever be monitored in the H -task. It is therefore unclear whether workers who are monitored at the bottom of the H -task would be viewed as at the bottom of a new job scale or at the top of the one into

The existence of a no-monitoring range follows by continuity. In the H -task, the value of monitoring arises solely from the ability to catch mistakes. Therefore, if monitoring occurs in the H -task, it will always be for the lowest values of θ associated with that task.

which they were hired. In the latter case, we again have the phenomenon of workers being hired into a job classification receiving regularly scheduled pay increases until they “hit the top of the scale” and only receiving further pay increases following promotion.

6 Discussion and Conclusion

If we allow for partial monitoring, our model permits the following monitoring stages as a function of θ :

No Monitoring
 Partial Monitoring
 Full Monitoring
 Partial Monitoring
 No Monitoring
 High Task

Not all stages need exist. For the no monitoring range to exist, we require that $b'(0) > 0$, so that is that it is costly to do even a little monitoring. For partial monitoring, we require $b'' > 0$, and for the existence of full monitoring, we require that $b'(1)$ be sufficiently small.

Therefore, the precise nature of the internal labor market depends on the monitoring technology. We should not be surprised by variation in internal labor markets across companies and types of workers.

If monitoring is very expensive, wages are likely to be determined largely by observable proxies for productivity such as education and seniority. If monitoring is inexpensive and crises are very informative, there is likely to be little wage growth within job assignment. At intermediate monitoring costs, wages may rise formulaically within some job assignment until some maximum wage. With partial monitoring, they climb formulaically except for “fast-trackers” who get a boost from resolving a crisis.

Our model can be contrasted with that of Gibbons and Waldman (1999, 2006). The major difference is that in our model, learning comes in large chunks and when it does not come, either there is no learning or updating depends only on the passage of time. In contrast, in Gibbons and Waldman, firms continuously receive information about workers which allows them to distinguish among them. We do not view these approaches as strict alternatives. Clearly, information can come in both forms.

Our approach has both advantages and disadvantages. On the positive side, as discussed above, in many jobs wages are determined solely by objec-

tive measures such as tenure and education that are only very imperfectly related to productivity. In our model, wages are explained *perfectly* by $\tilde{\theta}_0$, task assignment and seniority. If we consider education to be an imperfect proxy for θ_0 , the model is strongly consistent with this regularity. However, on the negative side, it is too strong. There are many settings in which wages are determined in part by subjective performance evaluations even though much of the variation in wages is explained by education, seniority and tier in the hierarchy.

Our model is also consistent with both the steady increase in wages (at least up to some maximum) that often accompanies seniority and the large jumps in wages often associated with promotions. The strong association between large wage increases and promotions does not arise naturally in Gibbons/Waldman.

In both models demotions are rare, albeit for different reasons. In Gibbons and Waldman workers acquire human capital over time. The worker will be demoted only if new information is sufficiently negative to outweigh the growth in human capital and if the worker's productivity previously placed him just above the cutoff between two levels of the hierarchy. In contrast, demotions are rare in our model because negative information usually causes a separation. In Gibbons/Waldman demotions should be concentrated among recently promoted workers. In our model, demotions should be particularly rare among workers with low seniority. We know of no empirical findings on these two issues.

In contrast, the Gibbons/Waldman model is better able to explain the frequency of real wage decreases. In our model, in the absence of macroeconomic shocks, real wage decreases are like demotions. Bad news is infrequent and generally results in a separation not in the worker remaining with the firm but with a lower estimate of θ . In our model, small real wage decreases happen only when there are negative macroeconomic shocks.

Finally, we note that technology has made monitoring easier. In almost any model including this one, this will make pay-for-performance more common. Consistent with this expectation, the proportion of British workers receiving performance pay rose from 16 to 32 percent of workers between 1988 and 1994 (Manning and Saidi, 2008). But our model suggests some less obvious effects. Reducing the cost of monitoring could shift the nature of the hierarchy. When monitoring is relatively expensive, as discussed above, we can have two apparent jobs at the low task, one comprised of workers in or below the full monitoring range and one comprised of workers above the full monitoring range, with both jobs leading directly to the high task and relatively little "lateral" movement. When monitoring becomes less ex-

pensive, particularly if it becomes easier to observe less informative crises, there will be more movement from the lower range of the low task into the upper range of the low task so that the low task now appears more like a single job in the hierarchy. Thus we believe the model could be used to help explain how hierarchical structures change over time.

7 References

Abnagale, Jr., Frank W. with Stan Redding, *Catch Me if You Can*, New York: Grosset & Dunlap, 1980.

