

Optimal Job Design and Career Dynamics in the Presence of Uncertainty*

Elena Pastorino

University of Pennsylvania

(PRELIMINARY AND INCOMPLETE)

June 2004

Abstract

The paper studies a learning model in which information about a worker's ability can be acquired symmetrically by the worker and a firm in any period by observing the worker's performance on a given task. Productivity at different tasks is assumed to be differentially sensitive to a worker's intrinsic talent: potentially more profitable tasks entail the risk of greater output destruction if the worker assigned to them is not of the ability required. We characterize the (essentially unique) optimal retention and task assignment policy for the class of sequential equilibria of this game, by showing that the equilibria of interest are strategically equivalent to the solution of an experimentation problem (a discounted multi-armed bandit with independent and dependent arms). These equilibria are all ex ante efficient but involve ex post inefficient task allocation and separation. While the ex post inefficiency of separations persists even as the time horizon becomes arbitrarily large, in the limit task assignment is efficient. When ability consists of multiple skills, low performing promoted workers are fired rather than demoted, if outcomes at lower level tasks, compared to those at higher level tasks, provide a sufficiently accurate measure of ability. We then examine the strategic effects of the dynamics of learning on a worker's career profile. We prove, in particular, that price competition among firms causes ex ante inefficient turnover and task assignment, independently of the degree of transferability of human capital. In a class of equilibria of interest it generates a wage dynamics consistent with properties observed in the data.

JEL Classification: D83, J24, J41.

Keywords: Retention, Job Assignment, Learning, Experimentation, Correlated Bandit.

*I am indebted to Steven Matthews, Kenneth Burdett, George Mailath and Jan Eeckhout for their generous advice and encouragement. I sincerely wish to thank Bengt Holmström and Dale Mortensen for their precious suggestions. I am grateful to Hector Chade and the other participants in the Midwest Economic Theory and International Trade Conference at Indiana University (October 2003) for their insightful comments.

1 Introduction

An issue firms face, when allocating employees to jobs, is how workers can be made to perform different tasks in a way that does not create competition in their attention, if a firm values these tasks as complementary. In the principal-agent literature examples of this are the trade-off between quantity and quality in output production or between teaching of basic and higher order thinking skills in education (Holmström and Milgrom [1991]). These models commonly interpret job design as an instrument to *control incentives* in presence of moral hazard. An alternative approach, which will be suggested in the following, is to interpret instead job and career design as a problem of *optimal design of experiment*. When agents' productive characteristics are not perfectly observable, jobs can be valuable as a source of information for the principal. In particular, complementarity of tasks in terms of *information* production can be as relevant as their complementarity in terms of *output* production.

When a firm assigns a newly hired worker to a position, it frequently does not know with certainty his ability and attitude to the various jobs. Then, allocating an employee a given task entails an opportunity cost in terms of the foregone profit at an alternative task, at which the employee might be more productive if he is talented for that job. A young investment banker or consultant, for instance, can always be made to carry out an easy task (i.e., provide support to senior workers on routine duties) at which performance can be easily monitored, but the ability the firm is typically uncertain about is related to the worker's productivity at other tasks, such as elaborating a restructuring project for a client or increasing the firm's managed portfolio, for which monitoring is less accurate and outcomes more variable. Indeed, the accomplishment of these richer tasks, with a greater impact on a firm's profit, commonly requires the participation of other workers (i.e., a consultancy project is a typical team activity) and success is usually influenced by factors outside the worker's immediate control (i.e., the demand for the firm's services or the current phase of business cycle). In this case the firm thus faces a trade-off between observing the worker's performance at a more informative task, less profitable if the worker's ability is high, and assigning him immediately a more profitable task, which might generate a loss for the firm if the worker is not talented for that job.

The interpretation we will provide to the way personnel policies solve this trade-off is that, when jobs affect the firm's information about workers' abilities, the change in a worker's task assignment over his career can be understood as the outcome of a sequential screening process. In particular, as suggested by Holmström and Tirole [1989], in presence of uncertainty about workers' productive characteristics, the firm's allocation problem cannot be reduced to "match[ing] workers with jobs in a myopic fashion based on currently available information, but [it] also [requires] to consider the

implications of current assignments on what might be learned for the benefit of future assignments”. This endogenous process of information acquisition on the part of the firm, as hinted, has to balance the profit loss the firm might incur, if a worker is assigned a task at which he is ill-suited, with the benefit of acquiring new information, if the worker’s performance at that task is sufficiently informative about his ability. This situation, therefore, creates a tension for the firm between the objective of profit-maximization and the need for experimentation, by trying a worker at different jobs in order to learn about his productivity. As a result, retention, task assignment and promotion policies can be naturally characterized as the optimal solution to a *sequential design problem*, in a precise statistical sense, with the firm experimenting on workers’ ability by employing them at different tasks.

Abstracting from incentive issues, the focus of the paper is on how informational concerns can account for certain aspects of job design, in particular for the existence of stable patterns of career advancements in firms (i.e., a hierarchy or ‘job ladder’, with workers moving along it as their tenure at a firm increases). One of the insights from the analysis is that, when talent is correlated across tasks, a firm typically benefits by assigning a worker, early in his career, a task at which he has a static comparative disadvantage, if that task provides accurate information about his productivity. In particular, even if workers of different skills have different patterns of comparative advantage across tasks, the possibility of drawing more accurate inferences about ability can be valuable enough for the firm to justify assigning a worker a task at which he is expected to generate a loss. Since, as time passes, the firm has more and more chances to observe the worker’s performance, eventually the worker will be either permanently retained, and assigned the task for which he is most talented, or fired, if his ability is perceived to be inadequate for the firm’s needs.

An intuition for why this screening procedure is optimal is that, even if tasks are intertemporally *substitutes* in terms of output production, information on performance at different tasks is *complementary* from a dynamic perspective: observing a worker’s performance at one task can be useful in assessing his ability at other tasks. We will show then that the existence of this informational link makes it profitable for the firm to assign tasks sequentially in order of *decreasing informativeness*, with easier-to-monitor tasks assigned first and harder-to-measure task allocated only when uncertainty about ability has been partially resolved.¹

Formally, the model features one firm which, in each period of an infinite horizon, is at most matched with one worker, whose ability can be either ‘high’ or ‘low’. Both the firm and the worker

¹Note that this intuition contrasts with the conclusion from multitask incentive models with risk averse agents. In those, in fact, separating among workers tasks at which measurement errors are correlated is optimal, since it reduces the risk premium incurred in providing incentives. In presence of uncertainty, instead, it is exactly this correlation, and its impact on the process of inference about ability, to make sequential assignment profitable.

do not observe the worker’s ability and are risk neutral; for the firm, though, it is only profitable to employ a high ability worker. Once employed, a worker supplies effort inelastically, which is observed by the firm and determines a stochastic output signal, whose distribution depends on the worker’s actual ability and on the task he is assigned. The output realization in a period, being public, can then be used by the firm and the worker to draw inferences about the worker’s unobserved productivity.

Consider first the case in which only one task is available. Given this informational structure, the firm’s decision problem can be interpreted as a two-armed bandit problem, in which the arm associated with employing the worker in a period delivers a stochastic reward, while the arm represented by the firm’s outside option generates a known profit.² The firm’s optimal strategy in this framework is a reservation-belief strategy, with the worker being employed by the firm as long as its updated belief about the worker’s productivity exceeds a cut-off value, endogenously determined, and dismissed (forever) otherwise. Under this optimal strategy a worker of high ability has (at least) the same probability of being retained at the firm as a low ability worker. Still, as the time horizon becomes arbitrarily large, a high ability worker is retained forever with strictly positive probability, while a low ability worker is fired almost surely. Then, the firm typically risks firing a high ability worker even if it is potentially allowed to sample an infinite number of observations. This incompleteness in the firm’s learning about a worker’s ability, characteristic of experimentation models, is due to the fact that, even in the long run, with positive probability the firm can observe a sequence of mediocre performances sufficiently long to convince it that the worker’s ability is low, even if he is truly of high productivity.

To analyze the impact of learning on job assignment and career design, we then explore the case in which, once employed, a worker can be assigned one of two different tasks, either task y or task z .³ Analogously to the above, in all the sequential equilibria of the task-assignment game the firm’s best response can be characterized as the solution to a three-armed bandit problem, with two dependent arms (task y and task z), being ability correlated across them, and one independent arm, corresponding to the firm’s outside option. Suppose that a low ability worker has a comparative advantage at task z while a high ability worker at task y , but that the firm prefers to collect its outside option rather than to employ a low ability worker. Then, if task z is more informative about a worker’s productivity than task y (i.e., the distribution of the firm’s posterior that the worker’s ability is high under task z is a mean-preserving spread of the corresponding distribution under task y), the firm’s optimal employment strategy is a belief-interval strategy, under which a worker

²It can be interpreted, for instance, as the value of re-sampling a worker.

³The two task case is explored in the paper for expositional simplicity. The extension to a finite number of tasks preserves all the results of interest.

is assigned task z as long as the firm’s posterior lies in an intermediate belief range. If the firm’s posterior is sufficiently high, the worker is thus allocated task y , while a worker whose assessed ability is relatively low is permanently fired. In particular, the option of experimenting across multiple tasks has the benefit of reducing the probability that a high ability worker is eventually fired. More specifically, by assigning a newly hired worker task z when most uncertain about his ability, the firm can reduce the limiting likelihood of a type II error, i.e., dismissing forever a high ability worker, without increasing the likelihood of a type I error, i.e., retaining forever a low ability worker.

We then investigate the case in which ability consists of two dimensions of skill. Suppose, for instance, that good performance at either task requires a specific and *independent* dimension of skill, but, say, task z is equally informative about ability at z and y , while task y is less informative than z about the talent required to perform satisfactorily at y . In this context it can be shown that the information collected by observing the worker’s performance at task z can help reduce the firm’s uncertainty about a worker’s *overall* skill and, therefore, make its use profitable, even if a newly hired worker is always more profitable at z than at y . In this scenario an endogenous grouping of tasks into job levels emerges. As a consequence, in contrast with the one-dimensional ability case, if a worker performs unsatisfactorily at task y , the firm is better off by firing him than assigning him back the more informative task. The intuition for this result is that a worker whose performance is revealed unsuccessful at a higher level task, at which production outcomes are not very informative, must be even less likely to be profitable at a lower level task, if performance measures at the two tasks are correlated. Notice that this finding can help explaining the empirical puzzle that promoted workers are more likely to be fired than demoted, as documented, for instance, by Baker, Gibbs and Holmström [1994a].

The above results overall offer an interpretation of career advancements alternative to the Peter Principle: workers are not necessarily promoted to their level of incompetence, as the principle states, but rather gradually assigned to the positions at which they can best contribute to a firm’s value. Namely, they are promoted to their level of *perceived* competence. In particular, a good record at a *current* job efficiently supports promotion to a *different* job, even when a firm does not need to provide workers with incentives for performance. This feature is in stark contrast with the result familiar from bandit problems (i.e., each arm is informative only about the type of the arm being played), according to which a worker should be continuously assigned to the same job as long as his performance at it is good, when information about ability at different tasks is generated independently.⁴

⁴It is the well-known result of the *stay-on-the-winner* characteristic of the optimal policy in independent bandit problems.

We finally show that the same characterization results hold in presence of price competition in the external labor market. By an argument similar to the one used by Bergemann and Välimäki [1996], it is in fact possible to show that the efficient experimentation solution can be characterized as the solution to the employment problem of a single firm, in which the worker's outside option is normalized to zero. In this case turnover is typically inefficient and the inefficiency is independent of the degree of task- or firm-specificity of a worker's human capital.⁵

The paper is organized as follows. Section 2 introduces the baseline multitask model and presents the main results. Section 3 explores the case in which ability consists of multiple skills, while Section 4 addresses the issue of outside labor market competition. Section 5 illustrates the relevant related literature. Finally, Section 6 briefly concludes and discusses possible extensions of the model, as well as directions of future research.

2 A Two Task Model

Consider a firm that employs one worker. Time is discrete and has an infinite horizon, with dates $t = 1, 2, \dots$. The worker's productivity θ is unobserved to both the firm and the worker, with support $\{\underline{\theta}, \bar{\theta}\}$, $\bar{\theta} > \underline{\theta}$. The firm and the worker's prior distribution over the worker's productivity is $\Pr(\underline{\theta}) = 1 - \pi$ and $\Pr(\bar{\theta}) = \phi_0$, $\phi_0 \in (0, 1)$. In every period of employment the worker inelastically supplies one unit of (observable and verifiable) effort, which determines output according to the task the worker is assigned.⁶ In particular, if the worker performs task y , his output in period t is $y_t = \{y, \bar{y}\}$, $\bar{y} > y > 0$, distributed according to the conditional density $\Pr(\tilde{y}_t = \bar{y} \mid \theta = \bar{\theta}) = \alpha_1$ and $\Pr(\tilde{y}_t = y \mid \theta = \underline{\theta}) = \beta_1$, with $\alpha_1, \beta_1 \in (0, 1)$ and $\alpha_1 > \beta_1$. If, instead, the worker is assigned task z , his output is $z_t = \{z, \bar{z}\}$, $\bar{z} > z > 0$, with conditional density $\Pr(\tilde{z}_t = \bar{z} \mid \theta = \bar{\theta}) = \alpha_2$ and $\Pr(\tilde{z}_t = z \mid \theta = \underline{\theta}) = \beta_2$, where $\alpha_2, \beta_2 \in (0, 1)$ and $\alpha_2 > \beta_2$. It is assumed that y_t and z_t are observable to the firm and the worker.

If employed at date t the worker receives from the firm a payment w_t . Correspondingly, the firm's realized payoff at the end of period t is $y_t - w_t$ or $z_t - w_t$, depending on the task performed by the worker, while the worker's is w_t . Notice that the worker's disutility of effort is normalized to zero. The firm's reservation profit is given by $\bar{\Pi} \geq 0$, while the worker's reservation utility by

⁵An empirical estimation of job and wage dynamics is carried out in a companion paper using the same firm-level dataset first used in Baker, Gibbs and Holmström [1994a, 1994b].

⁶Note that the verifiability of output only matters for the case in which ability is private information to the worker. It is possible to show that, under limited liability, if output is nonverifiable and the distance between the two type-dependent outside options not too small, nonverifiability prevents the use of revelation schemes and the same characterization results as for the symmetric learning case hold.

$\bar{U} \geq 0$. Let

$$(A1) : y(\bar{\theta}) > \bar{\Pi} + \bar{U}$$

$$(A2) : y(\bar{\theta}) > z(\bar{\theta}), z(\underline{\theta}) > y(\underline{\theta})$$

where $y(\theta) \equiv E(y | \theta)$ and $z(\theta) \equiv E(z | \theta)$ denote, respectively, the one period expected revenue to the firm from assigning the worker task z or task y . Let also $y(\phi) \equiv \phi y(\bar{\theta}) + (1 - \phi)y(\underline{\theta})$ and $z(\phi) \equiv \phi z(\bar{\theta}) + (1 - \phi)z(\underline{\theta})$. Notice that (A2) is satisfied whenever the distribution of output under task z is generated through a mean-preserving decrease in the spread of the probability mass of the cumulative density function of y , for each θ , with $\underline{y} < \underline{z}$ and $\bar{y} > \bar{z}$. The assumption is meant to capture the feature that ability is more valuable at activities which can in principle contribute more to the firm's value or, equivalently, that the incremental impact of ability is greater at potentially more profitable tasks. The restriction also implies that task y entails the risk of greater output destruction if the worker assigned to it is not of the ability required. Both the firm and the worker discount future payoffs according to a common discount factor $\delta \in [0, 1)$.

Timing is as follows. At the beginning of any period $t \geq 1$ the firm decides whether to propose employment to the worker at wage w_t . Wage offers are made by the firm on a take-it-or-leave-it basis. If in a period the firm does not make a proposal or the worker rejects the firm's offer, both the firm and the worker receive their reservation payoffs. They then meet again the following period. If the worker accepts the firm's offer, he is assigned a task and effort at this task stochastically determines output y_t or z_t . Finally, the wage w_t is paid.

Let $H = \bigcup_{t \geq 1} H_t$ represent the set of all the possible histories and H_t the set of the period- t histories (up to but not including period t). An element h_t of H_t contains all the past wage offers, w_τ , $1 \leq \tau < t$, the worker's acceptance decisions, d_τ , $1 \leq \tau < t$, his task assignments, $j_\tau \in \{y, z\}$, $1 \leq \tau < t$, and the random realizations of output, $k_\tau \in \{y_\tau, z_\tau\}$, $1 \leq \tau < t$. Hence, an element h_t of H_t , $t \geq 2$, is given by

$$h_t = (w_1, d_1, j_1, k_1, \dots, w_{t-1}, d_{t-1}, j_{t-1}, k_{t-1})$$

with $H_1 = \{\emptyset\}$ and, for $1 \leq \tau < t$,

$$k_\tau(j_\tau) = \begin{cases} k_\tau(y) \in \{\underline{y}, \bar{y}\}, & \text{if } j_\tau = y \\ k_\tau(z) \in \{\underline{z}, \bar{z}\}, & \text{if } j_\tau = z. \end{cases}$$

A pure strategy for the firm in period t is given by the function ω_t , which determines the compensation offered to the worker, and the function e_t , which defines the worker's task assignment in the period. The function ω_t maps the history into the positive real numbers or the empty set (the firm

has always the possibility of not offering employment)

$$\omega_t : H_t \rightarrow \mathbb{R} \cup \{\emptyset\}$$

with $w_t \in \mathbb{R} \cup \{\emptyset\}$. The function e_t , describing the firm's task assignment decision, maps the history and the firm's wage offer into the task space $\{y, z\}$,

$$e_t : H_t \times \mathbb{R} \rightarrow \{y, z\}$$

with $j_t \in \{y, z\}$. A pure strategy for the worker in period t is a function from the history and the firm's current wage offer into the decision space $\{A, R\}$, with A denoting the worker's acceptance of the firm's offer while R his refusal,

$$a_t : H_t \times \mathbb{R} \rightarrow \{A, R\}$$

where $d_t \in \{A, R\}$.

Given the existence of symmetric uncertainty about the worker's ability, let the firm's and the worker's posterior belief at the beginning of date $t \geq 2$ that the worker's ability is $\bar{\theta}$ be denoted by ϕ_t , with $\phi_1 \equiv \pi$. In order to focus the equilibrium analysis on the strategic effects of the firm's learning, we restrict attention to Markov Perfect equilibria (MPE's) of the game for which ϕ_t is the state variable. In this framework a strategy is Markovian if it depends on the past only through the payoff relevant history as summarized by ϕ_t . Stationary MPE strategies are MPE strategies which are time-invariant. In any MPE the firm maximizes profit, the worker maximizes utility, the firm and the worker's expectations about the worker's ability are correct and they use Bayes' rule to update their posteriors about the worker's productivity. Let ϕ (omitting the subscript t) be then the state variable of the corresponding complete information game (CONTINUE).

Apart from issues of sequential rationality, non-Markovian equilibria, typically in trigger strategies, involve a degree of coordination between the firm's and the worker's behavior which ignores the screening value of early periods of employment. On the contrary, since different job positions require specific qualifications and skills which are difficult to assess in newly hired workers, the value of the match to both parties is typically identified only after the worker has been employed for some time.⁷ In the following, therefore, we rule out these equilibria by requiring the parties' behavior to be Markovian. In our symmetric learning framework, though, the restriction to Markovian strategies is without loss of generality. As it will be shown, given the worker's acceptance behavior, the firm's *uniquely* optimal employment strategies (modulo the way indifference is solved) are Markovian, i.e., all sequential equilibria are stationary MPE's as defined.