Akerlof, George A. and Lawrence F. Katz, "Do Deferred Wages Dominate Involuntary Unemployment as a Worker Discipline Device?" National Bureau of Economic Research Working Paper No. 2025, 1986.

Autor, David H., Frank Levy and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118 (November 2003): 1279-1333.

Baker, George, Michael Gibbs, and Bengt Holmstrom. "The Internal Economics of the Firm: Evidence from Personnel Data." *Quarterly Journal of Economics*, 109 (November 1994a): 881-919.

Baker, George; Michael Gibbs, and Bengt Holmstrom, "The Wage Policy of a Firm," *Quarterly Journal of Economics*. 109 (November 1994b): 921-55.

Becker, Gary S. and George J. Stigler, "Law Enforcement, Malfeasance, and Compensation of Enforcers," *Journal of Legal Studies*, 3 (January 1974): 1-18

Bjerk, David, "Glass Ceilings or Sticky Floors? Statistical Discrimination in a Dynamic Model of Hiring and Promotion," *Economic Journal*, 118 (July 2008): 961-982.

Dickens, William T., Lawrence F. Katz, Kevin Lang and Lawrence H. Summers, "Employee Crime and the Monitoring Puzzle," *Journal of Labor Economics*, 7 (July 1989): 331-347.

Doeringer, Peter and Michael A. Piore, *Internal Labor Markets and Manpower Analysis*, Lexington, MA: DC Heath, 1971.

Dohmen, Thomas J., "Performance, Seniority and Wages: Formal Salary Systems and Individual Earnings," *Labour Economics*, 11 (December 2004): 741-63.

Dohmen, Thomas J., Ben Kriechel and Gerard A Pfann, "Monkey Bars and Ladders: The Importance of Lateral and Vertical Job Mobility," *Journal of Population Economics*, 17 (June 2004): 193-228.

Flabbi, Luca and Andrea Ichino, "Productivity, Seniority and Wages: New Evidence from Personnel Data," *Labour Economics*, 8 (June 2001): 359-87.

Gibbs, Michael and Wallace Hendricks, "Do Formal Salary Systems Really Matter?" *Industrial and Labor Relations Review*, 58 (October 2004): 71-92.

Gibbons, Robert and Michael Waldman, "A Theory of Wage and Promotion Dynamics Inside Firms," *Quarterly Journal of Economics*, 114 (November 1999a): 1321-58.

Gibbons, Robert and Michael Waldman, "Enriching A Theory of Wage and Promotion Dynamics Inside Firms," *Journal of Labor Economics*, 24 (January 2006): 59-108.

Gibbons, Robert and Michael Waldman, "Careers in Organizations: Theory and Evidence," in Orley C. Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol 3B., Amsterdam: Elsevier, 1999b, 2373-438.

Grund, Christian, "The Wage Policy of Firms: Comparative Evidence for the US and Germany from Personnel Data," *International Journal of Human Resource Management*, 16 (January 2005): 104-119.

Kahn, Shulamit and Kevin Lang, "Constraints on the Choice of Work Hours: Agency vs. Specific Capital," *Journal of Human Resources*, 27 (Fall 1992), 661-678.

Kahn, Shulamit and Kevin Lang, "Sources of Hours Constraints: Evidence from Canada," *Canadian Journal of Economics*, 28 (November 1995): 914-28.

Kwon, Illong, "Incentives, Wages and Promotions: Theory and Evidence," *Rand Journal of Economics*, 37 (Spring 2006): 100-20.

Lazear, Edward P., "Why Is There Mandatory Retirement?" *Journal of Political Economy*, 87 (December 1979): 1261-1284.

Lazear, Edward P., "Agency, Earnings Profiles, Productivity, and Hours Restrictions," *American Economic Review*, 71 (September 1981): 606-620.

Loewenstein, George and Nachum Sicherman, "Do Workers Prefer Increasing Wage Profiles?" *Journal of Labor Economics*, 9 (January 1991): 67-84.

Manning, Alan and Farzad Saidi, "Understanding the Gender Pay Gap: What's Competition Got to Do with It?" *Industrial & Labor Relations Review*, 63 (2010) No. 4, article 7.

Medoff, James L and Katharine G. Abraham, "Experience, Performance, and Earnings," *Quarterly Journal of Economics*, 95 (December 1980): 703-36.

Medoff, James L and Katharine G. Abraham, "Are those paid more really more productive?" *Journal of Human Resources*, 16 (Spring 1981): 186-216.

Murphy, Kevin M. and Robert H. Topel, "Efficiency Wages Reconsidered: Theory and Evidence," in Yoram Weiss and Gideon Fishelson, eds., *Advances in the Theory and Measurement of Unemployment*, New York : St. Martin's Press, 1990,204-40.

Treble, John, Edwin van Gameraan, Sarah Bridges, and Tim Barmby, "The Internal Economics of the Firm: Further Evidence from Personnel Data," *Labour Economics*, 8 (December 2001): 531-552.

A Proofs

A.1 Proof of theorem 3.1

Suppose a worker with prior θ_0 has been put in a job at time 0 and has not failed until time t . If the worker is good, then non-failure occurs with probability 1, and if he is bad then the probability of non-failure is equal to the probability that a crisis has not occurred by time t . Thus the employer's updated belief about the worker's type (i.e., the updated probability that the worker is good) is:

$$\theta(t, \theta_0) = \frac{\theta_0}{\theta_0 + [1 - p(t)](1 - \theta_0)} \quad (22)$$

which for future reference we rearrange as

$$1 - p(t) = \frac{\theta_0[1 - \theta(t)]}{\theta(t)[1 - \theta_0]} \quad (23)$$

Let $\bar{\theta}$ be the threshold such that a worker who was initially placed in an L-job is promoted to the H-job when $\theta(t) \geq \bar{\theta}$. We will show below that $\bar{\theta}$ is independent of θ_0 . Define $\bar{t}(\theta_0)$ such that $\theta(\bar{t}(\theta_0), \theta_0) = \bar{\theta}$. Below we will suppress the arguments in $\bar{t}(\cdot)$, $\theta(\cdot)$ etc.

If $\theta_0 < \bar{\theta}$, then the employer puts the worker in the L-job, and promotes him if he has not failed by time \bar{t} . Thus a good worker produces nothing between times 0 and \bar{t} , and thereafter produces a flow output of g . A bad worker fails before promotion with probability $p(\bar{t})$. With probability $[1 - p(\bar{t})]$ he produces nothing until \bar{t} , and thereafter produces g until the first crisis arrives, at which time he produces $-c_h$ and is fired. Hence the expected payoff from the N-strategy with prior θ_0 and threshold $\bar{\theta}$ is

$$\begin{aligned}
U(\theta_0, [1 - p(\bar{t})]) &= \theta_0 \int_{\bar{t}}^{\infty} g e^{-rt} dt \\
&+ (1 - \theta_0)[1 - p(\bar{t})] \left[\int_{\bar{t}}^{\infty} \lambda e^{-\lambda(t-\bar{t})} e^{-rt} [-c_h] dt + e^{-r\bar{t}} \frac{g}{\lambda + r} \right] \\
&- (1 - \theta_0)c_l \int_0^{\bar{t}} \lambda e^{-(\lambda+r)t} dt \tag{24}
\end{aligned}$$

$$\begin{aligned}
&= e^{-r\bar{t}} \left\{ \frac{1}{r} \theta_0 g - \frac{1}{\lambda + r} [1 - p(\bar{t})] (1 - \theta_0) (\lambda(c_h - c_l) - g) \right\} \\
&- (1 - \theta_0)c_l \frac{\lambda}{\lambda + r} \tag{25}
\end{aligned}$$

Note that $e^{-r\bar{t}} = e^{[-\lambda\bar{t}] \frac{r}{\lambda}} = [1 - p(\bar{t})]^{\frac{r}{\lambda}}$, which substituted in (24) yields

$$\begin{aligned}
U(\theta_0, [1 - p(\bar{t})]) &= \frac{1}{r} [1 - p(\bar{t})]^{\frac{r}{\lambda}} \theta_0 g \tag{26} \\
&- \frac{1}{\lambda + r} [1 - p(\bar{t})]^{\frac{\lambda+r}{\lambda}} (1 - \theta_0) (\lambda(c_h - c_l) - g) - (1 - \theta_0)c_l \frac{\lambda}{\lambda + r}
\end{aligned}$$

The employer maximizes this payoff by choosing $\bar{\theta}$, or equivalently \bar{t} or $p(\bar{t})$. Maximizing $U(\theta_0, [1 - p(\bar{t})])$ in (26) with respect to $[1 - p(\bar{t})]$ we obtain the first order condition:

$$\begin{aligned}
0 &= \frac{1}{r} \frac{r}{\lambda} [1 - p(\bar{t})]^{\frac{r}{\lambda} - 1} \theta_0 g - \frac{1}{\lambda + r} \frac{\lambda + r}{\lambda} [1 - p(\bar{t})]^{\frac{r}{\lambda}} (1 - \theta_0) (\lambda(c_h - c_l) - g) \\
[1 - p(\bar{t})]^{-1} \theta_0 g &= (1 - \theta_0) (\lambda(c_h - c_l) - g) \tag{27}
\end{aligned}$$