⁷Indeed, evaluating ability in young employees appears to be a major force shaping firms' personnel policies. See, for a reference, Milgrom and Roberts [1992].

Given the structure of the game between the firm and the worker, in equilibrium the firm's wage offer must exactly match the worker's reservation utility. Intuitively, any wage higher than \bar{U} would only decrease the firm's profit, since it does not affect the probability of realization of a high output: by paying \bar{U} , the firm can always ensure the worker's participation. On the other hand, a wage payment strictly lower than \bar{U} would discourage a worker from accepting the firm's offer, given that the firm cannot credibly promise to pay more than \bar{U} in a period. In particular, accepting any offer at least equal to \bar{U} is for a worker a dominant strategy. This observation motivates the following:

Lemma 1. *In any MPE the firm's cost-minimizing wage policy consists in paying the worker the wage \bar{U} for any period of employment.*

In order to investigate the informational value of experimenting a worker at different tasks, as a benchmark we first consider the case in which the firm can only allocate the worker task y .

2.1 The One Task Case

If only task y is available, assumptions (A1) and (A2) read as

$$(A1') : y(\bar{\theta}) > \bar{\Pi} + \bar{U} > y(\underline{\theta}).$$

For fixed wage policy (and worker's acceptance behavior), the firm's problem in any period reduces to deciding whether or not to employ the worker at wage \bar{U} . Since the worker's productive ability is unknown, an employment strategy that yields the highest total expected profit can be characterized as the solution to an independent *one-armed bandit problem*. The term bandit derives from modelling this class of sequential decision problems as an n -armed bandit, $n \geq 1$, or a slot machine with n arms, each generating rewards according to one of a finite number of distributions (the types of the arm). Independence refers to the fact that each arm is only informative about its unknown type. Specifically, at each point in time the firm's decision is between the stochastic arm (i.e., employing the worker), which generates rewards according to a Bernoulli distribution with probability of success α or β , or playing the deterministic arm (i.e., not employing the worker), which delivers the constant reservation profit $\bar{\Pi}$.

Observe first that, given the informational structure of the problem and the equilibrium restriction, an optimal strategy for the firm (i.e., a wage proposal strategy which is a best response to the workers's acceptance behavior) must be one in which, whenever the firm does not employ the worker in a period, it will never employ him in any subsequent period. Equivalently, in our environment separations can only be permanent.⁸ The argument goes as follows. If the firm employs

⁸This feature seems consistent with the observed finding that in human capital intensive sectors like finance,

the worker for $t - 1$ periods, at the beginning of period t it has to decide whether to employ him for at least one additional period. If it chooses not to, the firm obtains a flow payoff of $\bar{\Pi}$, but it does not receive any additional information about the worker's ability. Therefore, if the firm's belief at date t is such that not employing the worker is optimal, the same choice must be optimal at $t + 1$, given that the belief is unchanged. The firm's hiring problem in a generic period t can then be described by the value function

$$V(\phi) = \max\{\bar{\Pi}, V_y(\phi)\} = \max\{\bar{\Pi}, (1 - \delta)[y(\phi) - \bar{U}] + \delta E_y V(\phi)\}. \quad (1)$$

Payoffs are normalized so as to be expressed as per period averages. Properties of V^f are summarized in the following Lemma.

Lemma 2. *Under assumption (A1'):*

- (i) $V(\phi)$ is well-defined, continuous, increasing and convex in ϕ .
- (ii) $V(0) = \bar{\Pi}$ and $V(1) = y(\bar{\theta}) - \bar{U} = \alpha_1 \bar{y} + (1 - \alpha_1) \underline{y} - \bar{U}$.

Proof: See the Appendix. □

Since $y(\phi)$ is strictly increasing in ϕ and $V(\phi)$ increasing, the right-hand side of (1) is strictly increasing in ϕ . From this, together with (A1') and (A2'), it follows that there exists a unique value $\tilde{\phi}_y \in (0, 1)$ satisfying $\bar{\Pi} = V_y(\phi)$. This implies that the firm's optimal strategy consists in employing the worker in a period, if the resulting present expected discounted profit exceeds the value of the firm's outside option or, equivalently, if $\phi \geq \tilde{\phi}_y$. A part for the specification of the wage offered when $\phi < \tilde{\phi}_y$ (any offer $w(\phi) < \bar{U}$ is payoff-equivalent given that it induces refusal on the part of the worker), this strategy is essentially the firm's uniquely optimal one, for given prior belief π . More formally:

Proposition 1. *Let (A1') hold. Then:*

- (i) *A strategy profile (ω^*, a^*) is an MPE if and only if*

$$\omega^*(\phi) = \begin{cases} \bar{U}, & \text{if } \phi \geq \tilde{\phi}_y \\ w', & \text{otherwise} \end{cases} \quad \text{and} \quad a^*(\phi, w) = \begin{cases} A, & w \geq \bar{U} \\ R, & \text{otherwise} \end{cases}$$

where $w' \in [0, \bar{U})$, together with beliefs determined according to Bayes' rule.

- (ii) *All MPE's are payoff-equivalent to both parties and constrained Pareto efficient.*
- (iii) *Every MPE is outcome equivalent to some sequential equilibrium.*

insurance and real estate, temporary layoffs are relatively rare, as compared, for instance, with manufacturing. For a reference, see Anderson and Meyer [1994]. In our framework, though, separations are always initiated by the firm.

Proof: (i) By Lemma 1, given the firm's proposal strategy and the common belief about the worker's ability, the worker maximizes utility by accepting any (and only) a wage offer at least equal to \bar{U} for any ϕ . For given ϕ , by the above argument the firm maximizes its profit by offering the worker the wage \bar{U} when $\phi \geq \tilde{\phi}_y$ and a wage strictly less than \bar{U} otherwise. Given the output realization in a period, the firm and the worker update their beliefs about the worker's ability using Bayes' rule. (ii) Notice first that any other equilibrium strategy profile in which the firm's offer is less than \bar{U} , when $\phi < \tilde{\phi}_y$, achieves the same payoff, given the worker's acceptance behavior. Consider now the case of non-stationary strategies. Observe that, by the argument in Lemma 1, there cannot be any equilibrium in which the firm's wage offer is greater than \bar{U} for some ϕ . The only admissible MPE's in non-stationary strategies are therefore those in which, whenever $\phi < \tilde{\phi}_y$, the firm chooses to offer a wage less than \bar{U} , possibly different from period to period. Again, all these MPE are payoff-equivalent. Given the worker's acceptance behavior and, correspondingly, the firm's equilibrium wage strategy, profit maximization is equivalent to surplus maximization. Finally, (iii) by the equivalence between the original non-Markovian problem and the Markovian stochastic dynamic programming problem (with state space given by the set of probability distribution over Θ), it is possible to restrict attention to the the solution of the problem in recursive form. \square

Intuitively, for the firm to be willing to employ the worker in a period, the corresponding one-period profit must be at least equal to $\bar{\Pi}$, if the firm ignores the benefit of the additional information about the worker's ability generated by his performance. This implies that the value of the belief for which the firm is indifferent between employing and not employing the worker must be higher when the firm acts myopically (for $\delta = 0$) than when it internalizes the informational gain associated with observing the worker's output in a period. As a result, the firm's optimal employment strategy is an *experimentation strategy*, i.e., the firm (efficiently, as shown) trade-offs in expected terms the current period profit from employing the worker against the benefit of improving on its assessment about the worker's productivity. In particular, it is optimal for the firm to employ the worker even when he is expected to generate a one-period loss. Formally, define $\phi_{m,y}$ to be the (unique) value of the posterior belief for which a myopic firm is indifferent between offering and not offering the worker employment. Let $\tilde{V}^f(\phi)$ be the firm's payoff when it behaves myopically. The following Proposition can then be proved:

Proposition 2. *Let (A1') hold and $\phi_{m,y}$ satisfy $\tilde{V}(\phi_{m,y}) = \bar{\Pi}$. Then, $\phi_{m,y} > \tilde{\phi}_y$.*

Proof: See the Appendix. \square

By solving period by period an informationally constrained Pareto problem, the firm induces employment only when ex ante efficient. A consequence of separation being ex ante efficient is

that the firm can not benefit by contracting on the worker’s participation for more than one period (*short-term contracting*). By committing to employ the worker for a fixed number of periods at wage \bar{U} , the firm can at most replicate the outcome of the optimal experimentation strategy, since, by the principle of optimality, it is the optimal strategy among the class of all (Markovian and non-Markovian) commitment and non-commitment strategies.⁹ A fortiori, no long-term contract can improve on a random sequence of spot contracts when only participation is contractible. Therefore:

Corollary 1. *Under assumptions (A1’), any optimal long-term contract is equivalent to the random sequence of spot contracts associated with the corresponding MPE.*

Notice, though, that MPE’s are typically ex post inefficient. The above characterization results imply that, after observing a history along which few high output signals have realized, it is optimal for the firm to fire the worker. In particular, even when the worker is of high productivity, the firm may employ him only a finite number of periods, if it becomes sufficiently pessimistic about his talent. Indeed, as it will be shown, histories that induce the firm to believe it is employing a low ability worker, despite the worker’s productivity being high, have strictly positive probability even when the time horizon becomes arbitrarily large.

2.1.1 Tenure and the Dynamics of Learning

To ensure that it is always profitable for the firm to hire the worker in the first period, given the common belief π that his ability is $\bar{\theta}$, assume

$$(A2') : \pi > \tilde{\phi}_y.$$

The sequence of output realizations observed in equilibrium can be interpreted as the outcome of consecutive Bernoulli trials, with probability of success equal to α , if the worker’s ability is high, or β , if the worker’s ability is low. Accordingly, the cardinality of a history h_t , $c(h_t)$, denotes the number of realizations of high output along it and can be expressed as $c(h_t) = \sum_{\tau=1}^{t-1} 1_\tau$, where $1_\tau \in \{0, 1\}$ is the indicator function of a success (high output) in period τ . Since the signal y_t is exchangeable, the belief ϕ is then just a function of the total number of high output signals realized up to time $t - 1$, regardless of their order. As such, it is increasing in $c(h_t)$.

To analyze the impact of learning on the worker’s tenure, define $\rho_t(\theta)$ to be the probability that a worker of type θ is continuously employed by the firm at least until period t .¹⁰ Notice that the

⁹See, for instance, Hinderer [1970] for a formal argument.

¹⁰The discussion in the present Subsection builds on an analogous argument in Araujo and Camargo [2002]. I am indebted to Braz Camargo for pointing me out the similarity.

firm's retention rule $\phi \geq \tilde{\phi}_y$ can be expressed equivalently in terms of the cardinality of the period- t histories associated with ϕ , as $c(h_t) \geq \lfloor \lambda(t-1) - \gamma \rfloor$, where¹¹

$$\lambda = \frac{\ln\left(\frac{1-\alpha_2}{1-\alpha_1}\right)}{\ln\left(\frac{\alpha_1(1-\alpha_2)}{\alpha_2(1-\alpha_1)}\right)} \in (0, 1) \quad \text{and} \quad \gamma = \frac{\ln\left(\frac{\pi(1-\tilde{\phi}_y)}{\tilde{\phi}_y(1-\pi)}\right)}{\ln\left(\frac{\alpha_1(1-\alpha_2)}{\alpha_2(1-\alpha_1)}\right)} > 0.$$

The firm, then, continues to employ the worker at least until period t as long as a history h_τ with cardinality $c_{\tau-1} \equiv c(h_\tau) \geq \lfloor \lambda(\tau-1) - \gamma \rfloor$ realizes, where $2 \leq \tau \leq t$. Since, by assumption (A3'), a worker is always employed in period 1, $\rho_1(\theta) = 1$. The probability that a worker of type θ will still be employed at the beginning of period 2 is given by the probability that a history with cardinality of at least $\lfloor \lambda - \gamma \rfloor$ realizes, or

$$\rho_2(\theta) = \sum_{c_1 \geq \lfloor \lambda - \gamma \rfloor} p(\theta)^{c_1} (1 - p(\theta))^{1-c_1}$$

with $p(\bar{\theta}) = \alpha_1$ and $p(\underline{\theta}) = \alpha_2$. Proceeding similarly, the probability that a worker of type θ will be continuously employed at least until period t can be computed as

$$\rho_t(\theta) = \sum_{(c_1, \dots, c_{t-1}) \in C_{t-1}} p(\theta)^{c_1 + \dots + c_{t-1}} (1 - p(\theta))^{t-1-c_1-\dots-c_{t-1}}$$

where

$$C_{t-1} = \{(c_1, \dots, c_{t-1}) : c_\tau \geq \lfloor \lambda\tau - \gamma \rfloor - c_1 - \dots - c_{\tau-1}, \text{ for } \tau = 1, \dots, t-1\}.$$

Proposition 3 summarizes the properties of a worker's tenure prospect at the firm.

Proposition 3. *Let (A1')-(A2') hold. Then, for all t , $\rho_t(\bar{\theta}) - \rho_{t+1}(\bar{\theta}) \leq \rho_t(\underline{\theta}) - \rho_{t+1}(\underline{\theta})$. In particular, $\rho_t(\underline{\theta}) \leq \rho_t(\bar{\theta})$. Moreover, for all γ there exists a \bar{t} such that, for all $t > \bar{t}$, these inequalities are strict.*

Proof: See the Appendix. □

Even if the probability of retention of any type of worker is decreasing over time, its decrease is higher for a low type worker: the probability that a worker is fired exactly after t periods of employment, $\rho_t(\theta) - \rho_{t+1}(\theta)$, is higher for a low than for a high ability worker. This follows from the fact that a worker is dismissed only after the firm has observed a sufficiently long sequence of low output realizations, leading it to believe that the worker is of low productivity. This evidence, in turn, is more likely to occur when the worker's actual ability is low rather than high. Therefore,

¹¹The symbol $\lfloor \cdot \rfloor$ denotes the greatest integer lower than the indicated number.

at any point in time the prospect of retention is more favorable for a worker of type $\bar{\theta}$ than for a worker of type $\underline{\theta}$.

The bandit structure of the firm's problem has further implications for a worker's tenure in the long run. It is immediate to see that the Bernoulli distribution governing output realizations satisfies the monotone likelihood ratio property (MLRP) for $\alpha_1 > \alpha_2$. Suppose to order the types of an arm in terms of their expected one-period return. In independent bandit problems the best type of an arm, which becomes optimal at some point in time, will survive forever and continuously with strictly positive probability (see Banks and Sundaram [1992a]).¹² In our model this amounts to the fact that the ex ante probability of permanent tenure of a high ability worker (i.e., the 'best type' of the stochastic arm) is strictly positive. In this framework, in addition, the firm will always be able to screen out a low ability worker in the limit.

The characteristics of the firm's experimentation process have therefore different implications for the average duration of employment of workers of different productivity. In particular, while a worker of low productivity is expected to be employed only a finite number of periods, a high ability worker experiences on average permanent tenure at the firm. More formally, define $p_{t+1;t}(\theta)$ to be the probability that a worker of type θ , who has been continuously employed by the firm for at least t periods, will be employed at least for an additional period. Let $T(\theta)$ indicate the random length of tenure at the firm of a worker of productivity θ and $E_1[T(\theta)]$ the expected value of $T(\theta)$ at the beginning of period 1. Then:

Proposition 4. *Under assumptions (A1')-(A2') the following results hold:*

- (i) $\lim_{t \rightarrow \infty} \rho_t(\underline{\theta}) = 0$, while $0 < \lim_{t \rightarrow \infty} \rho_t(\bar{\theta}) = \bar{\rho} < 1$.
- (ii) $\lim_{t \rightarrow \infty} p_{t+1;t}(\bar{\theta}) = 1$.
- (iii) $E_1[T(\underline{\theta})] < \infty$ and $E_1[T(\bar{\theta})] = \infty$.

Proof: See the Appendix. □

As illustrated by Aghion, Bolton, Harris and Jullien [1991], as long as the firm does not know all the relevant aspects of its objective function, when deciding which action to choose it has to balance two conflicting objectives, the maximization of the informational content of the choice of the current action, on one hand, and the maximization of its expected current period profit, on the other. Since sufficient information about the unknown parameters of interest is needed to select the appropriate arm, the optimal experimentation strategy can be interpreted as adjusting the myopic strategy (i.e., the one which maximizes the one-period expected reward) by allowing for

¹²Notice that, in the absence of MLRP, the problem of identifying in the limit an arm of the best type is not trivial, since the best type of an arm may last only finitely long with probability one (Example 5.1 in Banks and Sundaram [1992a]).

active experimentation as long as it is profitable (i.e., for $\phi \geq \tilde{\phi}_y$). The fact that this trade-off does not disappear in the long run is the source of the incompleteness in the firm's limiting learning: for a sufficiently long sequence of low output realizations, even a high ability worker risks being permanently fired. This is also the rationale behind separations being ex post inefficient even in the long run: the firm does not necessarily know the worker's true level of productivity when it is optimal for it to fire him, even if, were the firm to sample an infinite number of output observations, its estimate about θ would eventually converge to its true value.

2.2 The Two Task Case

Consider now the case in which both tasks y and z are available. By the same argument as in Lemma 1 the cost-minimizing wage policy for the firm is to pay the worker the wage \bar{U} in each period. Since, similarly to the above, under the Markovian restriction on equilibrium behavior not employing the worker is an absorbing state, the firm's value function can be expressed as

$$V(\phi) = \max\{\bar{\Pi}, V_z(\phi), V_y(\phi)\} \\ \max\{\bar{\Pi}, (1 - \delta)[z(\phi) - \bar{U}] + \delta E_z V(\phi), (1 - \delta)[y(\phi) - \bar{U}] + \delta E_y V(\phi)\}. \quad (2)$$

For fixed firm's and worker's strategy, E denotes the (conditional) expectation over the period- $t+1$ value of the firm's posterior, given the period- t value of ϕ and the worker's current task assignment. When the firm does not employ the worker, its (normalized) present expected discounted profit is $\bar{\Pi}$. When, instead, employment takes place, its flow (normalized) profit is $(1 - \delta)[z(\phi) - \bar{U}]$ or $(1 - \delta)[y(\phi) - \bar{U}]$, depending on the task the worker performs. The firm then faces the same decision problem in the next period, with posterior belief determined according to Bayes' rule.

Lemma 3. *Under assumptions (A1)-(A2):*

- (i) $V(\phi)$ is well-defined and continuous, increasing and convex in ϕ .
- (ii) $V(0) = \bar{\Pi}$ and $V(1) = y(\bar{\theta}) - \bar{U} = \alpha_1 \bar{y} + (1 - \alpha_1) \underline{y} - \bar{U}$.

Proof: See the Appendix. □

Note that $y(\phi)$ and $z(\phi)$ are strictly increasing in ϕ . Since $V(\phi)$ is increasing in ϕ , it follows that both the (normalized) present expected discounted values from employing the worker at task z and y are strictly increasing in ϕ . Moreover, from the fact that $(1 - \delta)[z(\bar{\theta}) - \bar{U}] + \delta[y(\bar{\theta}) - \bar{U}] > \bar{\Pi}$ and $\bar{\Pi} > (1 - \delta)[z(\underline{\theta}) - \bar{U}] + \delta\bar{\Pi}$, it follows that there exists a unique value $\phi_z^* \in (0, 1)$ satisfying $\bar{\Pi} = V_z(\phi)$. As it will be shown in Proposition 5, ϕ_z^* equals $\tilde{\phi}_z$, where $\tilde{\phi}_z$ is defined as the threshold belief value for which the firm is indifferent between employing and not employing the worker when

only task z is available (and \tilde{V}_z is the corresponding value function).¹³ Denote this common value of ϕ_z^* and $\tilde{\phi}_z$ by ϕ^* . We will prove in the following that there also exists an interior threshold belief value ϕ^{**} , with $\phi^{**} > \phi^*$, defined implicitly by

$$(1 - \delta)[z(\phi^{**}) - \bar{U}] + \delta E_z V(\phi^{**}) = (1 - \delta)[y(\phi^{**}) - \bar{U}] + \delta E_y V(\phi^{**}) \quad (3)$$

such that it is profitable for the firm to hire the worker and assign him task z if $\phi \in [\phi^*, \phi^{**})$, but to allocate him task y whenever $\phi \in [\phi^{**}, 1]$. If, instead, the worker's assessed ability is less than ϕ^* , the firm is better off by not employing him. Formally, let $\phi_{z,m}$ be the value of the posterior for which a myopic firm is indifferent between employing the worker at task z and firing him, i.e., $z(\phi_{m,z}) - \bar{U} = \bar{\Pi}$ or $\phi_{m,z} = (\bar{\Pi} + \bar{U} - z(\underline{\theta})) / (z(\bar{\theta}) - z(\underline{\theta}))$. Recall that $\phi_{m,y}$ has been similarly defined so as to satisfy $y(\phi_{m,y}) - \bar{U} = \bar{\Pi}$, i.e., $\phi_{m,y} = (\bar{\Pi} + \bar{U} - y(\underline{\theta})) / (y(\bar{\theta}) - y(\underline{\theta}))$. It then follows:

Proposition 5. *Let (A1)-(A2) hold. Suppose that $\alpha_2\beta_1 \geq \alpha_1\beta_2$ and $(\alpha_2 - \beta_2) - (\alpha_1 - \beta_1) > \alpha_2\beta_1 - \alpha_1\beta_2$. Then, if $z(\bar{\theta}) > \bar{\Pi}$ (or, if $\bar{\Pi} > z(\bar{\theta})$, there is $\underline{\delta} \in (0, 1)$ such that for all $\delta > \underline{\delta}$) there exist $\phi^* \in (0, 1)$ and $\phi^{**} \in (0, 1)$, with $\phi^{**} > \phi^*$, so that the firm's optimal employment strategy in any MPE consists in assigning the worker task y , if $\phi \in [\phi^{**}, 1]$, task z , if $\phi \in [\phi^*, \phi^{**})$, and not employing him if $\phi \in [0, \phi^*)$. Under this strategy, $\phi^* < \tilde{\phi}_y < \phi^{**}$, i.e., the firm is more willing to experiment on the worker's ability than when only task y is available, and $\phi^{**} > \phi_m^*$, i.e., the firm experiments longer at z than in the static case.*

Proof: See the Appendix. □

The restrictions on the output distribution needed for the Proposition to hold have an immediate interpretation in terms of increasing riskiness of the distribution of the updated posterior, respectively $\varphi(\phi | z)$ at task z and $\varphi(\phi | y)$ task y . These conditions require that the distribution of the future values of the posterior under task z , conditional on its current period value, be a mean-preserving spread of the corresponding distribution under y . Intuitively, since the firm's value function is convex in ϕ , a riskier distribution of the update of ϕ is always preferred by the firm to a less risky one, given that it entails a faster speed of learning, if measured by the spread in the distribution of the one-step ahead posterior. Notice that the restriction $\alpha_2\beta_1 \geq \alpha_1\beta_2$ and $(\alpha_2 - \beta_2) - (\alpha_1 - \beta_1) > \alpha_2\beta_1 - \alpha_1\beta_2$ ensure that $\phi_z^h \geq \phi_y^h$ and $\phi_z^l < \phi_y^l$, i.e., that the support of the $t + 1$ -period values of ϕ is consistent with the desired second-order stochastic dominance result. In the special case in which output signals are symmetrically distributed, i.e., $\beta_i = 1 - \alpha_i$, $i \in \{1, 2\}$, these conditions simplify to $\alpha_2 \geq \alpha_1$. It is then sufficient that the distribution of output at task z first-order stochastically dominates the one at y .

¹³By the same argument as in the proof of Lemma 2, it is possible to show that this value function is well-defined and continuous, increasing and convex in ϕ .

On the other hand, the restriction $\phi_{m,z} < \phi_{m,y}$, or, equivalently, $y(\phi_{m,z}) - \bar{U} < \bar{\Pi}$, ensures that when task z starts being profitable it strictly dominates task y (in static terms). Under this condition there exists a range of values of ϕ for which the firm is strictly better off by assigning the worker task z rather than task y both in the static and in the dynamic case. In particular, if the signals were equally informative about ability, i.e., $\alpha = \mu$ and $\beta = \nu$, the condition $\phi_{m,z} < \phi_{m,y}$ would be sufficient to guarantee that z is strictly preferred over y for intermediate values of ϕ . Notice that $\phi_{m,z} < \phi_{m,y}$ can also be rewritten as

$$[\bar{\Pi} + \bar{U}]\{y(\bar{\theta}) - y(\underline{\theta}) - z(\bar{\theta}) + z(\underline{\theta})\} < y(\bar{\theta})z(\underline{\theta}) - y(\underline{\theta})z(\bar{\theta}).$$

By assumptions (A1)-(A2), the left-hand side of the above is strictly positive. This implies that a necessary condition for $\phi_{m,z} < \phi_{m,y}$ to hold is

$$y(\bar{\theta})z(\underline{\theta}) - y(\underline{\theta})z(\bar{\theta}) > 0 \iff \frac{y(\bar{\theta})}{z(\underline{\theta})} > \frac{y(\underline{\theta})}{z(\bar{\theta})}$$

i.e., a worker of ability $\bar{\theta}$ must have a comparative advantage at task y , while a worker of ability $\underline{\theta}$ at task z .

The greater accuracy of the inference process about ability at z than at y makes it profitable for the firm to assign the worker task z exactly when it is most uncertain about his ability, i.e., its posterior belongs to the intermediate belief range $[\phi^*, \phi^{**}]$. In these instances, the informational gain from observing the worker's performance at task z offsets the profit loss the firm incurs by not allocating the worker task y , in case he is of high productivity, or not firing him, in case he is of low productivity. The value of learning at these intermediate states is also large enough to overcome the efficiency loss the firm suffers by violating the pattern of comparative advantages. Moreover, since $\phi^* < \tilde{\phi}_y$, employment is profitable for values of ϕ for which the firm is better off by not employing the worker if only task y is available. What lies at the heart of our result is the trade-off the firm faces between maximizing the value of information, on one hand, and short-run profit, on the other, characteristic of incomplete information frameworks.

Along the lines of the argument used for the case in which only task y is available, it is also possible to show that the equilibrium in which the firm employs the worker at wage \bar{U} and assigns him task z or y , respectively, if $\phi \in [\phi^*, \phi^{**})$ or $\phi \in [\phi^{**}, 1]$, and does not employ him otherwise, is constrained Pareto efficient and it is essentially unique. Clearly, this equilibrium, entailing the assignment of task z in some states, Pareto improves on the corresponding equilibrium achievable when the firm can only assign the worker task y .

Proposition 6. *Let (A1)-(A2) hold. Suppose also that $\alpha_2\beta_1 \geq \alpha_1\beta_2$ and $(\alpha_2 - \beta_2) - (\alpha_1 - \beta_1) > \alpha_2\beta_1 - \alpha_1\beta_2$. Then:*

(i) Any strategy profile (ω^*, e^*, a^*) is an MPE if and only if

$$\omega^*(\phi) = \begin{cases} \bar{U}, & \text{if } \phi \in [\phi^*, 1] \\ w', & \text{otherwise} \end{cases} \quad e^*(\phi, w) = \begin{cases} y, & \text{if } \phi \in [\phi^{**}, 1] \\ z, & \text{if } \phi \in [\phi^*, \phi^{**}) \end{cases}$$

and

$$a^*(\phi, w) = \begin{cases} A, & w \geq \bar{U} \\ R, & \text{otherwise} \end{cases}$$

where $w' \in [0, \bar{U})$ and beliefs are determined according to Bayes' rule. Any such equilibrium is constrained Pareto efficient and ex ante Pareto dominates the corresponding equilibrium of the model in which the worker can only be assigned task y .

(ii) Every MPE is outcome equivalent to some sequential equilibrium.

Proof: By revealed preferences, the fact that any MPE of the model with two tasks dominate ex ante the corresponding equilibrium (i.e., for given firm's compensation strategy and worker's acceptance behavior) of the model with only task y is an immediate consequence of the fact that, for task z to be assigned in equilibrium, it must be that it is more profitable to the firm than y at some value of ϕ . The rest of the proof follows the same argument as the one used in the proof of Proposition 1. \square

2.2.1 Learning and Career Dynamics

To analyze the impact of learning on a worker's career profile, from now on we will assume

$$(A3') : \quad \phi^* < \pi < \phi^{**}$$

i.e., in the first period a worker is employed by the firm and assigned task z . Let $r_t^z(\theta)$ be the probability that a worker of type θ is continuously employed by the firm at task z at least for the first t periods. By (A3'), $r_1^z(\theta) = 1$. The probability that a worker will be continuously employed and assigned task z at least until period t can be computed as

$$r_t^z(\theta) = \sum_{(c_1^z, \dots, c_{t-1}^z) \in C_{t-1}^z} p_z(\theta)^{c_1^z + \dots + c_{t-1}^z} (1 - p_z(\theta))^{t-1-c_1^z - \dots - c_{t-1}^z}$$

where $p_z(\bar{\theta}) = \beta_1$, $p_z(\underline{\theta}) = \beta_2$ and

$$C_{t-1}^z = \{(c_1^z, \dots, c_{t-1}^z) : \lfloor \lambda_z \tau - \underline{\gamma} \rfloor - c_1^z - \dots - c_{\tau-1}^z \leq c_\tau^z < \lfloor \lambda_z \tau + \bar{\gamma} \rfloor - c_1^z - \dots - c_{\tau-1}^z, \\ \text{for } \tau = 1, \dots, t-1\}$$

with

$$\lambda_z = \frac{\ln\left(\frac{1-\beta_2}{1-\beta_1}\right)}{\ln\left(\frac{\beta_1(1-\beta_2)}{\beta_2(1-\beta_1)}\right)} \in (0, 1), \quad \underline{\gamma} = \frac{\ln\left(\frac{\pi(1-\phi^*)}{\phi^*(1-\pi)}\right)}{\ln\left(\frac{\alpha_1(1-\alpha_2)}{\alpha_2(1-\alpha_1)}\right)} \quad \text{and} \quad \bar{\gamma} = \frac{\ln\left(\frac{\pi(1-\phi^{**})}{\phi^{**}(1-\pi)}\right)}{\ln\left(\frac{\beta_1(1-\beta_2)}{\beta_2(1-\beta_1)}\right)}.$$

Define analogously $r_{t+1,t'}^y(\theta)$ to be the probability that a worker of type θ , who is assigned task y at the beginning of period $t+1$, when the firm's posterior is ϕ_{t+1} , will be continuously allocated task y for at least t' consecutive periods. This probability is

$$r_{t+1,t'}^y(\theta) = \sum_{(c_1^y, \dots, c_{t'-1}^y) \in C_{t'-1}^y} p_y(\theta)^{c_1^y + \dots + c_{t'-1}^y} (1 - p_y(\theta))^{t'-1 - c_1^y - \dots - c_{t'-1}^y}$$

where $p_y(\bar{\theta}) = \alpha_1$, $p_y(\underline{\theta}) = \alpha_2$ and

$$C_{t+1,t'-1}^y = \{(c_1^y, \dots, c_{t'-1}^y) : c_\tau^y \geq \lfloor \lambda_y \tau - \gamma(\phi_{t+1}) \rfloor - c_1^y - \dots - c_{\tau-1}^y, \text{ for } \tau = 1, \dots, t'-1\}.$$

where

$$\lambda_y = \frac{\ln\left(\frac{1-\alpha_2}{1-\alpha_1}\right)}{\ln\left(\frac{\alpha_1(1-\alpha_2)}{\alpha_2(1-\alpha_1)}\right)} \in (0, 1) \quad \text{and} \quad \gamma(\phi_{t+1}) = \frac{\ln\left(\frac{\phi_{t+1}(1-\phi^{**})}{\phi^{**}(1-\phi_{t+1})}\right)}{\ln\left(\frac{\alpha_1(1-\alpha_2)}{\alpha_2(1-\alpha_1)}\right)} \geq 0$$

for $\phi_{t+1} \geq \phi^{**}$. In the definition of $r_{t+1,t'}^y$ dates have been relabelled so that 1 is the first period the worker is assigned task y . The following Proposition summarizes properties of the intertemporal profile of a worker's tenure and task assignment.

Lemma 4. *Under assumptions (A1)-(A3) the following results hold:*

(i) *for all t , $r_{t+1,t'}^y(\bar{\theta}) - r_{t+1,t'+1}^y(\bar{\theta}) \leq r_{t+1,t'}^y(\underline{\theta}) - r_{t+1,t'+1}^y(\underline{\theta})$. Then, $r_{t+1,t'}^y(\underline{\theta}) \leq r_{t+1,t'}^y(\bar{\theta})$. For all $\gamma(\phi_{t+1})$ there exists a \bar{t} such that, for all $t > \bar{t}$, these inequalities are strict;*

(ii) *for all t , the probability of being assigned task y after task z is higher (strictly, for t sufficiently large) for a type $\bar{\theta}$ than for a type $\underline{\theta}$ worker. Conditional on being assigned task z , the probability of being fired is higher (strictly, for t sufficiently large) for a type $\underline{\theta}$ than for a type $\bar{\theta}$ worker;*

(iii) *for all t' , the probability of being assigned task z after task y is higher (strictly, for t' sufficiently large) for a type $\underline{\theta}$ than for a type $\bar{\theta}$ worker.*

Proof: See the Appendix. □

Similarly to the case in which only task y is available, the probability of being continuously employed at task y is decreasing over time for each type of worker. But the decrease in the probability of retention at task y , $r_{t+1,t'}^y(\theta) - r_{t+1,t'+1}^y(\theta)$, is smaller when the worker is of high productivity

than when he is of low productivity. As a result, the prospect of continuous employment at task y is more favorable for a type $\bar{\theta}$ than for a type $\underline{\theta}$ worker (Lemma 4 (i)). While a high type worker has a higher probability of being allocated task y after task z than a low type worker, the probability of being fired at task z is higher for a low than for a high ability worker (Proposition 4 (ii)). A low ability worker is also more likely to be allocated task z , after having performed task y , than a high ability worker (Lemma 4 (iii)).

The fact that a worker performing task z is assigned task y once $\phi \geq \phi^{**}$, but is dismissed altogether if $\phi < \phi^*$, corresponds to the use of task z as a sequential screening device, specifically as a version of Wald's *sequential probability ratio test* (SPRT). The SPRT is a sequential procedure for testing a simple hypothesis H_0 against a simple alternative H_1 . It prescribes sampling to continue as long as an appropriate statistic of the observations falls in a predetermined interval and to accept or reject H_0 (i.e., $\theta = \underline{\theta}$) depending on whether the test statistics falls in the pre-specified acceptance or rejection ('critical') region. The SPRT satisfies a number of optimality properties: it minimizes the expected sample size over all fixed-sample-size tests having the same significance level (i.e., the probability of a type I error or, equivalently, the rejection of H_0 when H_0 is true) and power (i.e., the probability of rejecting the hypothesis under consideration when $\theta = \underline{\theta}$ and $\theta = \bar{\theta}$). It is also a uniformly most-powerful test in the sense of Neyman-Pearson, i.e., the power of the test associated with the chosen critical region is at least as large as the power of the test associated with any other critical region of the same significance level. In this sense, screening the worker at z is information wise efficient.

The trade-off between the type I error (i.e., assigning a low ability worker task y) and the type II error (i.e., firing a high ability worker) is apparent. The longer the period in which a worker is continuously assigned task z , the more accurate the inference about the worker's ability but the larger the loss the firm incurs by employing a worker whose actual productivity might be low or the larger profit at y if $\theta = \bar{\theta}$. On the other hand, the shorter the time a worker is allocated task z , the higher the risk the firm faces of dismissing a high ability worker or assigning y to a low ability worker, but the lower the loss in case of employment of a low productivity worker. Notice that the indeterminacy in the choice of the size of the two errors, characteristic of the SPRT, is solved under the firm's optimal employment strategy. By construction, in fact, it balances efficiently the need for maximization of the returns from the sampling process and the need to infer the value of the unknown parameter of interest, θ . As we will show, though, even as the time horizon becomes arbitrarily large the probability of a type II error is still strictly positive.

2.2.2 Limiting Retention and Task Assignment

Because of the trade-off between learning and short-run profit maximization, the benefit of assigning a worker task z in our framework is only temporary, being related to the the greater informativeness of task z with respect to task y in the inference process about ability. The intuition is as follows. The profitability to the firm of allocating a worker task z derives from the possibility of testing whether the worker's ability is indeed high at a lower cost, in terms of output destruction, than by assigning him task y . By the Law of Large Numbers, though, the average number of successes observed along a history in which a worker continuously performs task z will eventually converge to the value of the worker's actual ability. If the worker is revealed as being of high type, therefore, the firm will assign him permanently task y , the most profitable one if his productivity is $\bar{\theta}$, but, if his performance at task z is unsatisfactory, the firm will fire him.

As before, because of the opportunity cost incurred by the firm in experimenting, only with strictly positive probability a worker will be retained forever at the firm. Still, the conditional probability of retention of a high productivity worker converges to 1 as t grows arbitrarily large. Formally, let $p_{t'+1;t'}^y(\theta)$ denote the conditional probability that a worker of type θ who has been continuously employed for the first t' periods at task y will be employed at y in period $t' + 1$ as well. Then:

Proposition 7. *Under assumptions (A1)-(A3) the following limiting results hold:*

- (i) $\lim_{t \rightarrow \infty} r_t^z(\underline{\theta}) = \lim_{t \rightarrow \infty} r_t^z(\bar{\theta}) = 0$.
- (ii) $0 = \lim_{t' \rightarrow \infty} r_{t'+1;t'}^y(\underline{\theta}) < \lim_{t' \rightarrow \infty} r_{t'+1;t'}^y(\bar{\theta}) = \bar{r} < 1$ and $\bar{r} \geq \bar{p}$.
- (iii) $\lim_{t' \rightarrow \infty} p_{t'+1;t'}^y(\bar{\theta}) = 1$.

Proof: See the Appendix. □

Notice that in the limit the risk of assigning a low ability worker task y is zero, i.e., the probability of a type I error is zero. Since the purpose of having a worker perform task z is to screen out low types, over time the firm always acquires enough information to convince it to either assign the worker task y or dismiss him. As before, though, the firm's learning is not complete: even if a low ability worker will be fired almost surely, the probability of retaining in the limit a high ability worker is bounded away from one. In other words, even in the long run, then, the type II error is strictly positive. Still, assigning task z improves (at least weakly) the prospect of permanent retention of a high productivity worker, compared to the one-task case, because of the firm's increased willingness to experiment (CONTINUE).