Let (27) be solved at $\bar{t} = t^*$, and correspondingly $\bar{\theta} = \theta^*$ etc. Using (23), (27) simplifies to

$$\begin{aligned}
g &= (\lambda(c_h - c_l) - g) \frac{[1 - \theta^*]}{\theta^*} \\
\theta^* &= \frac{\lambda(c_h - c_l) - g}{\lambda(c_h - c_l)}
\end{aligned}$$

It can be checked that the second derivative of $U(\theta_0, [1 - p(\bar{t})])$ in (26)

is strictly negative at the solution, as follows:

$$\begin{aligned}
\frac{\partial^2 U}{\partial [1-p(\bar{t})]^2} &= \frac{r-\lambda}{\lambda^2} [1-p]^{(\frac{r}{\lambda}-2)} \theta_0 g - \frac{r}{\lambda^2} (1-\theta_0) [1-p]^{(\frac{r}{\lambda}-1)} (\lambda(c_h - c_l) - g) \\
&= \lambda^{-2} [1-p]^{(\frac{r}{\lambda}-1)} (-\lambda \theta_0 g + r ([1-p]^{-1} \theta_0 g - (1-\theta_0) (\lambda(c_h - c_l) - g))) \\
&= -\frac{\theta_0 g}{\lambda [1-p]} \\
&< 0
\end{aligned}$$

so this is indeed a strict maximum.

Note also that the optimal threshold θ^* is independent of the prior θ_0 , from which it follows that a worker entering with prior $\theta_0 \geq \theta^*$ will be placed directly in the H-job. At the optimum, the employer's expected payoff from a new worker with prior $\theta_0 \leq \theta^*$ can be obtained by making the appropriate substitutions in (26) to give:

$$U^*(\theta_0) = \frac{\lambda}{r(\lambda+r)} \theta_0 g \left[\frac{\theta_0}{1-\theta_0} \frac{g}{\lambda(c_h - c_l) - g} \right]^{\frac{r}{\lambda}} - (1-\theta_0) c_l \frac{\lambda}{\lambda+r} \quad \text{for } \theta_0 \leq \theta^* \quad (28)$$

It follows directly that U^* is increasing in θ_0 . For $\theta_0 \geq \theta^*$. It is straightforward to check that the expected payoff is then

$$U^*(\theta_0) = \frac{1}{r} \theta_0 g - \frac{(1-\theta_0)(\lambda c_h - g)}{\lambda+r} > U^*(\theta^*) \quad \text{for } \theta_0 > \theta^*$$

■

A.2 Proof of Theorem 3.2

When the first crisis arises, the firm gets $(\theta_0 \frac{g}{r} - (1-\theta_0) c_l)$. Expected discounting is

$$\int_0^\infty e^{-rt} \lambda e^{-\lambda t} dt = \frac{\lambda}{\lambda+r}.$$

Expected discounted monitoring costs are

$$\begin{aligned}
b \int_0^\infty e^{-rt} e^{-\lambda t} dt &= \frac{b}{\lambda+r} \\
\tilde{U}(\theta_0) &= \left(\theta_0 \frac{g}{r} - (1-\theta_0) c_l \right) - \frac{b}{\lambda+r}.
\end{aligned}$$

Rearranging terms yields (4).

A.3 Proof of Theorems 3.3 and 3.4

A.3.1 Preliminaries:

Given a prior θ , it is better to monitor the worker than not monitor if

$$\begin{aligned} \hat{U}(\theta) &\geq U^*(\theta) \\ \Rightarrow \frac{1}{r(\lambda+r)} [\lambda\theta g - rb] &\geq \frac{1}{r(\lambda+r)} \lambda\theta g \left[\frac{\theta}{1-\theta} \frac{g}{\lambda(c_h - c_l) - g} \right]^{\frac{r}{\lambda}} \\ \Rightarrow \lambda\theta g \left[1 - \left(\frac{\theta}{1-\theta} \right)^{\frac{r}{\lambda}} \left(\frac{g}{\lambda(c_h - c_l) - g} \right)^{\frac{r}{\lambda}} \right] &\geq rb \end{aligned} \quad (29)$$

Name the left-hand-side of (29) $Z(\theta)$:

$$Z(\theta) = \lambda\theta g \left[1 - \left(\frac{\theta}{1-\theta} \right)^{\frac{r}{\lambda}} \left(\frac{g}{\lambda(c_h - c_l) - g} \right)^{\frac{r}{\lambda}} \right] \quad (30)$$

A.3.2 Theorem 3.3

If $\theta = 0$ or $\theta = \theta^*$, $Z(\theta) = 0$, which proves the existence of the lower and upper no-monitoring ranges.