2.3 Costly Switch of Tasks

Suppose now that the firm incurs the cost $c_z > 0$ when assigning the worker task z if he was allocated task y in the previous period and, symmetrically, the cost $c_y > 0$ when assigning task y if the worker has previously performed task z . In the presence of switching costs, the state of the firm's Bandit problem in a period cannot be adequately described by the updated posterior about the worker's ability, except at date 1. Rather, it also includes the arm (i.e., task) which was in use in the period immediately preceding.¹⁴ Let $x \in \{y, z\}$ denote the incumbent arm. The firm's employment problem admits then a new value function, $V : [0, 1] \times \{y, z\} \rightarrow \mathbb{R}$, which satisfies the following Bellman optimality equations at each (ϕ, x)

$$\begin{aligned} V(\phi, z) &= \max\{\bar{\Pi}, V_z(\phi, z), V_y(\phi, z)\} \\ &= \max\{\bar{\Pi}, (1 - \delta)[z(\phi) - \bar{U}] + \delta E_z V(\phi, z), (1 - \delta)[y(\phi) - \bar{U} - c_y] + \delta E_y V(\phi, y)\} \\ V(\phi, y) &= \max\{\bar{\Pi}, V_z(\phi, y), V_y(\phi, y)\} \\ &= \max\{\bar{\Pi}, (1 - \delta)[z(\phi) - \bar{U} - c_z] + \delta E_z V(\phi, z), (1 - \delta)[y(\phi) - \bar{U}] + \delta E_y V(\phi, y)\} \end{aligned}$$

where $E_z V(\phi, z) \equiv E[V(\varphi(\phi | z), z) | \phi]$ and $E_y V(\phi, y) \equiv E[V(\varphi(\phi | y), y) | \phi]$. Any measurable selection from the correspondence of maximizers of the equations above constitutes a stationary Markovian optimal strategy for the firm. Let

$$(A4) : \quad y(\bar{\theta}) - (1 - \delta)c_y > z(\bar{\theta}) > z(\underline{\theta}) > y(\underline{\theta})$$

which ensures that, if the worker is of the highest ability and z was the arm in use in the previous period, the firm is better off by assigning him task y rather than task z . Define $\phi^*(z)$ and $\phi^{**}(z)$ to be, respectively, the cut-off belief values which make the firm indifferent between employing the worker at task z or not employing him and between assigning him task z or task y , when z is the incumbent arm. Let, similarly, $\phi^*(y)$ and $\phi^{**}(y)$ be the cut-off belief values for which the firm is indifferent, respectively, between employing the worker at task z or not employing him and between assigning him task z or task y , when y is the incumbent arm. Notice that the largest value of c_y for which task y is assigned in equilibrium is the one for which, when $\phi = 1$, the two tasks are equally profitable, given our assumption that, when indifferent, the firm allocates the worker task y instead of task z . Solving for c_y the equation $V_z(1, z) = V_y(1, z)$, it is immediate that this value is given by $[y(\bar{\theta}) - z(\bar{\theta})]/[1 - \delta]$. In fact, the first inequality in assumption (A4) ensures $0 < c_y < [y(\bar{\theta}) - z(\bar{\theta})]/(1 - \delta)$.

¹⁴See Banks and Sundaram [1994] for a detailed discussion.

Proposition 8. *Suppose (A4) holds, $\alpha_2\beta_1 \geq \alpha_1\beta_2$ and $(\alpha_2 - \beta_2) - (\alpha_1 - \beta_1) > \alpha_2\beta_1 - \alpha_1\beta_2$. Then, there exists $\underline{\delta} \in (0, 1)$ such that, for all $\delta > \underline{\delta}$, $0 < \phi^*(z) < \phi^{**}(z) < 1$ and $0 < \phi^*(y) < \phi^{**}(y) < 1$.*

Moreover:

(i) *for any fixed $c_y > y(\bar{\theta}) - z(\bar{\theta})$, there exists a $\bar{c}_z > 0$ such that, for all $0 < c_z < \bar{c}_z$, $V(\phi, y) > V(\phi, z)$. Then, $\phi^* \leq \phi^*(z) < \phi^*(y) < \phi^{**}(y) < \phi^{**} < \phi^{**}(z)$;*

(ii) *if $c_z \geq \bar{\Pi} + \bar{U} - y(\underline{\theta})$, a worker assigned task y can be fired but is never allocated task z .*

Proof: See the Appendix. □

Suppose that y is the incumbent arm. It is immediate that, with respect to the model without switching costs, the threshold belief value $\phi^*(y)$ which makes the firm indifferent between collecting the outside option and employing the worker at y increases whenever $c_y > 0$, i.e., $\phi^* < \phi^*(y)$. It is also intuitive that, as the firm incurs a cost by changing the worker job assignment in a period, its willingness to switch the worker across tasks decreases. In our framework this implies that, once the worker is allocated z (respectively, y), the range of posteriors for which task y (respectively, z) is profitable is smaller with respect to the no-cost case, i.e., $\phi^{**} < \phi^{**}(z)$ (respectively, $\phi^{**}(y) < \phi^{**}$). Once costs are present, these assignment frictions change the relative profitability of the two tasks, so that, when the worker performs task z , y is relatively less attractive to the firm and, similarly, when the worker is assigned task y , z is profitable for a smaller range of posteriors, i.e., $\phi^*(z) < \phi^*(y) < \phi^{**}(y) < \phi^{**}(z)$.

By interpreting tasks as job positions, the claim under (ii) is an instance of the ‘no demotion’ result (see, for instance, Gibbons and Waldman [1999b]), according to which low performing workers, once promoted, can be fired but are almost never demoted back to the position from which they have been promoted out. Notice that the result also proves a dynamic test of the learning versus the human capital explanation of promotion dynamics: in the latter case, in fact, according to the interpretation that promotion rewards the accumulation of nonverifiable firm-specific human capital, demotions cannot occur in equilibrium, but firings do not take place either. In a number of firm-level studies, instead, as documented for instance by Baker, Gibbs and Holmström [1994a] (henceforth, BGH), even if demotions are rare, exit is almost uniform across levels of a firm’s hierarchy.

2.4 Learning on the Job

Suppose now that ability is still one-dimensional but imperfectly correlated across tasks (CONTINUE).

3 Two-Dimensional Ability

Suppose that the worker's unknown productivity at the firm can now be described by the bundle (θ_1, θ_2) . As a difference with respect to the baseline case, the worker's ability is *not* perfectly correlated across tasks. Assume, in particular, that the worker can be one of three types, $(\bar{\theta}, \bar{\theta})$ (type 1), $(\bar{\theta}, \underline{\theta})$ (type 2), and $(\underline{\theta}, \underline{\theta})$ (type 3). Let ϕ^1 denote the firm and the worker's posterior that the worker is of type 1 and ϕ^2 that he is of type 2. As before, the performance signal at each task, where, as before, only high or low output can realize, reveals information about the worker's underlying ability. Specifically, $\Pr(\bar{y} \mid \bar{\theta}, \bar{\theta}) = \alpha_1$, $\Pr(\bar{y} \mid \bar{\theta}, \underline{\theta}) = \alpha_2$ and $\Pr(\bar{y} \mid \underline{\theta}, \underline{\theta}) = \alpha_3$, with $\alpha_1, \alpha_2, \alpha_3 \in (0, 1)$ and $\alpha_1 > \alpha_2 > \alpha_3$.¹⁵ At task z , similarly, $\Pr(\bar{z} \mid \bar{\theta}, \bar{\theta}) = \beta_1$, $\Pr(\bar{z} \mid \bar{\theta}, \underline{\theta}) = \beta_2$, and $\Pr(\bar{z} \mid \underline{\theta}, \underline{\theta}) = \beta_3$. Let the one-period expected profit satisfy

$$(A5) : y(\bar{\theta}, \bar{\theta}) > \bar{\Pi} + \bar{U} > z(\bar{\theta}, \bar{\theta}) > z(\bar{\theta}, \underline{\theta}) > y(\bar{\theta}, \underline{\theta})$$

$$(A6) : z(\underline{\theta}, \underline{\theta}) > y(\underline{\theta}, \underline{\theta}).$$

The firm's value function for the new problem is given by

$$V(\phi^1, \phi^2) = \max\{\bar{\Pi}, V_z(\phi^1, \phi^2), V_y(\phi^1, \phi^2)\} \quad (4)$$

where

$$V_z(\phi^1, \phi^2) = (1 - \delta)[\phi^1 z(\bar{\theta}, \bar{\theta}) + \phi^2 z(\bar{\theta}, \underline{\theta}) + (1 - \phi^1 - \phi^2)z(\underline{\theta}, \underline{\theta}) - \bar{U}] + \delta E_z V(\phi^1, \phi^2)$$

$$V_y(\phi^1, \phi^2) = (1 - \delta)[\phi^1 y(\bar{\theta}, \bar{\theta}) + \phi^2 y(\bar{\theta}, \underline{\theta}) + (1 - \phi^1 - \phi^2)y(\underline{\theta}, \underline{\theta}) - \bar{U}] + \delta E_y V(\phi^1, \phi^2).$$

Using standard arguments it is possible to show that the firm's value function V is continuous, increasing and convex in (ϕ^1, ϕ^2) . Increasingness, in particular, is a direct consequence of the fact that, at z , $z(\bar{\theta}, \bar{\theta}) - z(\underline{\theta}, \underline{\theta})$ and $z(\bar{\theta}, \underline{\theta}) - z(\underline{\theta}, \underline{\theta})$, while, at y , $y(\bar{\theta}, \bar{\theta}) - y(\underline{\theta}, \underline{\theta})$ and $y(\bar{\theta}, \underline{\theta}) - y(\underline{\theta}, \underline{\theta})$. Moreover, $V(0, 0) = V(0, 1) = \bar{\Pi}$ and $V(1, 0) = y(\bar{\theta}, \bar{\theta}) - \bar{U}$. As in the baseline case, the relative size of the probability of success at either task has implications for the informativeness of the output signal at z and y . In particular:

Lemma 5. *Let (A5) and (A6) hold. If $\alpha_2\beta_1 = \alpha_1\beta_2$ and $(\alpha_3 - \beta_3) - (\alpha_1 - \beta_1) > \alpha_3\beta_2 - \alpha_1\beta_3 > 0$, then the conditional distribution of the updated posterior on ϕ^1 and ϕ^2 at task z is a mean-preserving spread of the corresponding distributions at y .*

Proof: Notice that $\varphi(\phi^1 \mid \bar{z}) = \beta_1\phi^1 / [\beta_1\phi^1 + \beta_2\phi^2 + \beta^3(1 - \phi^1 - \phi^2)]$ and $\varphi(\phi^1 \mid \bar{y}) = \alpha_1\phi^1 / [\alpha_1\phi^1 + \alpha_2\phi^2 + \alpha^3(1 - \phi^1 - \phi^2)]$, while $\varphi(\phi^2 \mid \bar{z}) = \beta_2\phi^1 / [\beta_1\phi^1 + \beta_2\phi^2 + \beta^3(1 - \phi^1 - \phi^2)]$ and $\varphi(\phi^2 \mid \bar{y}) =$

¹⁵Note that this implies that the output distribution of type 1 first-order stochastically dominates the output distribution of type 2, which in turn first-order stochastically dominates the one of type 3.

$\alpha_2\phi^1/[\alpha_1\phi^1+\alpha_2\phi^2+\alpha^3(1-\phi^1-\phi^2)]$. Then, $\alpha_2\beta_1 \geq \alpha_1\beta_2$ and $\alpha_3\beta_1 > \alpha_1\beta_3$ imply $\varphi(\phi^1 | \bar{z}) > \varphi(\phi^1 | \bar{y})$, while $\alpha_1\beta_2 \geq \alpha_2\beta_1$ and $\alpha_3\beta_2 > \alpha_2\beta_3$ imply $\varphi(\phi^2 | \bar{z}) > \varphi(\phi^2 | \bar{y})$. Since $\alpha_3\beta_1 > \alpha_3\beta_2$, given $\beta_1 > \beta_2$, and $\alpha_1\beta_3 > \alpha_2\beta_3$, given $\alpha_1 > \alpha_2$, for both inequality to hold it is enough $\alpha_2\beta_1 = \alpha_1\beta_2$ and $\alpha_3\beta_2 > \alpha_1\beta_3$. \square

To guarantee that task z is therefore ‘more informative’ than task y , let then

$$(A7): \quad \alpha_2\beta_1 = \alpha_1\beta_2, \quad \alpha_3\beta_2 > \alpha_1\beta_3.$$

Intuitively, when the firm is sufficiently convinced that the worker is of type 3, the informational gain from experimentation is too low to make employment profitable. In this case, then, the outside option is for the firm the dominant alternative. On the other hand, task z is as profitable as the outside option for the firm whenever $\bar{\Pi} = V_z(\phi^1, \phi^2)$ or, equivalently,

$$\bar{\Pi} = (1 - \delta)[\phi^1 z(\bar{\theta}, \bar{\theta}) + \phi^2 z(\bar{\theta}, \underline{\theta}) + (1 - \phi^1 - \phi^2)z(\underline{\theta}, \underline{\theta}) - \bar{U}] + \delta E_z V(\phi^1, \phi^2). \quad (5)$$

The conditions which guarantee that the equality is satisfied amount to requiring that the posterior ϕ^1 is high enough to offset the (flow) profit loss the firm would incur by employing a type 2 or a type 3 worker.

Lemma 6. *Let (A5)-(A7) hold. Suppose $\delta \geq \underline{\delta} = [\bar{\Pi} - z(\bar{\theta}, \bar{\theta})]/[y(\bar{\theta}, \bar{\theta}) - z(\bar{\theta}, \bar{\theta})]$. Then, for any $\phi^2 \in [0, \bar{\phi}^2]$ there exists a unique $\phi^1 \in [\phi_z^*(\bar{\phi}^2), \phi_z^*(0)]$, determined implicitly by (5), such that $\bar{\Pi} = V_z(\phi^1, \phi^2)$. Equivalently, for any $\phi^3 \in [0, \bar{\phi}^3]$ there exists a unique $\phi^1 \in [\phi_z^*(0), \phi_z^*(\bar{\phi}^3)]$ such that $\bar{\Pi} = V_z(\phi^1, \phi^3)$.*

Proof: See the Appendix. \square

Observe that the state of the game for both the firm and the worker can be described by their common posterior about the worker’s ability, in this case the updated distribution on (θ_1, θ_2) . Then, the characterization of the locus of points where task z is as profitable as the firm’s outside option can be equivalently stated in terms of ϕ^1 and ϕ^2 or in terms of ϕ^1 and ϕ^3 . Still, to verify that, in the three type case as well, task z is valuable only for intermediate belief values, i.e., once y is allocated to the worker, it continues to be optimal for any higher values of ϕ^1 , it is necessary to show that a single crossing property holds for both the difference $V_y(\phi^1, \phi^2) - V_z(\phi^1, \phi^2)$ and the difference $V_y(\phi^1, \phi^3) - V_z(\phi^1, \phi^3)$.

Lemma 7. *Under (A5)-(A7), the following holds:*

(i) *there exists $\underline{\delta} \in (0, 1)$ such that, for all $\delta > \underline{\delta}$, at any (ϕ^1, ϕ^2) for which $\bar{\Pi} = V_z(\phi^1, \phi^2)$, the firm strictly prefers task z to task y ;*

(ii) if $y(\bar{\theta}, \bar{\theta}) - y(\bar{\theta}, \underline{\theta}) > z(\bar{\theta}, \bar{\theta}) - z(\underline{\theta}, \underline{\theta})$, the differences $V_y(\phi^1, \phi^2) - V_z(\phi^1, \phi^2)$ and $V_y(\phi^1, \phi^3) - V_z(\phi^1, \phi^3)$ are increasing in ϕ^1 . This implies that, for given ϕ^2 or ϕ^3 , there exists a unique value of ϕ^1 , respectively $\phi_{zy}^*(\phi^2)$ or $\phi_{zy}^*(\phi^3)$, for which task z and y are equally profitable. In particular, given (i), $\phi_{zy}^*(\phi^2) > \phi_z^*(\phi^2)$ and $\phi_{zy}^*(\phi^3) > \phi_z^*(\phi^3)$.

Proof: See the Appendix. □

The above Lemma implies that there exists a set of prior distributions for which the worker is assigned task z when first employed. Task z acts as a training ground: once the worker has proved successful at it, i.e., once a sufficiently long sequence of high output, he is allocated task y , where he is retained as long as his performance is satisfactory. These features of the optimal policy are summarized in the following Proposition.

Proposition 9. *Suppose (A5)-(A7) hold. Under the optimal employment policy, solution to the problem described in (4), a worker assigned task z is assigned task y after a sufficient number of high output signals realize and fired after a sufficient number of low output signals. A worker assigned task y is retained at it as long as high output occurs. After a long enough sequence of low output realizations, he can either be assigned task z or fired.*

Proof: A consequence of Lemma 6 and 7. □

Intuitively, the extent to which a worker assigned task y , who has previously performed task z , is allocated z again depends on the set of posterior beliefs for which z is more profitable than y (and of the firm's outside option) and on the informativeness of the output signal at y . Interpreting the assignment of task y , after task z , as a 'promotion', from the above Proposition 'demotions' back to z have positive frequency in equilibrium. However, if the performance signal at y was relatively uninformative, even after poor performance at y a worker would not be assigned z , but he would not be fired either, while a robust finding in the data is that firm-level exit rates are significant at high levels of a hierarchy as well.

As the analysis to follow suggests, one way to reconcile the fact that low performing promoted workers can be fired but are almost never demoted is to think at low level tasks as providing accurate information along some dimensions of interest. Once the informational rationale from assigning them disappears, it is no longer profitable for the firm to allocate them. Still, if there exists residual uncertainty about the worker *overall* skill, once assigned a higher level task a worker might be fired, if he produces a long enough record of bad output realizations.