A.3.3 Theorem 3.4

Monitoring is more profitable than no-monitoring at θ if $Z(\theta) \geq rb$. First we prove that Z is concave. We have

$$\frac{dZ}{d\theta} = \lambda g - g \left(\lambda + \frac{r}{1-\theta} \right) \left(\frac{\theta}{1-\theta} \right)^{\frac{r}{\lambda}} \left(\frac{g}{\lambda(c_h - c_l) - g} \right)^{\frac{r}{\lambda}} \quad (31)$$

and

$$\frac{d^2Z}{d\theta^2} = -(r + \lambda) g r \left(\frac{\theta}{1-\theta} \right)^{\frac{r}{\lambda}} \frac{\left(\frac{g}{\lambda(c_h - c_l) - g} \right)^{\frac{r}{\lambda}}}{\lambda(1-\theta)^2 \theta} < 0.$$

Next we establish conditions under which $Z(\theta)$ exceeds rb , that is there is a range in which monitoring is preferred to no-monitoring. $Z(\theta)$ is a maximum when $dZ/d\theta = 0$, which implies

$$\left(\frac{\theta}{1-\theta} \right) \left(\frac{g}{\lambda(c_h - c_l) - g} \right) = \left[\frac{\lambda(1-\theta)}{\lambda(1-\theta) + r} \right]^{\frac{\lambda}{r}} \quad (32)$$

Let $\theta = \hat{\theta}$ solve (32). Note that the right-hand side of (32) is less than unity, which in turn implies that

$$\begin{aligned} \hat{\theta} \left(\frac{\lambda(c_h - c_l)}{\lambda(c_h - c_l) - g} \right) &< 1 \\ \Rightarrow \hat{\theta} &< \frac{\lambda(c_h - c_l) - g}{\lambda(c_h - c_l)} = \theta^* \end{aligned}$$

substituting (32) in (30) gives

$$Z(\hat{\theta}) = \lambda \hat{\theta} g \left[\frac{r}{\lambda(1 - \hat{\theta}) + r} \right]$$

which in (29) yields

$$\begin{aligned} \tilde{U}(\hat{\theta}) &\geq U^*(\hat{\theta}) \\ \Leftrightarrow \frac{\lambda \hat{\theta} g}{\lambda(1 - \hat{\theta}) + r} &> b \\ \Leftrightarrow \hat{\theta} &> \frac{b(\lambda + r)}{\lambda(g + b)}. \end{aligned} \tag{33}$$

To prove the theorem, note that the left-hand-side of (32) is increasing in θ while the right-hand-side is decreasing in θ . Thus (33) will be satisfied if and only if the lhs of (32) is less than the rhs at $\theta = \frac{b(\lambda+r)}{\lambda(g+b)}$. This condition on rearrangement yields the theorem. ■

A.4 Proof of Theorem 5.1

We prove a somewhat expanded version of the theorem, which follows after some preliminary constructions and a lemma.

Let $p(\theta)$ be the optimal monitoring program, and let $U(\theta)$ be the value of the optimal program starting from θ . We assume that $p(\theta)$ is right-continuous in θ . Continue to assume that the worker is promoted at $\theta^* = \frac{\lambda(c_h - c_l) - g}{\lambda(c_h - c_l)}$. Observe that $p(\theta) = 0$ must hold for $\theta \geq \theta^*$. Denote the flow cost of monitoring at intensity p by $b(p)$. Assume $b'(p) > 0$, $b''(p) > 0 \forall p \in [0, 1]$.