3.1 The No-Demotion Case

Suppose that task z perfectly reveals whether the worker is of type 3. Namely, $\Pr(\bar{z} | \bar{\theta}, \bar{\theta}) = \Pr(\bar{z} | \bar{\theta}, \underline{\theta}) = 1$, while $\Pr(\bar{z} | \underline{\theta}, \underline{\theta}) = 0$. Let $\varphi(\phi^1 | \bar{z})$ denote the updated posterior, under Bayes' rule, that the worker is of type 1 after a high output realization at z . Then, $\varphi(\phi^1 | \bar{z}) = \phi^1 / (\phi^1 + \phi^2) \equiv \tilde{\phi}^1$. Similarly, $\varphi(\phi^2 | \bar{z}) = \phi^2 / (\phi^1 + \phi^2) \equiv 1 - \tilde{\phi}^1$, while $\varphi(\phi^3 | \bar{z}) = 0$. Then, $z(\bar{\theta}, \bar{\theta}) = z(\bar{\theta}, \underline{\theta})$. The value to the firm of using task z can now be expressed as

$$V_z(\phi^1, \phi^2) = (1 - \delta)[(\phi^1 + \phi^2)z(\bar{\theta}, \bar{\theta}) + (1 - \phi^1 - \phi^2)z(\underline{\theta}, \underline{\theta}) - \bar{U}] \\ + \delta(\phi^1 + \phi^2)\tilde{V}(\tilde{\phi}^1) + \delta(1 - \phi^1 - \phi^2)\bar{\Pi}$$

where $\tilde{V}(\tilde{\phi}^1) = \max\{\bar{\Pi}, \tilde{V}_y(\tilde{\phi}^1)\}$ and

$$\tilde{V}_y(\tilde{\phi}^1) = (1 - \delta)[\tilde{\phi}^1 y(\bar{\theta}, \bar{\theta}) + (1 - \tilde{\phi}^1)y(\bar{\theta}, \underline{\theta}) - \bar{U}] + \delta E_y \tilde{V}(\tilde{\phi}^1).$$

The informativeness of the output signal at task z has natural implications for the equilibrium assignment decision. Once a worker is assigned task z in a period, the firm's remaining uncertainty is only as to whether he is actually fit for y . Should the worker not perform satisfactorily at it, the firm is then better off by firing him. Formally:

Lemma 8. *Suppose (A5) and (A6) hold. If the worker is assigned task z and a low output signal realizes, the worker is permanently fired. If the worker is assigned task z and a high output signal occurs, he is allocated task y . The retention rule at y is a reservation-belief strategy, i.e., the worker is employed at task y if and only if $\phi^1 / (\phi^1 + \phi^2) \geq \phi_y^*$ and dismissed (forever) otherwise.*

Proof: The first part is due to the fact that $\varphi(\phi^1 | \underline{z}) = \varphi(\phi^2 | \underline{z}) = 0$, while $\varphi(\phi^3 | \underline{z}) = 1$, and $\bar{\Pi} > y(\underline{\theta}, \underline{\theta})$. The second part is an immediate consequence of the fact that the firm's employment problem at task y , if the worker has performed task z in the previous period, can be described by the value function $V(\varphi(\phi^1 | \bar{z}), \varphi(\phi^2 | \bar{z})) \equiv V(\tilde{\phi}^1)$ (since $\varphi(\phi^2 | \bar{z}) = \varphi(\phi^1 | \bar{z})$), given by

$$\tilde{V}(\tilde{\phi}^1) = \max\{\bar{\Pi}, \tilde{V}_y(\tilde{\phi}^1)\} \\ = \max\{\bar{\Pi}, (1 - \delta)[\tilde{\phi}^1 y(\bar{\theta}, \bar{\theta}) + (1 - \tilde{\phi}^1)y(\bar{\theta}, \underline{\theta}) - \bar{U}] + \delta E_y \tilde{V}(\tilde{\phi}^1)\} \quad (6)$$

as in the standard one-armed bandit. As shown in the baseline model for the one-task case, the cut-off belief strategy under which the worker is employed at y if and only if $\tilde{\phi}^1 \geq \phi_y^*$, where ϕ_y^* is determined so as to equate the reservation value $\bar{\Pi}$ to $\tilde{V}_y(\tilde{\phi}^1)$, and terminated otherwise is the unique optimal solution. The necessity part of the statement derives from the fact that in this framework a sequential equilibrium is a Markov perfect equilibrium. \square

As before, the conditions under which task z is as profitable for the firm as the outside option amount to requiring that the informational gain from trying the worker's ability at z is sufficient to compensate the firm for employing possibly a low skill worker. As a difference from the previous three type case, though, the trade-off at task z is such that, provided the discount factor is sufficiently high, an increase in ϕ^2 *decreases* its profitability, since it makes experimentation at y riskier, i.e., it decreases $\tilde{V}(\tilde{\phi}^1)$. The appropriate version of Lemma 6 is then:

Lemma 9. *Suppose (A5) and (A6) hold. Then, there exists $\underline{\delta} \in (0, 1)$ such that, for all $\delta > \underline{\delta}$, the following holds: for any $\phi^2 \in [0, \bar{\phi}^2]$ there exists a unique $\phi^1 \in [\phi_z^*(0), \phi_z^*(\bar{\phi}^2)]$ such that $\bar{\Pi} = V_z(\phi^1, \phi^2)$. Equivalently, for any $\phi^3 \in [0, \bar{\phi}^3]$ there exists a unique $\phi^1 \in [\phi_z^*(\bar{\phi}^3), \phi_z^*(0)]$ such that $\bar{\Pi} = V_z(\phi^1, \phi^3)$.*

Proof: See the Appendix. □

Notice that Lemma 7 applies to the present case as well. Then, the optimal policy can be characterized as follows:

Proposition 10. *Suppose (A5) and (A6) hold. Under the optimal employment policy, solution to the problem described in (4), a worker assigned task z is assigned task y after a high output realization and fired after a low one. A worker assigned task y , after having performed task z in the period immediately preceding, can be fired after a sufficiently long sequence of low output realizations but he is never assigned task z again.*

Proof: A consequence of Lemma 9 and 7. □

This characterization result implies a ‘no-demotion’ feature of the optimal policy. Once the worker is assigned task y he can only be continually allocated task y or fired, consistently with the observation that workers promoted out of a given job level are almost never re-assigned to it, even if, in case their performance becomes unsatisfactory at the new level, they can still be terminated.¹⁶

4 Labor Market Competition

Suppose now that there exists a market in which 2 firms compete for the labor services of the worker by announcing simultaneously a wage and a task assignment at the beginning of each period t . Both firms have the same unit costs normalized to zero. Upon entering the labor market, the worker is randomly matched with one of the two firms. Assume also that the worker's performance is publicly

¹⁶Notice that here a promotion does not bring a wage increase to the worker, but it entails a change in task, under the interpretation that different tasks reflect different job contents.

observable to all market participants. For simplicity, let both the firms and the worker's outside option be equal to zero. We will show in the following that the interaction of the strategic aspects of outside labor market competition with firm-level experimentation does alter the efficiency of each firm's assignment strategies and, as a consequence, of turnover.

Bergemann and Välimäki [1996] (henceforth BV) analyze a very similar set up in which two sellers price compete to provide a good of uncertain quality to a single buyer and they proved that all MPE of this game are efficient. In our framework, their result would imply efficiency of turnover in case each firm's decision consisted *only* in whether or not to hire the worker. Their intuition is that strategic competition can sustain efficient learning (and, therefore, dynamically efficient trade) if the exchange of the costs and benefits of experimentation is frictionless both intertemporally and interpersonally. Efficiency, and multiplicity of equilibria, derives from the fact that equilibrium prices only determine different allocations of the surplus among the parties.

Similarly, in our case as well the firm employing the worker is selected efficiently in equilibrium. But, as a difference with respect to BV, competition renders the assignment decision inefficient, since the two firms might not agree on the relative profitability of the two tasks, while the equilibrium assignment only internalizes the preference over the two tasks of the firm currently employing the worker. This in turn implies that MPE are typically inefficient. Equilibria are efficient only under the restriction that technologies are sufficiently similar, i.e., if the probability of a high output signal is the same across firms, or, equivalently, if firms' marginal valuations of ability are sufficiently congruent. Define, analogously to BV, an MPE (in stationary strategies) to be *cautious* if the firm not employing the worker is indifferent between hiring and not hiring the worker. Notice that now assumptions (A1)-(A2) are required to hold at each firm, with probability of success at either task given, respectively, by α_x^k , for type $\bar{\theta}$ and β_x^k , for type $\underline{\theta}$, with $k \in \{1, 2\}$ and $x \in \{z, y\}$. Correspondingly, $z^k(\theta)$ and $y^k(\theta)$ are, respectively, the expected product of a worker of type θ at task z or y in firm k . Denote by i the firm employing the worker and by j the firm not employing the worker. Then:

Proposition 11. *If assumptions (A1)-(A2) are satisfied at each firm, then:*

(i) *all MPE entail inefficient turnover and task assignment, unless $\alpha_x^j = \alpha_x^i$ and $\beta_x^j = \beta_x^i$, for $x \in \{z, y\}$, and $z^j = az^i + b$, $y^j = ay^i + b$, with $a > 0$ and for each $z^i \in \{\bar{z}, \underline{z}\}$, $y^i \in \{\bar{y}, \underline{y}\}$;*

(ii) *in the cautious equilibrium $w^i(\phi) = x^j(\phi)$, where $x^j \in \{z^j, y^j\}$, $i, j \in \{1, 2\}$, and the wage policy of the non-employing firm is a supermartingale, i.e., $w^j(\phi) \geq \delta Ew^j(\varphi(\phi) | x^i)$.*

Proof: See the Appendix. □

Notice that the inefficiency of turnover, as stated in (i), does not depend on the worker's human capital being general or firm-specific, given the assumption of public observability of job assignment

and performance outcomes. On the other hand, efficiency holds only in case technologies merely differ in the size of the *realized* output at either task. In this case, provided the firm are not symmetric, it also follows that the firm employing the firm obtain a non-zero flow profit, as long as $x^i(\phi^i) \neq x^j(\phi^j)$.¹⁷

The characteristics of the wage dynamics generated under the cautious equilibrium closely resemble dynamic patterns widely documented in the data (see, for instance, Gibbons and Waldman [1999a, 1999b]). The four facts which seem to be most supported by the data on wage and promotion dynamics inside firms are: (i) real wage decreases are frequent but demotions are rare, (ii) promotion rates are serially correlated, (iii) wage increases at promotions are significant but small compared to the differences between average wages across levels of a job ladder, and (iv) workers who receive large wage increases early in their stay at one level of the job ladder are promoted more quickly to the next, i.e., wage increases ‘predict’ promotions. While the first cannot be properly addressed in the present context of a one-good economy, serial correlation in promotion rates, i.e., the fact that promotion rates *decrease* with tenure in the current job, depends in our model on how far apart the prior π is from the cut-off belief value for which a firm is indifferent between task z and y . For low levels of π , it can be typically *increasing* in tenure on z , given that a sufficient number of high output realizations must realize for the worker to be assigned y . In other words, only for the *most* able at z , i.e., the workers on z with the highest posteriors, serial correlation holds.¹⁸

The fact that promotions are associated with wage increases, but that wage increases are small relative to the difference between average wages across levels of the job ladder, is an immediate consequence of the output signal at each task being discrete and of the characteristics of the firm’s assignment policy. For a worker to be assigned y , in fact, the last output realized at z *must* have been a high one, which implies an increase in *paid* wage between the last period in z and the first period in y (even if, as shown in Proposition 11, being the belief process a martingale, the *expected* wage is always equal to the current wage). On the other hand, the fact that the assignment of either task occurs for an interval of posterior values, whose lengths depend on how farther apart, ϕ^* is from 0 and ϕ^{**} is from 1, the size of this wage increase can be small compared to the difference between average wages at z and at y (AND THE DISPERSION OF WAGES AT EITHER TASK/JOB DEPENDS ON THE FREQUENCY OF A SWITCH OF FIRM). The fact that workers who receive large wage increases early in their stay at one level of the job ladder are promoted quickly to the next, i.e., the probability of promotion in $t + 1$ is an increasing function of the wage received in t , in our framework derives from the fact that the paid wage is linear in ϕ and that the probability of being

¹⁷Observe that the result only depends on expected output being linear in the common posterior. It then carries over to more general specifications of technology.

¹⁸Notice that, on the contrary, in BGH this serial correlation is especially pronounced at low levels of prior tenure.

assigned task y , for a worker currently on z , increases in ϕ .¹⁹

5 Related Literature

There are several related strands of literature. Closest in spirit to ours are learning and human-capital acquisition models which analyze job assignment and career choices. Harris and Weiss [1984] study a two-job matching model of occupational choice in which workers are finitely lived. While in ‘primary’ jobs a worker’s productivity is unknown, so that learning about ability takes place through the observation of output over time, in ‘secondary’ jobs a worker’s ability is known with certainty. As a result of this assumed independence in the structure of uncertainty, when workers are risk neutral they all begin their careers in primary jobs and switch to secondary jobs only after a sufficiently long history of low performances in the primary job. Given the finite retirement age, all workers who achieve at least a certain cumulative output record by a certain age will remain in the primary job until retirement.

Waldman [1984b] analyzes a two-period two-job assignment problem in presence of learning and human capital accumulation. He examines an environment in which a worker’s ability is perfectly observable only by a worker’s current employer after the first period of employment, but that a worker’s job assignment is publicly observable. Both under spot and long-term contracting (a contract consisting in a wage-job assignment pair), and in the latter case when the firm’s contract offer is constrained by what other firms would offer old workers in the second period, equilibrium wage rates tend to be more closely associated with jobs rather than ability levels. Under spot contracting a worker not promoted in the second period of employment experiences a downward sloping age-earnings profile. Moreover, under these assumptions there exists inefficient allocation of workers to jobs, whose severity is negatively correlated with the level of firm-specific human capital in the economy. The presence of human capital acquisition in his model rules out the degenerate equilibrium outcome in which only workers of the highest ability are assigned the highest job in the second period. The optimal employment policy in his model is similar to ours, in the sense that job assignment is determined by threshold ability levels. Since ability is perfectly revealed to the first-period employer after one period of production, the dynamics of learning is ignored. Moreover, in our dynamic bandit setting if output is non-contractible there is no benefit to long-term contracting.

¹⁹Another fact frequently cited is the existence of correlation in wage increases, i.e., the fact that the expected wage increase, conditional on the current wage increase and the last period wage, is an increasing function of the current wage increase: $E\{w_{t+1} - w_t \mid w_t - w_{t-1}, w_{t-1}\}$. In our model, given that wage is linear in ϕ and the belief process is a martingale, the expected wage increase is always zero, both conditionally and unconditionally. Still in BGH, as an example, this feature of wage dynamics does not seem robust.

Gibbons and Waldman [1999b] analyze a model on job assignment, learning and human capital acquisition which accounts for a broad pattern of evidence on wage and career dynamics inside firms. They assume that there exist an output interaction between learning and human capital acquisition which determines the worker’s expected product in a period. Since human capital is accumulated by experience, all workers eventually reach the highest job position in the hierarchy.²⁰ Because of learning on the part of the firm and the accumulation of skills on the part of the worker, demotions are rare. In their model as well the hierarchy is inferred from the pattern of job-to-job transitions and results from the assumption that higher ability is more valuable in higher level jobs. As a difference from theirs, though, in our framework workers move up the job ladder over time purely as a consequence of the firm’s improved estimate of their ability, not because human capital acquisition causes effective ability to increase. In addition mobility to higher-level jobs is specifically identified in our model by the transition from job positions characterized by low output risk to job position associated with greater output risk.

A similar characterization of job transitions is obtained by Jovanovic and Nyarko [1997]’s ‘stepping-stone’ model of occupational mobility, which combines as well learning and human-capital acquisition. The accumulation of skills in their model occurs in the form of acquisition of information on how to perform a given task. Specifically, one period of production resolves the uncertainty about a worker’s production function in a certain occupation and reduces the uncertainty about the profitability of the alternative occupation. As a result, since experimentation is costly in an opportunity cost sense, workers in equilibrium always move from lower to higher paying occupations, i.e., from low to high variance jobs. The model bears a number of similarities with ours. For instance, they predict that the complexity of a job, measured in terms of the size of output destruction in case of a mistake, should increase over a worker’s lifetime. Therefore, activities which are informationally close, in the sense that part of the skills acquired on a given task can be transferred to another occupation, will be part of an occupational ladder in which safer jobs are tried first. The bandit version of their model, unlike ours, displays though the opposite feature.

6 Conclusion

The paper has investigated the impact of firm’s learning about workers’ unknown characteristics on the intertemporal profile of a worker’s tenure and job assignment. Properties of a worker’s career prospect and of the dynamics of job transitions have been derived under alternative assumptions about a firm’s monitoring and production technology and (price) competition in the outside labor

²⁰This occurs even if workers are finitely lived because of their restriction on the length of the time horizon, i.e., $T \geq 5$.

market. In particular, under the hypothesis that the firm can generate different performance signals, depending on the worker’s task or job position, and that firms compete for the labor services of the worker, the characterization of a firm’s optimal employment policy has been shown to generate predictions which are consistent with certain empirical findings concerning wage, promotion and career systems inside firms.

Aspects of the analysis which deserve further investigation relate, for instance, to the trade-off between short-run profit-maximization and learning when workers accumulate human capital on the job (for instance, through an unobservable learning-by-doing component of effort) or to the case in which there is competition for workers in the outside labor market, but firms other than the worker’s current employer can *only* observe the worker’s assignment. These extensions would reinforce our account for the empirically strong link between wage and promotion dynamics (see, for a reference, Baker, Gibbs and Holmström [1994a]), and for the effect of a worker’s job level on the timing (and correlation in the rate) of promotions and on the size of the associated wage increase. They would also enable us to assess the relative predictive power of human capital versus learning models, in presence of a non trivial task assignment problem.

The model also abstracts from the consideration that jobs have scarcity value, arising from output complementarity and capacity constraints (as in athletic teams) or from heterogeneity among cooperating inputs in the production process (as in the assignment of workers of different unknown productivity to pieces of equipment which require at most one operator). These features have a clear impact on sorting and mobility - within and among organizations - and on their effect on wages and occupational mobility.²¹ The investigation of these issues will constitute the object of future research.

Appendix

Proof of Lemma 2: (i) The equivalence between the firm’s optimal stopping problem and the dynamic programming problem defining V^f is established in Hinderer [1970]. The properties of V^f follow from a straightforward application of the Contraction Mapping Theorem. In particular, the increasingness of $E[V^f(\varphi(\phi | y)) | \phi]$ in ϕ derives from the fact that $\phi' < \phi''$ implies that $\varphi(\phi'' | y)$ first-order stochastically dominates $\varphi(\phi' | y)$. From this, in fact, it is possible to conclude that for f increasing in ϕ , $E[f(\varphi(\phi | y)) | \phi]$ is also increasing in ϕ . The convexity of V^f in ϕ can be proved by the same argument used by Banks and Sundaram [1992a] in the proof of Lemma 3.1 part (i). (ii) An immediate consequence of (A1') and (A2'). □

²¹See Davis [1997] for a discussion.

Proof of Proposition 2: The proof builds on the same argument used in the proof of Proposition 2 in Araujo and Camargo [2002]. Let $f(\phi) \equiv \tilde{V}^f(\phi) = y(\phi) - \bar{U}$. By the Bayes' updating rule and the linearity of $\tilde{V}^f(\phi)$ in ϕ , it follows $E[f(\varphi(\phi | y)) | \phi] = y(\phi) - \bar{U}$, so that

$$Tf(\phi) = \max\{\bar{\Pi}, (1 - \delta)\tilde{V}^f(\phi) + \delta E[f(\varphi(\phi | y)) | \phi]\} = \max\{\bar{\Pi}, \tilde{V}^f(\phi)\} \geq f(\phi).$$

By standard arguments it is possible to show that T is a monotone contraction mapping. Since T is monotone, $T^{n-1}f(\phi) \leq T^n f(\phi)$ for all $n \in \mathbb{N}$ and $\phi \in [0, 1]$, which yields $f(\phi) \leq \lim_{n \rightarrow \infty} T^n f(\phi)$. Since T is a contraction, $T^n f$ converges uniformly to its unique fixed point V^f . Then, since uniform convergence implies pointwise convergence, $f(\phi) \leq V^f(\phi)$ for all $\phi \in [0, 1]$. From the definition of $\tilde{\phi}_y$ and the fact that $f(\phi) \leq V^f(\phi)$, it follows

$$\begin{aligned} \bar{\Pi} &= (1 - \delta)[y(\tilde{\phi}_y) - \bar{U}] + \delta E[V^f(\varphi(\tilde{\phi}_y | y)) | \tilde{\phi}_y] \\ &\geq (1 - \delta)[y(\tilde{\phi}_y) - \bar{U}] + \delta E[f(\varphi(\tilde{\phi}_y | y)) | \tilde{\phi}_y] = \tilde{V}^f(\tilde{\phi}_y) \end{aligned}$$

Since $\tilde{V}^f(\phi_{m,y}) = \bar{\Pi}$, the above implies $\tilde{V}^f(\phi_{m,y}) \geq \tilde{V}^f(\tilde{\phi}_y)$. With $\tilde{V}^f(\phi)$ increasing in ϕ , it must be $\phi_{m,y} \geq \tilde{\phi}_y$. Observe that

$$Tf(\phi) = \max\{\bar{\Pi}, f(\phi)\} > f(\phi)$$

for all $\phi \in [0, \phi_{m,y})$, by definition of $\phi_{m,y}$. Since there exists a $\phi \in \text{supp}(\varphi(\phi_{m,y} | y))$ such that $\phi < \phi_{m,y}$, it follows

$$E[Tf(\varphi(\phi_{m,y} | y)) | \phi_{m,y}] > E[f(\varphi(\phi_{m,y} | y)) | \phi_{m,y}] = f(\phi_{m,y})$$

which implies

$$\begin{aligned} (1 - \delta)[y(\phi_{m,y}) - \bar{U}] + \delta E[Tf(\varphi(\phi_{m,y} | y)) | \phi_{m,y}] &> (1 - \delta)[y(\phi_{m,y}) - \bar{U}] + \delta f(\phi_{m,y}) \\ &= y(\phi_{m,y}) - \bar{U} = f(\phi_{m,y}) \end{aligned}$$

and, therefore,

$$\begin{aligned} T^2 f(\phi_{m,y}) &= \max\{\bar{\Pi}, (1 - \delta)[y(\phi_{m,y}) - \bar{U}] + \delta E[Tf(\varphi(\phi_{m,y} | y)) | \phi_{m,y}]\} \\ &> y(\phi_{m,y}) - \bar{U} = f(\phi_{m,y}) \end{aligned}$$

from which it is possible to conclude $V^f(\phi_{m,y}) = \lim_{n \rightarrow \infty} T^n f(\phi_{m,y}) > f(\phi_{m,y})$. Thus,

$$V^f(\phi_{m,y}) > f(\phi_{m,y}) = \bar{\Pi} = V^f(\tilde{\phi}_y)$$

which yields $\phi_{m,y} > \tilde{\phi}_y$. □

Proof of Proposition 3: Notice first that $\rho_t(\theta) \leq \rho_{t+1}(\theta)$ for all $t \geq 1$ and $\theta \in \{\underline{\theta}, \bar{\theta}\}$. The decrease in the probability of retention of a type- θ worker from date t to date $t + 1$ is given by $\psi_t(\theta) = \rho_t(\theta) - \rho_{t+1}(\theta)$, which is equal to the probability that a worker of type θ will be employed for exactly t periods (i.e., he will be fired at the beginning of period $t + 1$). Since $\rho_1(\theta) = 1$, it is possible to show that, for all $t \geq 2$,

$$\psi_t(\theta) \equiv \rho_t(\theta) - \rho_{t+1}(\theta) = \sum_{(c_1, \dots, c_t) \in \bar{C}_t} p(\theta)^{c_1 + \dots + c_t} (1 - p(\theta))^{t - c_1 - \dots - c_t}$$

where $\bar{C}_1 = \{c_1 : c_1 \leq \lfloor \lambda - \gamma \rfloor - 1\}$ and $\bar{C}_t = \{(c_1, \dots, c_t) : (c_1, \dots, c_{t-1}) \in C_{t-1}, c_t \leq \lfloor \lambda t - \gamma \rfloor - c_1 - \dots - c_{t-1} - 1\}$ for $t \geq 2$. Therefore,

$$\psi_t(\bar{\theta}) - \psi_t(\underline{\theta}) = \sum_{(c_1, \dots, c_t) \in \bar{C}_t} \alpha_1^{c_1 + \dots + c_t} (1 - \alpha_2)^{t - c_1 - \dots - c_t} \left\{ \left(\frac{1 - \alpha_1}{1 - \alpha_2} \right)^{t - c_1 - \dots - c_t} - \left(\frac{\alpha_2}{\alpha_1} \right)^{c_1 + \dots + c_t} \right\}.$$

Note that

$$\left(\frac{1 - \alpha_1}{1 - \alpha_2} \right)^{t - c_1 - \dots - c_t} < \left(\frac{\alpha_2}{\alpha_1} \right)^{c_1 + \dots + c_t} \iff c_1 + \dots + c_t < \frac{\ln \left(\frac{1 - \alpha_2}{1 - \alpha_1} \right) \cdot t}{\ln \left(\frac{\alpha_1 (1 - \alpha_2)}{\alpha_2 (1 - \alpha_1)} \right)} \iff c_1 + \dots + c_t < \lambda t$$

but $(c_1, \dots, c_t) \in \bar{C}_t$ implies that $c_1 + \dots + c_t \leq \lfloor \lambda t - \gamma \rfloor - 1 \leq \lambda t$, so that the above holds for all $(c_1, \dots, c_t) \in \bar{C}_t$. Since \bar{C}_t might be empty (for γ high, for example), it follows $\psi_t(\bar{\theta}) - \psi_t(\underline{\theta}) \leq 0$. Also, from

$$\psi_t(\bar{\theta}) - \psi_t(\underline{\theta}) = \rho_t(\bar{\theta}) - \rho_{t+1}(\bar{\theta}) - \rho_t(\underline{\theta}) + \rho_{t+1}(\underline{\theta}) \leq 0$$

it is possible to conclude that, for all $t \geq 1$, $\rho_t(\bar{\theta}) - \rho_{t+1}(\bar{\theta}) \leq \rho_t(\underline{\theta}) - \rho_{t+1}(\underline{\theta})$. Observe that $\rho_1(\underline{\theta}) = \rho_1(\bar{\theta}) = 1$, together with $\rho_1(\bar{\theta}) - \rho_2(\bar{\theta}) \leq \rho_1(\underline{\theta}) - \rho_2(\underline{\theta})$, implies that $\rho_2(\underline{\theta}) \leq \rho_2(\bar{\theta})$. Given

$$\rho_2(\bar{\theta}) - \rho_3(\bar{\theta}) \leq \rho_2(\underline{\theta}) - \rho_3(\underline{\theta}) \leq \rho_2(\bar{\theta}) - \rho_3(\underline{\theta})$$

it also follows $\rho_3(\underline{\theta}) \leq \rho_3(\bar{\theta})$. Proceeding similarly, we obtain $\rho_t(\underline{\theta}) \leq \rho_t(\bar{\theta})$ for all $t \geq 1$. Now, to show that these inequalities hold strictly, it is just sufficient to show that \bar{C}_t is non-empty for t sufficiently large. Let \bar{t} be the first t such that $c_1 + \dots + c_{\bar{t}} < \lfloor \lambda \bar{t} - \gamma \rfloor$, i.e., $\lfloor \lambda \bar{t} - \gamma \rfloor > 0$. Notice that, since γ is finite for any possible choice of π , such \bar{t} always exists. But this implies that, for $t < \bar{t}$, $c_1 + \dots + c_t \geq \lfloor \lambda t - \gamma \rfloor$, so that $C_{\bar{t}-1}$ is non-empty, and, for $t = \bar{t}$, $c_1 + \dots + c_{\bar{t}} < \lfloor \lambda \bar{t} - \gamma \rfloor$, from which it follows that $\bar{C}_{\bar{t}}$ is non-empty either. Therefore, $\bar{C}_{t'}$ is non-empty for all $t' > \bar{t}$. \square

Proof of Proposition 4: (i) Observe that

$$\rho_t(\underline{\theta}) = \sum_{(c_1, \dots, c_{t-1}) \in C_{t-1}} \alpha_2^{c_1 + \dots + c_{t-1}} (1 - \alpha_2)^{t-1 - c_1 - \dots - c_{t-1}} \leq \Pr\{X_{t-1}(\underline{\theta}) \geq \lfloor \lambda(t-1) - \gamma \rfloor\}$$

since $C_{t-1} \subset \{(c_1, \dots, c_{t-1}) : c_1 + \dots + c_{t-1} \geq \lfloor \lambda(t-1) - \gamma \rfloor\}$, where $X_{t-1}(\underline{\theta})$ denotes the number of successes in $t-1$ Bernoulli trials when the probability of success is α_2 . By the Law of Large Numbers for the Bernoulli distribution (see Feller [1965], Ch. 6), we know that, for all $\varepsilon > 0$, $\lim_{t \rightarrow \infty} \Pr\{X_{t-1}(\underline{\theta}) \geq (t-1)(\beta + \varepsilon)\} = 0$. Notice that

$$\lambda = \frac{\ln\left(\frac{1-\alpha_2}{1-\alpha_1}\right)}{\ln\left(\frac{\alpha_1(1-\alpha_2)}{\alpha_2(1-\alpha_1)}\right)} > \alpha_2 \iff (1-\alpha_2) \ln\left(\frac{1-\alpha_1}{1-\alpha_2}\right) + \alpha_2 \ln\left(\frac{\alpha_1}{\alpha_2}\right) < 0$$

where the latter inequality always holds by strict concavity of the logarithm. Since $\lambda > \alpha_2$, we know that there exists a \bar{t} such that, if $t \geq \bar{t}$, then $\lfloor \lambda(t-1) - \gamma \rfloor > \alpha_2(t-1)$. Let $\varepsilon = \frac{\lfloor \lambda(\bar{t}-1) - \gamma \rfloor}{\bar{t}-1} - \alpha_2$. Then, for any $t \geq \bar{t}$,

$$\Pr\{X_{t-1}(\underline{\theta}) \geq \lfloor \lambda(t-1) - \gamma \rfloor\} \leq \Pr\{X_{t-1}(\underline{\theta}) \geq (t-1)(\alpha_2 + \varepsilon)\}$$

since

$$\lfloor \lambda(t-1) - \gamma \rfloor \geq (t-1)(\alpha_2 + \varepsilon) = \frac{(t-1)\lfloor \lambda(\bar{t}-1) - \gamma \rfloor}{\bar{t}-1}$$

if $t \geq \bar{t}$. It is then possible to conclude that $\rho_t(\underline{\theta}) \rightarrow 0$ as $t \rightarrow \infty$. As for the asymptotic behavior of $\rho_t(\bar{\theta})$, observe that the sequence $\{\rho_t(\bar{\theta})\}$ is bounded and decreasing, so it has a limit. The above result, together with $\rho_t(\underline{\theta}) < \rho_t(\bar{\theta})$ for t sufficiently large (Proposition 3), implies that this limit must be strictly positive. In addition, the fact that $\rho_t(\underline{\theta}) \geq \rho_{t+1}(\underline{\theta})$ for all t and, with \bar{C}_t non-empty for t large enough (see Proposition 3), $\rho_{t'}(\underline{\theta}) > \rho_{t'+1}(\underline{\theta})$ for some $t' \geq 2$, this limiting probability is bounded away from one.

(ii) Since $p_{t+1;t}(\underline{\theta}) \equiv \frac{\rho_{t+1}(\underline{\theta})}{\rho_t(\underline{\theta})}$, it follows from the argument under (i) that

$$\lim_{t \rightarrow \infty} p_{t+1;t}(\bar{\theta}) = \lim_{t \rightarrow \infty} \frac{\rho_{t+1}(\bar{\theta})}{\rho_t(\bar{\theta})} = \frac{\bar{\rho}}{\bar{\rho}} = 1$$

as claimed. Notice that this convergence could be highly non-monotone even if the distributions of reward satisfy MLRP. See Banks and Sundaram [1991] for details.

(iii) Notice that $\Pr(T(\underline{\theta}) = \infty) \leq \lim_{t \rightarrow \infty} \rho_t(\underline{\theta}) = 0$, which implies

$$\begin{aligned} E[T(\underline{\theta})] &= \sum_{t=1}^{\infty} \psi_t(\underline{\theta})t = \sum_{t=1}^{\infty} [\rho_t(\underline{\theta}) - \rho_{t+1}(\underline{\theta})]t = \rho_1(\underline{\theta}) - \rho_2(\underline{\theta}) + 2\rho_2(\underline{\theta}) - 2\rho_3(\underline{\theta}) + \dots \\ &= \rho_1(\underline{\theta}) + \rho_2(\underline{\theta}) + \rho_3(\underline{\theta}) + \dots = \sum_{t=1}^{\infty} \rho_t(\underline{\theta}). \end{aligned}$$

By (i) and Bernstein inequality (see Shiryaev [1995], Ch. 1, Section 6, page 55), which characterizes how the empirical mean $\frac{X_{t-1}(\underline{\theta})}{t-1}$ converges to the expected value β , it follows

$$\Pr\{X_{t-1}(\underline{\theta}) \geq (t-1)(\alpha_2 + \varepsilon)\} \leq 2e^{-2\varepsilon^2(t-1)}$$

for fixed $\varepsilon > 0$. Notice that

$$E[T(\theta)] = \sum_{t=1}^{\infty} \rho_t(\theta) \leq \sum_{t=1}^{\infty} 2e^{-2\varepsilon^2(t-1)} = \sum_{t=1}^{\infty} 2e^{-2\varepsilon^2 t} + 2.$$

By the integral test for convergence of a series, the series $\sum_{t=1}^{\infty} 2e^{-2\varepsilon^2 t}$ converges if and only if $\int_1^{\infty} 2e^{-2\varepsilon^2 t} dt$ converges. By integration, it follows

$$\int_1^{\infty} 2e^{-2\varepsilon^2 t} dt = - \left[\frac{2e^{-2\varepsilon^2 t}}{2\varepsilon^2} \right]_{t=1}^{\infty} = \frac{e^{-2\varepsilon^2}}{\varepsilon^2} < \infty$$

which implies that

$$E[T(\theta)] = \sum_{t=1}^{\infty} \rho_t(\theta) \leq \sum_{t=1}^{\infty} 2e^{-2\varepsilon^2 t} + 2 = \frac{e^{-2\varepsilon^2}}{\varepsilon^2} + 2 < \infty$$

for fixed $\varepsilon > 0$. Let now $E_t[T(\bar{\theta})]$ denote the expected number of periods of employment of a worker of type $\bar{\theta}$ at the beginning of period $t \geq 1$. Since $\lim_{t \rightarrow \infty} p_{t+1;t}(\bar{\theta}) = 1$, it follows $\lim_{t \rightarrow \infty} E_t[T(\bar{\theta})] = \infty$. From $E_1[T(\bar{\theta})] \geq E_t[T(\bar{\theta})]$, we obtain $E_1[T(\bar{\theta})] = \infty$. \square

Proof of Lemma 3: (i) The fact that V^f is well-defined, continuous and increasing in ϕ can be showed by a Contraction Mapping argument analogous to the one used in Lemma 2. As for the convexity of V^f , the argument is again an immediate application of the proof of Lemma 3.1 part (i) in Banks and Sundaram [1992a]. (ii) An immediate consequence of (A1) and (A2). \square

Proof of Proposition 5: Recall that, by definition of ϕ_z^* ,

$$\bar{\Pi} = (1 - \delta)[z(\phi_z^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_z^* | z)) | \phi_z^*].$$

Then, the condition

$$(1 - \delta)[z(\phi_z^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_z^* | z)) | \phi_z^*] > (1 - \delta)[y(\phi_z^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_z^* | y)) | \phi_z^*] \quad (7)$$

together with

$$y(1) - \bar{U} > (1 - \delta)[z(1) - \bar{U}] + \delta E[V^f(\varphi(1 | z)) | 1] > \bar{\Pi} \quad (8)$$

yields

$$\begin{cases} \bar{\Pi} > (1 - \delta)[y(\phi_z^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_z^* | y)) | \phi_z^*], \\ y(1) - \bar{U} > \bar{\Pi} \end{cases}$$

and

$$\begin{cases} (1 - \delta)[z(\phi_z^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_z^* | z)) | \phi_z^*] \\ > (1 - \delta)[y(\phi_y^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_y^* | y)) | \phi_y^*], \\ y(1) - \bar{U} > (1 - \delta)[z(1) - \bar{U}] + \delta E[V^f(\varphi(1 | z)) | 1]. \end{cases}$$

From this, it is possible to conclude, respectively, that there exists a unique value $\phi_y^* \in (0, 1)$, with $\phi_y^* > \phi_z^*$, such that

$$\bar{\Pi} = (1 - \delta)[y(\phi_y^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_y^* | y)) | \phi_y^*]$$

and a unique value $\phi^{**} \in (0, 1)$, with $\phi^{**} > \phi_z^*$, satisfying²²

$$\begin{aligned} & (1 - \delta)[z(\phi^{**}) - \bar{U}] + \delta E[V^f(\varphi(\phi^{**} | z)) | \phi^{**}] \\ & = (1 - \delta)[y(\phi^{**}) - \bar{U}] + \delta E[V^f(\varphi(\phi^{**} | y)) | \phi^{**}]. \end{aligned}$$

Since $\phi^{**} > \phi_z^*$, by definition of ϕ_z^* it follows $V^f(\phi^{**}) > \bar{\Pi}$, so $\phi^{**} > \phi_y^*$. Therefore, $\phi_z^* < \phi_y^* < \phi^{**}$, if (7) is satisfied. We will now prove that, under the conditions stated in the Proposition, (7) holds true. Note first that $y(\phi_z^*) - z(\phi_z^*) < 0$ is equivalent to $\phi_z^* < \phi_{zy}^*$, where ϕ_{zy}^* is the belief value for which the firm is indifferent, in static terms, between task z and y . Recall that $\phi_{m,z}$ is the value of the posterior belief satisfying $z(\phi_{m,z}) - \bar{U} = \bar{\Pi}$. It is immediate that $\phi_z^* \leq \phi_{m,z}$. Since $\phi_{m,z} < \phi_{zy}^*$ is equivalent to $y(\phi_{m,z}) - \bar{U} < \bar{\Pi}$ or $\phi_{m,z} < \phi_{m,y}$, this latter restriction, combined with $\phi_z^* \leq \phi_{m,z}$, implies $\phi_z^* < \phi_{zy}^*$ or $y(\phi_z^*) - z(\phi_z^*) < 0$. Notice that, if $\alpha_2\beta_1 > \alpha_1\beta_2$, $\phi_z^h > \phi_y^h$ for ϕ interior. Similarly, $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$ imply, respectively, $\phi_y^h > \phi_z^h$ and $\phi_z^h > \phi_z^l$, if ϕ is interior. Moreover, if $(\alpha_2 - \beta_2) - (\alpha_1 - \beta_1) > \alpha_2\beta_1 - \alpha_1\beta_2$, it follows to $\phi_y^l > \phi_z^l$ for ϕ interior. Therefore, $\phi_z^h > \phi_y^h > \phi_y^l > \phi_z^l$, for $\phi \in (0, 1)$, if $\alpha_2\beta_1 > \alpha_1\beta_2$ and $(\alpha_2 - \beta_2) - (\alpha_1 - \beta_1) > \alpha_2\beta_1 - \alpha_1\beta_2$.

Consider now the distributions of the next period value of ϕ , ϕ' , conditional on the worker being assigned task z or y . Denote these two distributions, respectively, by $F(\phi'; z)$ and $G(\phi'; y)$. The mean of the two distributions is ϕ . Now, the fact that $\phi^{hz} > \phi^{hy} > \phi^{ly} > \phi^{lz}$ and $F(\phi'; z)$ and $G(\phi'; y)$ are two-outcome distributions imply that $F(\phi'; z)$ constitutes a mean-preserving spread of $G(\phi'; y)$. Therefore, $G(\phi'; y)$ second-order stochastically dominates $F(\phi'; z)$. By definition, for any two distributions $F(x)$ and $G(x)$ with the same mean, G second-order stochastically dominates F if $\int \psi(x)dF(x) \geq \int \psi(x)dG(x)$ for every non-decreasing convex function $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}$.²³ It then follows $E[V^f(\varphi(\phi_z^* | z)) | \phi_z^*] \geq E[V^f(\varphi(\phi_z^* | y)) | \phi_z^*]$, where $\psi = V^f$, by convexity of V^f . For $\delta \in (0, 1)$, this condition, together with $y(\phi_z^*) - z(\phi_z^*)$, implies that the inequality in (7) holds true.