Now consider the following program $\langle \theta_0, \hat{p}(\theta), t \rangle$ starting from $\theta_0 < \theta^*$: monitoring occurs at the constant rate \hat{p} for time t . After t we revert to the optimal program $p(\theta)$. Let θ_1 be the worker's updated assignment if he has not failed by time t . Note that

$$\begin{aligned}
\theta_1 &= \frac{\theta_0 e^{-\lambda \hat{p} t}}{\theta_0 e^{-\lambda \hat{p} t} + (1 - \theta_0) e^{-\lambda t}} \\
&= \frac{\theta_0}{\theta_0 + (1 - \theta_0) e^{-(1 - \hat{p}) \lambda t}}
\end{aligned} \tag{34}$$

Therefore the rate of change of θ_1 with respect to time is given by:

$$\frac{\partial \theta_1}{\partial t} = \frac{\lambda(1 - \hat{p})\theta_0(1 - \theta_0)e^{-(1 - \hat{p})\lambda t}}{[\theta_0 + (1 - \theta_0)e^{-(1 - \hat{p})\lambda t}]^2}$$

Taking limits as $t \rightarrow 0$ and considering arbitrary θ , this implies

$$\frac{d\theta}{dt} = \lambda(1 - \hat{p})\theta(1 - \theta) \tag{35}$$

which depends only on θ and the monitoring intensity at θ .

For small t , the value of the program $\langle \theta_0, \hat{p}(\theta), t \rangle$ is given by

$$\begin{aligned}
\hat{U}(\theta_0, \hat{p}, t) &\cong \theta_0 [\hat{p}\lambda t U(1) + (1 - \hat{p}\lambda t) U(\theta_1)] \\
&\quad + (1 - \theta_0) [(1 - \lambda t) U(\theta_1)] - b(\hat{p})t
\end{aligned} \tag{36}$$

Note that the function $\hat{U}(\theta_0, \hat{p}, t)$ differs from the function $U(\theta)$ to the extent that it incorporates the perturbation implied by $\langle \theta_0, \hat{p}(\theta), t \rangle$. So

$$\begin{aligned}
\hat{U}(\theta_0, \hat{p}, t) - U(\theta_0) &\cong \theta_0 \hat{p} \lambda t U(1) - [\theta_0 \hat{p} \lambda t + (1 - \theta_0) \lambda t] U(\theta_1) \\
&\quad + [U(\theta_1) - U(\theta_0)] - b(\hat{p})t
\end{aligned} \tag{37}$$

Dividing both sides by t and taking limits as $t \rightarrow 0$ we get

$$\begin{aligned}
\lim_{t \rightarrow 0} \frac{\hat{U}(\theta_0, \hat{p}, t) - U(\theta_0)}{t} &\equiv \theta_0 \hat{p} \lambda U(1) - [\theta_0 \hat{p} \lambda + (1 - \theta_0) \lambda] U(\theta_0) \\
&\quad + U'(\theta_0) \frac{d\theta}{dt} \Big|_{\theta=\theta_0} - b(\hat{p})
\end{aligned} \tag{38}$$

since $\theta_1 \rightarrow \theta_0$ as $t \rightarrow 0$. The equivalence in (38) reflects the fact that the approximation in (36) holds for all values of \hat{p} . Using (35), (38) reduces to

$$\begin{aligned}
\lim_{t \rightarrow 0} \frac{\hat{U}(\theta_0, \hat{p}, t) - U(\theta_0)}{t} &\equiv \theta_0 \hat{p} \lambda U(1) - [\theta_0 \hat{p} \lambda + (1 - \theta_0) \lambda] U(\theta_0) \\
&\quad + (1 - \hat{p}) \lambda \theta_0 (1 - \theta_0) U'(\theta_0) - b(\hat{p})
\end{aligned} \tag{39}$$

Since $U(\theta_0)$ is the value of the optimal monitoring program, $\hat{U}(\theta_0, \hat{p}, t) - U(\theta_0) \leq 0$ for all \hat{p} . Now consider $\hat{p} = p(\theta_0)$. Since $p(\theta)$ is right-continuous $\theta_1 > \theta_0$, and $\theta_1 \rightarrow \theta_0$ as $t \rightarrow 0$, it follows that, for t small enough, $p(\theta_0)$ is arbitrarily close to $p(\theta) \forall \theta \in [\theta_0, \theta_1]$. Hence $\hat{U}(\theta_0, \hat{p}, t) \rightarrow U(\theta_0)$ as $t \rightarrow 0$. Hence $\hat{p} = p(\theta_0)$ maximizes the left-hand side of (39), and by virtue of the equivalence must also maximize the right-hand side. This implies that the derivative of the right hand side with respect to \hat{p} must vanish at $\hat{p} = p(\theta_0)$ if $p(\theta_0)$ is interior, or satisfy appropriate conditions for a corner solution if not.