²²Uniqueness of ϕ^{**} derives from the monotonicity (in ϕ) of the difference in the dynamic values to the firm of using task z and y (INCLUDE THE ARGUMENT ADAPTED FROM KAKIGI [1983]).

²³See Definition 6.D.2 in Mas-Colell, Whinston and Green [1995].

We now verify that $\phi_z^* = \tilde{\phi}_z$, i.e., ϕ_z^* equals the threshold belief value $\tilde{\phi}_z$ for which the firm is indifferent between employing and not employing the worker if only task z is available. Denote by

$$V_z^f(\phi) = \max\{\bar{\Pi}, (1 - \delta)[z(\phi) - \bar{U}] + \delta E[V_z^f(\varphi(\phi | z)) | \phi]\}$$

the firm's value function for the problem in which an employed worker can only be assigned task z . Then, $\tilde{\phi}_z$ can be defined implicitly as the value of the firm's posterior satisfying

$$V_z^f(\tilde{\phi}_z) = \bar{\Pi} = (1 - \delta)[z(\tilde{\phi}_z) - \bar{U}] + \delta E[V_z^f(\varphi(\tilde{\phi}_z | z)) | \tilde{\phi}_z]. \quad (9)$$

Analogously, $\phi_z^* \in (0, 1)$ has been defined to satisfy

$$\bar{\Pi} = (1 - \delta)[z(\phi_z^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_z^* | z)) | \phi_z^*] \quad (10)$$

where V^f indicates the firm's value function for the problem in which tasks y and z are available. Observe also that, under (7), $V^f(\phi_z^*) = \bar{\Pi}$. By inspection of (9) and (10), it thus follows $\phi_z^* = \tilde{\phi}_z$. For the last part of the claim, observe first that $V^f(\phi) \geq V_y^f(\phi)$ for all ϕ , where V_y^f indicates now the firm's value function for the problem in which only task y is available. Recall that $\tilde{\phi}_y$ is the corresponding threshold belief value for which the firm is indifferent between employing and not employing the worker, defined implicitly by

$$\bar{\Pi} = (1 - \delta)[y(\tilde{\phi}_y) - \bar{U}] + \delta E[V_y^f(\varphi(\tilde{\phi}_y | y)) | \tilde{\phi}_y].$$

Recall from above that $\phi^* \equiv \phi_z^* < \phi_y^* < \phi^{**}$, where ϕ_y^* is defined implicitly by

$$\bar{\Pi} = (1 - \delta)[y(\phi_y^*) - \bar{U}] + \delta E[V^f(\varphi(\phi_y^* | y)) | \phi_y^*].$$

Since $V^f(\phi) \geq V_y^f(\phi)$ for all ϕ and both V^f and V_y^f are monotonic in ϕ , it follows $\phi_y^* \leq \tilde{\phi}_y$, which yields $\phi^* < \tilde{\phi}_y$. From

$$\begin{aligned} (1 - \delta)[z(\phi^{**}) - \bar{U}] + \delta E[V^f(\varphi(\phi^{**} | z)) | \phi^{**}] &= (1 - \delta)[y(\phi^{**}) - \bar{U}] + \delta E[V^f(\varphi(\phi^{**} | y)) | \phi^{**}] \\ &> \bar{\Pi} = (1 - \delta)[y(\tilde{\phi}_y) - \bar{U}] + \delta E[V_y^f(\varphi(\tilde{\phi}_y | y)) | \tilde{\phi}_y] \end{aligned}$$

which implies $\tilde{\phi}_y < \phi^{**}$ since, for $\phi \in [\phi^{**}, 1]$,

$$\begin{aligned} V^f(\phi) &= (1 - \delta)[y(\phi) - \bar{U}] + \delta E[V^f(\varphi(\phi | y)) | \phi] \\ &= (1 - \delta)[y(\phi) - \bar{U}] + \delta E[V_y^f(\varphi(\phi | y)) | \phi] = V_y^f(\phi) \end{aligned}$$

given that, when $\phi \in [\phi^{**}, 1]$,

$$\begin{aligned} V^f(\phi) &= (1 - \delta)[y(\phi) - \bar{U}] + \delta E[V^f(\varphi(\phi | y)) | \phi] \\ &\geq \max\{\bar{\Pi}, (1 - \delta)[z(\phi) - \bar{U}] + \delta E[V^f(\varphi(\phi | z)) | \phi]\} \\ &= (1 - \delta)[z(\phi) - \bar{U}] + \delta E[V^f(\varphi(\phi | z)) | \phi]. \quad \square \end{aligned}$$

Proof of Lemma 4: (i) By definition of $r_{t+1,t'}^y(\theta)$, the probability that a worker of type θ will perform task y for exactly t' consecutive periods, given that he is assigned task y for the first time at the beginning of period $t + 1$, is given by

$$\psi_{t'}^y(\theta) \equiv r_{t+1,t'}^y(\theta) - r_{t+1,t'+1}^y(\theta) = \sum_{(c_1^y, \dots, c_{t'}^y) \in \bar{C}_{t+1,t'}^y} p(\theta)^{c_1^y + \dots + c_{t'}^y} (1 - p(\theta))^{t' - c_1^y - \dots - c_{t'}^y}$$

$$\bar{C}_{t+1,t'}^y = \{(c_1^y, \dots, c_{t'}^y) : (c_1^y, \dots, c_{t'-1}^y) \in C_{t+1,t'-1}^y, c_{t'}^y < \lfloor \lambda_y t' - \gamma(\phi_{t+1}) \rfloor - c_1^y - \dots - c_{t'-1}^y\}$$

where $p(\bar{\theta}) = \alpha$, $p(\underline{\theta}) = \beta$. Then,

$$\psi_{t'}^y(\bar{\theta}) - \psi_{t'}^y(\underline{\theta}) = \sum_{(c_1^y, \dots, c_{t'}^y) \in \bar{C}_{t+1,t'}^y} \alpha^{c_1^y + \dots + c_{t'}^y} (1 - \beta)^{t' - c_1^y - \dots - c_{t'}^y} \left[\left(\frac{1 - \alpha}{1 - \beta} \right)^{t' - c_1^y - \dots - c_{t'}^y} - \left(\frac{\beta}{\alpha} \right)^{c_1^y + \dots + c_{t'}^y} \right].$$

Notice that

$$\left(\frac{1 - \alpha_1}{1 - \alpha_2} \right)^{t' - c_1^y - \dots - c_{t'}^y} < \left(\frac{\alpha_2}{\alpha_1} \right)^{c_1^y + \dots + c_{t'}^y} \iff c_1^y + \dots + c_{t'}^y < \frac{\ln \left(\frac{1 - \alpha_2}{1 - \alpha_1} \right) \cdot t'}{\ln \left(\frac{\alpha_1(1 - \alpha_2)}{\alpha_2(1 - \alpha_1)} \right)}. \quad (11)$$

But $(c_1^y, \dots, c_{t'}^y) \in \bar{C}_{t+1,t'}^y$ implies that $c_1^y + \dots + c_{t'-1}^y < \lfloor \lambda_y t' - \gamma(\phi_{t+1}) \rfloor \leq \lambda_y t'$, so that the inequality in (11) holds for all $(c_1^y, \dots, c_{t'-1}^y) \in \bar{C}_{t+1,t'}^y$. Since $\bar{C}_{t+1,t'}^y$ might be empty (for $\gamma(\phi_{t+1})$ high, for example), it follows $\psi_{t'}^y(\bar{\theta}) - \psi_{t'}^y(\underline{\theta}) \leq 0$. This in turn implies

$$r_{t+1,t'}^y(\bar{\theta}) - r_{t+1,t'+1}^y(\bar{\theta}) \leq r_{t+1,t'}^y(\underline{\theta}) - r_{t+1,t'+1}^y(\underline{\theta}).$$

Observe that, by construction, $r_{t+1,1}^y(\underline{\theta}) = r_{t+1,1}^y(\bar{\theta}) = 1$, together with $r_{t+1,1}^y(\bar{\theta}) - r_{t+1,2}^y(\bar{\theta}) \leq r_{t+1,1}^y(\underline{\theta}) - r_{t+1,2}^y(\underline{\theta})$, implies $r_{t+1,2}^y(\underline{\theta}) \leq r_{t+1,2}^y(\bar{\theta})$. Given

$$r_{t+1,2}^y(\bar{\theta}) - r_{t+1,3}^y(\bar{\theta}) \leq r_{t+1,2}^y(\underline{\theta}) - r_{t+1,3}^y(\underline{\theta}) \leq r_{t+1,2}^y(\bar{\theta}) - r_{t+1,3}^y(\underline{\theta})$$

it also follows $r_{t+1,3}^y(\underline{\theta}) \leq r_{t+1,3}^y(\bar{\theta})$. Proceeding similarly, we obtain $r_{t+1,t'}^y(\underline{\theta}) \leq r_{t+1,t'}^y(\bar{\theta})$ for all $t' \geq 1$. As before, this inequality becomes strict for t' sufficiently large. Let \bar{t} be the smallest integer such that $c_1^y + \dots + c_{\bar{t}}^y < \lfloor \lambda_y \bar{t} - \gamma(\phi_{t+1}) \rfloor$, i.e., $\lfloor \lambda_y \bar{t} - \gamma(\phi_{t+1}) \rfloor > 0$. Since $\gamma(\phi_{t+1})$ is finite for any possible $\phi_{t+1} \in [\phi^{**}, 1)$, such \bar{t} always exists. But this implies that, for $t < \bar{t}$, $c_1^y + \dots + c_t^y \geq \lfloor \lambda_y t - \gamma(\phi_{t+1}) \rfloor$, so that $C_{t+1,\bar{t}-1}^y$ is non-empty, and, for $t = \bar{t}$, $c_1^y + \dots + c_t^y < \lfloor \lambda_y \bar{t} - \gamma(\phi_{t+1}) \rfloor$, from which it follows that $\bar{C}_{t+1,\bar{t}}^y$ is non-empty either. Therefore, $\bar{C}_{t+1,t'}^y$ is non-empty for all $t' > \bar{t}$. (ii) (iii) are straightforward modifications of the argument in (i). \square

Proof of Proposition 7: (i) By definition

$$\begin{aligned} C_{t-1}^z &= \{(c_1^z, \dots, c_{t-1}^z) : \lfloor \lambda_z \tau - \underline{\gamma} \rfloor - c_1^z - \dots - c_{\tau-1}^z \leq c_\tau^z < \lfloor \lambda_z \tau + \bar{\gamma} \rfloor - c_1^z - \dots - c_{\tau-1}^z, \\ &\quad \tau = 1, \dots, t-1\} \\ &= \left\{ (c_1^z, \dots, c_{t-1}^z) : \frac{\lfloor \lambda_z \tau - \underline{\gamma} \rfloor}{\tau} \leq \frac{c_1^z + \dots + c_\tau^z}{\tau} < \frac{\lfloor \lambda_z \tau + \bar{\gamma} \rfloor}{\tau}, \tau = 1, \dots, t-1 \right\}. \end{aligned}$$

Notice that $\lim_{t \rightarrow \infty} \left(\frac{\lfloor \lambda_z \tau + \bar{\gamma} \rfloor}{\tau} - \frac{\lfloor \lambda_z \tau - \underline{\gamma} \rfloor}{\tau} \right) = 0$, which implies that, as t grows arbitrarily large, the set C_{t-1}^z becomes empty. This, in turn, yields

$$\lim_{t \rightarrow \infty} r_t^z(\underline{\theta}) = \lim_{t \rightarrow \infty} r_t^z(\bar{\theta}) = 0.$$

(ii) Recall that, for $p(\bar{\theta}) = \alpha_1$, $p(\underline{\theta}) = \alpha_2$ and $t \geq 1, t' \geq 2$,

$$\begin{aligned} r_{t+1, t'}^y(\theta) &= \sum_{(c_1^y, \dots, c_{t'-1}^y) \in C_{t'-1}^y} p(\theta)^{c_1^y + \dots + c_{t'-1}^y} (1 - p(\theta))^{t'-1 - c_1^y - \dots - c_{t'-1}^y} \\ C_{t+1, t'-1}^y &= \{(c_1^y, \dots, c_{t'-1}^y) : c_\tau^y \geq \lfloor \lambda_y \tau - \gamma(\phi_{t+1}) \rfloor - c_1^y - \dots - c_{\tau-1}^y, \text{ for } \tau = 1, \dots, t'-1\} \end{aligned}$$

implies that

$$r_{t+1, t'}^y(\underline{\theta}) \leq \Pr\{X_{t'-1}^y(\underline{\theta}) \geq \lfloor \lambda_y(t-1) - \gamma(\phi_{t+1}) \rfloor\}$$

since $C_{t+1, t'-1}^y \subset \{(c_1^y, \dots, c_{t'-1}^y) : c_1^y + \dots + c_{t'-1}^y \geq \lfloor \lambda_y(t'-1) - \gamma(\phi_{t+1}) \rfloor\}$, where $X_{t'-1}^y$ denotes the number of successes in $t'-1$ Bernoulli trials when the probability of success is α_2 . By the Law of Large Numbers for the Bernoulli distribution (see Feller [1965], Ch. 6), we know that, for all $\varepsilon > 0$, $\lim_{t' \rightarrow \infty} \Pr\{X_{t'-1}^y(\underline{\theta}) \geq (t'-1)(\alpha_2 + \varepsilon)\} = 0$. By the same argument as in Proposition 4 under (i), it is possible to conclude that $r_{t+1, t'}^y(\underline{\theta}) \rightarrow 0$ as $t' \rightarrow \infty$, while $\lim_{t' \rightarrow \infty} r_{t+1, t'}^y(\bar{\theta}) > 0$. As before, this limiting probability is bounded away from one. For the rest of the claim, observe that

$$\lim_{t \rightarrow \infty} \Pr\left\{ \frac{c_1 + \dots + c_{t-1}}{t-1} \geq \lambda - \frac{\gamma}{t-1} \mid \bar{\theta} \right\} = \lim_{t' \rightarrow \infty} \Pr\left\{ \frac{c_1^y + \dots + c_{t'-1}^y}{t'-1} \geq \lambda_y - \frac{\gamma(\phi_{t+1})}{t'-1} \mid \bar{\theta} \right\} = 1$$

from which, together with Proposition 4 under (ii), it follows $\lim_{t \rightarrow \infty} \rho_t(\bar{\theta}) \leq \lim_{t' \rightarrow \infty} r_{t+1, t'}^y(\bar{\theta})$

(CONTINUE). (iii) Since $p_{t'+1, t'}^y(\theta) = \frac{r_{t'+1, t'+1}^y(\theta)}{r_{t'+1, t'}^y(\theta)}$, it follows from the argument under (2) that

$$\lim_{t' \rightarrow \infty} p_{t'+1, t'}^y(\bar{\theta}) = \lim_{t' \rightarrow \infty} \frac{r_{t'+1, t'+1}^y(\bar{\theta})}{r_{t'+1, t'}^y(\bar{\theta})} \equiv \frac{\bar{r}}{\bar{r}} = 1$$

as desired. □

Proof of Proposition 8: By inspection of $V(\phi, z)$ and $V(\phi, y)$, it is immediate that, for $c_y, c_z > 0$, $\phi^*(z) < \phi^*(y)$ and $\phi^{**}(y) < \phi^{**}(z)$. If $y(\bar{\theta}) - c_y > \bar{\Pi} + \bar{U} > z(\underline{\theta}) > y(\underline{\theta})$ and $\delta \geq [\bar{\Pi} + \bar{U} - z(\bar{\theta}) + c_z]/[y(\bar{\theta}) - z(\bar{\theta})c_z - c_y]$, it follows $V_y(1, y) > V_y(1, z) > V_z(1, z) > V_z(1, y) > \bar{\Pi} > V_z(0, z) > V_y(0, z)$. Also, $\bar{\Pi} > V_z(0, y) > V_y(0, y)$. By monotonicity of $V(\phi, z)$ and the fact that $V_z(1, z) > \bar{\Pi} > V_z(0, z)$, it follows that there exists a unique value of the posterior, $\phi^*(z)$, which makes the firm indifferent between employing the worker at z and not hiring him. Similarly, by monotonicity of $V(\phi, y)$ and the fact that $V_z(1, y) > \bar{\Pi} > V_z(0, y)$, it follows that there exists a unique value of the posterior, $\phi^*(y)$, for which employing the worker at y is just as profitable as the firm's outside option. To prove that there exists a range of posterior beliefs for which task y is strictly preferred to task z , for each possible incumbent arm, requires showing

$$\begin{aligned}\bar{\Pi} &= (1 - \delta)[z(\phi^*(z)) - \bar{U}] + \delta E_z V(\phi^*(z), z) > (1 - \delta)[y(\phi^*(z)) - \bar{U} - c_y] + \delta E_y V(\phi^*(z), y) \\ \bar{\Pi} &= (1 - \delta)[z(\phi^*(y)) - \bar{U} - c_z] + \delta E_z V(\phi^*(y), z) > (1 - \delta)[y(\phi^*(y)) - \bar{U}] + \delta E_y V(\phi^*(y), y).\end{aligned}$$

Notice that, if there exist values of the discount factor for which the inequality holds at $\phi^*(y)$, there must exist values for which it holds at $\phi^*(z)$ as well. Consider then $\phi^*(y)$. At $\phi^*(y)$ either z is more profitable than y or viceversa. In the first case (as well as in the case in which the firm is indifferent between the two tasks), $V(\phi, y)$ has a kink at $\phi^*(y)$. Then, it is strictly convex at $\phi^*(y)$. In case, instead, y is strictly preferred to z at $\phi^*(y)$, it follows that $V(\phi^*(y), y) = V_y(\phi^*(y), y)$. Since V_y is strictly increasing in ϕ , $V_y(\varphi(\phi^*(y) | \bar{y}), y) > V_y(\varphi(\phi^*(y) | \underline{y}), y)$, since $\phi^*(y) \in (0, 1)$. Therefore, even in this case $V(\phi, y)$ is strictly convex at $\phi^*(y)$. This implies that, for fixed probability of success at either task, it is always possible to find a value of δ sufficiently large such that at either $\phi^*(z)$ and $\phi^*(y)$ task z is strictly preferred to task y . Thus, $\phi^*(y) < \phi^{**}(y)$ and $\phi^*(z) < \phi^{**}(z)$ and, by monotonicity of the differences $V_y(\phi, z) - V_z(\phi, z)$ and $V_y(\phi, y) - V_z(\phi, y)$ in ϕ , these cut-offs values are uniquely determined. (i) Since $c_z > 0$ and $V(\phi, y) \leq V(\phi)$, it is immediate that $\phi^* < \phi^*(y)$. Given $V(\phi, z) \leq V(\phi)$, it also follows $\phi^* \leq \phi^*(z)$. Due to the monotonicity of $V_y(\phi, z) - V_z(\phi, z)$ and $V_y(\phi, y) - V_z(\phi, y)$ in ϕ , $\phi^{**}(y) < \phi^{**}$ if and only if at ϕ^{**}

$$(1 - \delta)y(\phi^{**}) + \delta E_y V(\phi^{**}, y) > (1 - \delta)[z(\phi^{**}) - c_z] + \delta E_z V(\phi^{**}, z).$$

while $\phi^{**} < \phi^{**}(z)$ if and only if at ϕ^{**}

$$(1 - \delta)z(\phi^{**}) + \delta E_z V(\phi^{**}, z) > (1 - \delta)[y(\phi^{**}) - c_y] + \delta E_y V(\phi^{**}, y).$$