$$b'(p|_{p=p(\theta_0)}) = \lambda\theta_0[U(1) - \{U(\theta_0) + (1 - \theta_0) U'(\theta_0)\}] ; \quad 0 < p(\theta_0) < 1 \quad (40)$$

$$b'(p|_{p=p(\theta_0)}) \geq \lambda\theta_0[U(1) - \{U(\theta_0) + (1 - \theta_0) U'(\theta_0)\}] ; \quad p(\theta_0) = 0$$

$$b'(p|_{p=p(\theta_0)}) \leq \lambda\theta_0[U(1) - \{U(\theta_0) + (1 - \theta_0) U'(\theta_0)\}] ; \quad p(\theta_0) = 1$$

with the corresponding complementary slackness conditions at the boundaries. These conditions characterize the optimal monitoring function $p(\theta)$. Note that the right-hand side of (40) is independent of \hat{p} , since $U(1)$, $U(\theta_0)$, $U'(\theta_0)$ are all values corresponding to the optimal program. Since $b'(p)$ is strictly positive and increasing in p , the solution of (40) is unique for each θ .

The result we want is:

Theorem A.1. *If $b'(p) > 0 \forall p \in [0, 1]$, then there is θ_a, θ_b with $0 < \theta_a \leq \theta_b < \theta^*$ such that*

(i) $p(\theta) = 0$ in the interval $[0, \theta_a)$

(ii) $p(\theta) = 0$ in the interval $(\theta_b, \theta^*]$.

(iii) If $b'(0)$ is not too large, then $p(\theta) > 0$ for some $\theta \in [\theta_a, \theta_b]$.

Proof of (i): $U(\theta)$ and $U'(\theta)$ are clearly non-negative, hence the right-hand side of (40) is bounded above by $\lambda\theta_0[U(1)]$. But this tends to 0 as $\theta_0 \rightarrow 0$, and must fall below $b'(0)$ for $\theta > 0$ small enough. Hence for small enough θ we must have $p(\theta) = 0$.

Proof of (ii): Suppose there is $\bar{\theta}$ such that (40) holds with strict inequality at $p = 0 \forall \theta$ between $\bar{\theta}$ and θ^* . Then in this range $p(\theta) = 0$ and $U(\theta)$ is identical to $U^*(\theta)$ as defined in (2).

Differentiating (2), noting that $U(1) = \frac{g}{r}$, and performing the necessary manipulations we obtain

$$U(\theta_0) + (1 - \theta_0) U'(\theta_0) = U(1) \left[\frac{\theta}{1 - \theta} \frac{1 - \theta^*}{\theta^*} \right] \quad (41)$$

As $\theta \rightarrow \theta^*$ the rhs of (41) converges to $U(1)$, which implies that the rhs of the second condition in (40) converges to 0. Since $b'(0)$ is strictly positive, it follows that for θ sufficiently close to θ^* , we must have $p(\theta) = 0$.

Proof of (iii): Since the expression in (41) is strictly positive for $\theta < \theta^*$, the RHS of (40) is strictly positive for all $\theta \in (0, \theta^*)$ and therefore greater than $b'(0)$ for $b'(0)$ sufficiently small. ■

Theorem (5.1) in the text is a restatement of parts (i) and (ii) above.

A.5 Proof of theorem 5.2

Proof. of theorem (5.2): If a worker with θ_0 is observed to have failed to handle a crisis correctly, the updated probability becomes

$$\theta(\theta_0, failure) = \frac{\theta_0(1 - \gamma)}{1 - \theta_0\gamma} \quad (42)$$

where θ_0 refers to the firm's belief about θ just prior to failure. If the firm does not observe a failure and has not been monitoring the worker, then it updates according to

$$\theta(\tilde{\theta}_0, t) = \frac{\tilde{\theta}_0(e^{-\lambda t} + (1 - e^{-\lambda t})\gamma)}{\tilde{\theta}_0(1 - e^{-\lambda t})\gamma + e^{-\lambda t}}. \quad (43)$$

A little manipulation establishes that

$$\begin{aligned} \theta(\tilde{\theta}_0, t|failure) < \tilde{\theta}_0 &\Leftrightarrow \\ e^{-\lambda t} &> \frac{1 - \gamma}{2 - \gamma}. \end{aligned} \quad (44)$$

Solving for t gives condition (7). □

A.6 Proof of theorem 5.3

Proof. The existence of the lower range follows directly from the proof in the base case.

The existence of the monitoring range in the H -task is established in the text. Setting $\theta = 1$ in condition (21) proves that the upper end of this range is less than 1.

To prove the existence of upper no-monitoring range in the L -task, assume that no such range exists. Then the transition between tasks occurs when

$$-b - \lambda c_l (1 - \theta_n) = g - b - \lambda c_m (1 - \theta_n) \quad (45)$$

$$\theta_n = \frac{\lambda (c_m - c_l) - g}{\lambda (c_m - c_l)}. \quad (46)$$

Now consider a strategy of assigning a worker to the L -task and then promoting him to the H -task and monitoring him until a crisis arises after which he either separates or is known to be good and is not monitored. Consider the determination of θ^* in this case. Letting $p = 1 - e^{-\lambda t_1}$

$$\begin{aligned} U(\theta_0, 1 - p) &= (\theta_0 + (1 - \theta_0)(1 - p)) e^{-rt_1} \left(\frac{g - b}{\lambda + r} + \frac{\lambda \theta^* g}{r(\lambda + r)} - \frac{(1 - \theta^*) \lambda c_m}{\lambda + r} \right) \\ &\quad - (1 - \theta_0) c_l \int_0^{t_1} \lambda e^{-(\lambda+r)t} dt \quad (47) \end{aligned}$$

$$\begin{aligned} &= (\theta_0 + (1 - \theta_0)(1 - p)) e^{-rt_1} \left(\frac{g - b}{\lambda + r} + \frac{\lambda \theta^* g}{r(\lambda + r)} - \frac{(1 - \theta^*) \lambda c_m}{\lambda + r} \right) \\ &\quad - (1 - \theta_0) c_l \lambda \frac{1 - e^{-(\lambda+r)t_1}}{\lambda + r} \quad (48) \end{aligned}$$

But Bayesian updating implies

$$\theta^* = \frac{\theta_0}{\theta_0 + (1 - p)(1 - \theta_0)}$$

or

$$\theta_0 = \theta^* \frac{1 - p}{1 - \theta^* p}$$

and

$$1 - \theta_0 = \frac{1 - \theta^*}{1 - \theta^* p}$$

So we have

$$\begin{aligned}
U(\theta_0, 1-p) &= (\theta_0 + (1-\theta_0)(1-p)) \cdot \\
& e^{-rt_1} \left(\frac{g-b}{\lambda+r} + \frac{\lambda g}{r(\lambda+r)} \frac{\theta_0}{\theta_0 + (1-p)(1-\theta_0)} - \frac{\lambda c_m}{\lambda+r} \frac{(1-p)(1-\theta_0)}{\theta_0 + (1-p)(1-\theta_0)} \right) \\
& - (1-\theta_0)c_l \lambda \frac{1-e^{-(\lambda+r)t_1}}{\lambda+r} \\
& = (\theta_0 + (1-\theta_0)(1-p)) e^{-rt_1} \left(\frac{g-b}{\lambda+r} + \frac{\lambda g}{r(\lambda+r)} \frac{\theta_0}{\theta_0 + (1-p)(1-\theta_0)} \right) \\
& - \frac{\lambda c_m}{\lambda+r} (1-p)(1-\theta_0) e^{-rt_1} - (1-\theta_0)c_l \lambda \frac{1-e^{-(\lambda+r)t_1}}{\lambda+r}
\end{aligned}$$

Note that $e^{-rt_1} = e^{[-\lambda t_1]^{\frac{r}{\lambda}}} = (1-p)^{\frac{r}{\lambda}}$, which yields

$$= \theta_0 (1-p)^{\frac{r}{\lambda}} \frac{(\lambda+r)g-rb}{r(\lambda+r)} + (1-\theta_0)(1-p)^{\frac{r}{\lambda}+1} \frac{g+(c_l-c_m)\lambda-b}{\lambda+r} - \frac{(1-\theta_0)(c_l)\lambda}{\lambda+r}$$

Maximize wrt to $1-p$

$$\begin{aligned}
\frac{\theta^*}{1-\theta^*p} \frac{(\lambda+r)g-rb}{(\lambda+r)} &= - \left(1 - \theta^* \frac{1-p}{1-\theta^*p} \right) (g + (c_l - c_m)\lambda - b) \\
\theta^* \frac{(\lambda+r)g-rb}{(\lambda+r)} &= -(1-\theta^*) (g + (c_l - c_m)\lambda - b) \\
\theta^* &= \frac{\lambda+r}{\lambda} \frac{b-g-(c_l-c_m)\lambda}{b-(\lambda+r)(c_l-c_m)}.
\end{aligned}$$

It is readily verified that $\theta^* > \theta_n$. Therefore, $U(\theta_n, L, N) > U(\theta_n, H, M) = U(\theta_n, L, M)$ which proves the existence of an upper no-monitoring zone. \square