Then, to show $\phi^{**}(y) < \phi^{**} < \phi^{**}(z)$ is equivalent to prove

$$(1 - \delta)[y(\phi^{**}) - z(\phi^{**}) + c_z] > \delta E_z V(\phi^{**}, z) - \delta E_y V(\phi^{**}, y) > (1 - \delta)[y(\phi^{**}) - c_y - z(\phi^{**})]. \quad (12)$$

By definition of ϕ^{**} , $\delta E_z V(\phi^{**}) - \delta E_y V(\phi^{**}) = (1 - \delta)[y(\phi^{**}) - z(\phi^{**})]$, where $E_x V(\cdot)$ denotes the expected continuation value at task $x \in \{z, y\}$ for the problem without switching costs. If $c_y > y(\bar{\theta}) - z(\bar{\theta})$, then $y(\phi^{**}) - c_y - z(\phi^{**}) < 0$. By convexity of $V(\phi, z)$ and $V(\phi, y)$ and the fact that the conditional distribution of the updated posterior at z is a mean-preserving spread of the corresponding distribution at y (ADD CONDITIONS), the second inequality in (12) holds. Observe that, if c_y is sufficiently bigger than c_z , $V(\phi, y) > V(\phi, z)$, which in turn implies $E_z V(\phi) - E_z V(\phi, z) \geq E_y V(\phi) - E_y V(\phi, y)$. As a result, $\delta E_z V(\phi^{**}) - \delta E_y V(\phi^{**}) \geq \delta E_z V(\phi^{**}, z) - \delta E_y V(\phi^{**}, y)$, and thus, for any $c_z > 0$, the first inequality in (12) holds as well. (ii) For z not to be profitable in equilibrium once y is chosen, it must be

$$\bar{\Pi} = (1 - \delta)[y(\phi^*(y)) - \bar{U}] + \delta E_y V(\phi^*(y), y) \geq (1 - \delta)[z(\phi^*(y)) - \bar{U} - c_z] + \delta E_z V(\phi^*(y), z).$$

or

$$(1 - \delta)c_z \geq (1 - \delta)[z(\phi^*(y)) - y(\phi^*(y))] + \delta E_z V(\phi^*(y), z) - \delta E_y V(\phi^*(y), y). \quad (13)$$

Notice that, by definition of $\phi^*(y)$ and the fact that $V_z(\phi^*(y)) \geq V_z(\phi^*(y), z)$,

$$V_z(\phi^*(y)) - \bar{\Pi} \geq (1 - \delta)[z(\phi^*(y)) - y(\phi^*(y))] + \delta E_z V(\phi^*(y), z) - \delta E_y V(\phi^*(y), y).$$

Moreover, since $\phi^* < \phi^*(y)$, by monotonicity of $V_y(\phi) - V_z(\phi)$ and the definition of ϕ^* ,

$$(1 - \delta)[\bar{\Pi} + \bar{U} - y(\underline{\theta})] \geq \bar{\Pi} - (1 - \delta)[y(\phi^*) - \bar{U}] - \delta \bar{\Pi} \geq \bar{\Pi} - V_y(\phi^*) \geq V_z(\phi^*(y)) - \bar{\Pi}$$

a sufficient condition for (13) to hold is $c_z \geq \bar{\Pi} + \bar{U} - y(\underline{\theta})$. \square

Proof of Lemma 6: Since $E_z V(\phi^1, \phi^2)$ is an increasing function of ϕ^1 , while the flow profit from z is a strictly increasing function of ϕ^1 , it follows that, for each ϕ^2 , there exists a unique value of ϕ^1 for which (5) is satisfied. Given that $\bar{\Pi} > z(\bar{\theta}, \bar{\theta}) = z(\bar{\theta}, \underline{\theta})$, it must be $\max\{\bar{\Pi}, V_y(\phi^1, \bar{z})\} > \bar{\Pi}$ for (5) to hold. The minimal value of ϕ^2 for which (5) holds is for $\phi^2 = 0$, with corresponding $\phi^1 = \phi_z^*(0)$. Since $V_z(1, 0) = (1 - \delta)z(\bar{\theta}, \bar{\theta}) + \delta y(\bar{\theta}, \bar{\theta}) > \bar{\Pi}$ for $\delta > \underline{\delta}$, it follows $\phi_z^*(0) \in (0, 1)$. Similarly, given that $V_z(0, 1) = (1 - \delta)z(\bar{\theta}, \underline{\theta}) + \delta \bar{\Pi} < \bar{\Pi}$, the maximal value of ϕ^2 for (5) to hold, $\bar{\phi}^2$, must be strictly bounded away from 1. Note that V_z is a strictly increasing function of ϕ^1 and ϕ^2 . Then, as ϕ^2 increases, ϕ^1 must decrease for (5) to be satisfied, so $\phi_z^*(\bar{\phi}^2) < \phi_z^*(0)$. Since the maximal increase in ϕ^2 consistent with (5) is for $\phi^3 = 0$, it follows $\phi_z^*(\bar{\phi}^2) = 1 - \bar{\phi}^2 \in (0, 1)$. Observe that, since

$$\begin{aligned} z(\phi^1, \phi^3) &= z(\bar{\theta}, \underline{\theta}) - \bar{U} + \phi^1[z(\bar{\theta}, \bar{\theta}) - z(\bar{\theta}, \underline{\theta})] - \phi^3[z(\bar{\theta}, \underline{\theta}) - z(\underline{\theta}, \underline{\theta})] \\ y(\phi^1, \phi^3) &= y(\bar{\theta}, \underline{\theta}) - \bar{U} + \phi^1[y(\bar{\theta}, \bar{\theta}) - y(\bar{\theta}, \underline{\theta})] - \phi^3[y(\bar{\theta}, \underline{\theta}) - y(\underline{\theta}, \underline{\theta})] \end{aligned}$$

and $z(\bar{\theta}, \underline{\theta}) > z(\underline{\theta}, \underline{\theta})$ and $y(\bar{\theta}, \underline{\theta}) > y(\underline{\theta}, \underline{\theta})$, V is decreasing in ϕ^3 . In particular, V_z is a strictly decreasing function of ϕ^3 . Therefore, the same argument holds for $\bar{\Pi} = V(\phi^1, \phi^3)$, with the difference that $\phi_z^*(\bar{\phi}^3) > \phi_z^*(0)$. \square

Proof of Lemma 7: (i) This amounts to showing

$$\begin{aligned} \bar{\Pi} &= (1 - \delta)[\phi_z^*(\phi^2)z(\bar{\theta}, \bar{\theta}) + \phi^2z(\bar{\theta}, \underline{\theta}) + (1 - \phi_z^*(\phi^2) - \phi^2)z(\underline{\theta}, \underline{\theta})] + \delta E_z V(\phi_z^*(\phi^2), \phi^2) \\ &> (1 - \delta)[\phi_z^*(\phi^2)y(\bar{\theta}, \bar{\theta}) + \phi^2y(\bar{\theta}, \underline{\theta}) + (1 - \phi_z^*(\phi^2) - \phi^2)y(\underline{\theta}, \underline{\theta})] + \delta E_y V(\phi_z^*(\phi^2), \phi^2) \end{aligned} \quad (14)$$

for any given ϕ^2 . Notice that V is strictly convex at $(\phi_z^*(\phi^2), \phi^2)$. This follows from considering three cases. At $(\phi_z^*(\phi^2), \phi^2)$ either V_z is greater than V_y or viceversa. Suppose $V_z(\phi_z^*(\phi^2), \phi^2) \geq V_y(\phi_z^*(\phi^2), \phi^2)$. Then, since at $(\phi_z^*(\phi^2), \phi^2)$ the value function has a kink, it must be strictly convex. Consider now the case $V_z(\phi_z^*(\phi^2), \phi^2) < V_y(\phi_z^*(\phi^2), \phi^2)$. This implies that, at $(\phi_z^*(\phi^2), \phi^2)$, $V(\phi_z^*(\phi^2), \phi^2) = V_y(\phi_z^*(\phi^2), \phi^2)$ and it is strictly greater than $\bar{\Pi}$. But then, since $V_y(\varphi(\phi_z^*(\phi^2) | \bar{y}), \varphi(\phi^2 | \bar{y})) > V_y(\varphi(\phi_z^*(\phi^2) | \underline{y}), \varphi(\phi^2 | \underline{y}))$, by strict monotonicity of V_y , it must be $EV(\phi_z^*(\phi^2), \phi^2) > V(\phi_z^*(\phi^2), \phi^2)$. By strict convexity of V at $(\phi_z^*(\phi^2), \phi^2)$ and the assumptions on the informativeness of the signals at z and y , i.e., by Lemma 5, with V convex

$$\begin{aligned} EV_z(\phi^1, \phi^2) &\equiv EV(\varphi(\phi^1 | z), \varphi(\phi^2 | z)) \geq EV(\varphi(\phi^1 | y), \varphi(\phi^2 | z)) \\ &\geq EV(\varphi(\phi^1 | y), \varphi(\phi^2 | y)) \equiv EV_y(\phi^1, \phi^2) \end{aligned}$$

there always exists a value of δ sufficiently large for (14) to be satisfied. (ii) Observe that, when $\phi^1 = 1$, $V_y(1, 0) = y(\bar{\theta}, \bar{\theta}) - \bar{U} > V_z(1, 0) = (1 - \delta)[z(\bar{\theta}, \bar{\theta}) - \bar{U}] + \delta[y(\bar{\theta}, \bar{\theta}) - \bar{U}]$, while, when $\phi^2 = 1$, $V_z(0, 1) = (1 - \delta)[z(\bar{\theta}, \underline{\theta}) - \bar{U}] + \delta\bar{\Pi} > V_y(0, 1) = (1 - \delta)[y(\bar{\theta}, \underline{\theta}) - \bar{U}] + \delta\bar{\Pi}$ or, similarly, when $\phi^3 = 1$, $V_z(0, 0) = (1 - \delta)[z(\underline{\theta}, \underline{\theta}) - \bar{U}] + \delta\bar{\Pi} > V_y(0, 0) = (1 - \delta)[y(\underline{\theta}, \underline{\theta}) - \bar{U}] + \delta\bar{\Pi}$. To show that, for given ϕ^2 or ϕ^2 , there exists at most a unique value of ϕ^1 for which z and y are equally profitable, it is enough to show that the difference $V_y(\phi^1, \phi^2) - V_z(\phi^1, \phi^2)$, for fixed ϕ^2 , and the difference $V_y(\phi^1, \phi^3) - V_z(\phi^1, \phi^3)$, for fixed ϕ^3 , are increasing in ϕ^1 .²⁴ Consider first the case in which ϕ^2 is fixed. Define the operators T_y and T_z as

$$\begin{aligned} T_y f(\phi^1, \phi^2) &= (1 - \delta)[\phi^1 y(\bar{\theta}, \bar{\theta}) + \phi^2 y(\bar{\theta}, \underline{\theta}) + (1 - \phi^1 - \phi^2)y(\underline{\theta}, \underline{\theta})] + \delta E_y f(\phi^1, \phi^2) \\ T_z f(\phi^1, \phi^2) &= (1 - \delta)[\phi^1 z(\bar{\theta}, \bar{\theta}) + \phi^2 z(\bar{\theta}, \underline{\theta}) + (1 - \phi^1 - \phi^2)z(\underline{\theta}, \underline{\theta})] + \delta E_z f(\phi^1, \phi^2). \end{aligned}$$

Let $Uf = \max\{\bar{\Pi}, T_y f(\phi^1, \phi^2), T_z f(\phi^1, \phi^2)\}$. Consider the difference

$$\begin{aligned} T_y(Uf) - T_z(Uf) &= [T_y(Uf) - T_y(\max\{\bar{\Pi}, T_z f\})] \\ &\quad + [T_z(\max\{\bar{\Pi}, T_y f\}) - T_z(Uf)] + [T_y(\max\{\bar{\Pi}, T_z f\}) - T_z(\max\{\bar{\Pi}, T_y f\})] \end{aligned} \quad (15)$$

²⁴Observe that, given the restrictions on payoffs, ϕ^1 has to be sufficiently high for y to be assigned in equilibrium.

as a function of ϕ^1 . Suppose that, for any real-valued function f on $[0, 1]^2$, $T_y f - T_z f$ increasing in ϕ^1 (for fixed ϕ^2) implies that $T_y(Uf) - T_z(Uf)$ is increasing in ϕ^1 (for fixed ϕ^2). Then, V , the unique fixed point of U , must be increasing in ϕ^1 (for fixed ϕ^2) as well. Therefore, for the statement in the Lemma to hold true it is enough to show that each term in (15) is increasing in ϕ^1 . Notice that the difference in the flow profit from y and z , as a function of ϕ^1 , is given by

$$\begin{aligned} y(\phi^1, \phi^2) - z(\phi^1, \phi^2) &\equiv \phi^1[y(\bar{\theta}, \bar{\theta}) - y(\underline{\theta}, \underline{\theta}) - z(\bar{\theta}, \bar{\theta}) + z(\underline{\theta}, \underline{\theta})] \\ &\quad + \phi^2[y(\bar{\theta}, \underline{\theta}) - y(\underline{\theta}, \underline{\theta}) - z(\bar{\theta}, \underline{\theta}) + z(\underline{\theta}, \underline{\theta})] + y(\underline{\theta}, \underline{\theta}) - z(\underline{\theta}, \underline{\theta}) \end{aligned}$$

and it is monotone increasing in ϕ^1 , for given ϕ^2 , if $y(\bar{\theta}, \bar{\theta}) - y(\underline{\theta}, \underline{\theta}) > z(\bar{\theta}, \bar{\theta}) - z(\underline{\theta}, \underline{\theta})$. Similarly, for given ϕ^3 , the difference in flow profit from y and z is

$$\begin{aligned} y(\phi^1, \phi^3) - z(\phi^1, \phi^3) &\equiv \phi^1[y(\bar{\theta}, \bar{\theta}) - y(\bar{\theta}, \underline{\theta}) - z(\bar{\theta}, \bar{\theta}) + z(\bar{\theta}, \underline{\theta})] \\ &\quad - \phi^3[y(\bar{\theta}, \underline{\theta}) - y(\underline{\theta}, \underline{\theta}) - z(\bar{\theta}, \underline{\theta}) + z(\underline{\theta}, \underline{\theta})] + y(\bar{\theta}, \underline{\theta}) - z(\bar{\theta}, \underline{\theta}) \end{aligned}$$

which is monotone increasing in ϕ^1 , for fixed ϕ^3 , if $y(\bar{\theta}, \bar{\theta}) - y(\bar{\theta}, \underline{\theta}) > z(\bar{\theta}, \bar{\theta}) - z(\bar{\theta}, \underline{\theta})$. Since $\alpha_2 > \alpha_3$, then $y(\bar{\theta}, \underline{\theta}) > y(\underline{\theta}, \underline{\theta})$, while $\beta_2 > \beta_3$ implies $z(\bar{\theta}, \underline{\theta}) > z(\underline{\theta}, \underline{\theta})$. Then, a sufficient condition for $y(\bar{\theta}, \bar{\theta}) - y(\underline{\theta}, \underline{\theta}) > z(\bar{\theta}, \bar{\theta}) - z(\underline{\theta}, \underline{\theta})$ and $y(\bar{\theta}, \bar{\theta}) - y(\bar{\theta}, \underline{\theta}) > z(\bar{\theta}, \bar{\theta}) - z(\bar{\theta}, \underline{\theta})$ to be satisfied is $y(\bar{\theta}, \bar{\theta}) - y(\bar{\theta}, \underline{\theta}) > z(\bar{\theta}, \bar{\theta}) - z(\underline{\theta}, \underline{\theta})$. The rest of the argument is an immediate extension of the one already used in the baseline case to prove that the difference $V_y(\phi) - V_z(\phi)$ is increasing in ϕ . \square

Proof of Lemma 9: Since

$$E_z V(\phi^1, \phi^2) = (\phi^1 + \phi^2) \max \left\{ \bar{\Pi}, \tilde{V}_y \left(\frac{\phi^1}{\phi^1 + \phi^2} \right) \right\} + (1 - \phi^1 - \phi^2) \bar{\Pi}$$

is an increasing function of ϕ^1 , while the flow profit from z is a strictly increasing function of ϕ^1 , it follows that, for each ϕ^2 , there exists a unique value of ϕ^1 for which (5) is satisfied. Given that $\bar{\Pi} > z(\bar{\theta}, \bar{\theta}) = z(\bar{\theta}, \underline{\theta})$, it must be $\max\{\bar{\Pi}, \tilde{V}_y(\tilde{\phi}^1)\} > \bar{\Pi}$ for (5) to hold. Since $V_z(1, 0) = (1 - \delta)z(\bar{\theta}, \bar{\theta}) + \delta y(\bar{\theta}, \bar{\theta}) > \bar{\Pi}$ for $\delta > \frac{\bar{\Pi} - z(\bar{\theta}, \bar{\theta})}{y(\bar{\theta}, \bar{\theta}) - z(\bar{\theta}, \bar{\theta})}$, the minimal value of ϕ^2 for (5) to hold is $\phi^2 = 0$ and it is for $\phi^1 = \phi_z^*(0) \in (0, 1)$. Similarly, given that $V_z(0, 1) = (1 - \delta)z(\bar{\theta}, \underline{\theta}) + \delta \bar{\Pi} < \bar{\Pi}$, the maximal value of ϕ^2 for (5) to hold, $\bar{\phi}^2$, must be strictly bounded away from 1. Note that V_z is a strictly increasing function of ϕ^1 and, for δ sufficiently large, strictly *decreasing* of ϕ^2 . Consider

$$\frac{\partial V_z(\phi^1, \phi^2)}{\partial \phi^2} = (1 - \delta)[z(\bar{\theta}, \bar{\theta}) - z(\underline{\theta}, \underline{\theta})] + \delta \tilde{V}_y \left(\frac{\phi^1}{\phi^1 + \phi^2} \right) - \delta \left\{ \frac{\phi^1}{\phi^1 + \phi^2} \tilde{V}'_y \left(\frac{\phi^1}{\phi^1 + \phi^2} \right) + \bar{\Pi} \right\}.$$

Since $\tilde{V}_y(\tilde{\phi}^1)$ is continuously differentiable to the right of ϕ_y^* and, as noted, $\tilde{V}_y(\tilde{\phi}^1) > \bar{\Pi}$, at any such $\tilde{\phi}^1$ for which $\bar{\Pi}$ and z are equally profitable, it follows $\tilde{V}_y(x + z) \geq \tilde{V}_y(x) + \tilde{V}'_y(x) \cdot z$, by convexity

of \tilde{V}_y . Let $x = \phi^1/(\phi^1 + \phi^2)$ and $z = -\varepsilon$, where $\varepsilon = \phi^1/(\phi^1 + \phi^2) - \phi_y^* > 0$. Then,

$$\tilde{V}_y \left(\frac{\phi^1}{\phi^1 + \phi^2} - \varepsilon \right) + \varepsilon \tilde{V}'_y \left(\frac{\phi^1}{\phi^1 + \phi^2} \right) \geq \tilde{V}_y \left(\frac{\phi^1}{\phi^1 + \phi^2} \right)$$

which implies

$$\frac{\phi^1}{\phi^1 + \phi^2} \tilde{V}'_y \left(\frac{\phi^1}{\phi^1 + \phi^2} \right) + \bar{\Pi} > \left(\frac{\phi^1}{\phi^1 + \phi^2} - \phi_y^* \right) \tilde{V}'_y \left(\frac{\phi^1}{\phi^1 + \phi^2} \right) + \tilde{V}_y(\phi_y^*) \geq \tilde{V}_y \left(\frac{\phi^1}{\phi^1 + \phi^2} \right)$$

since \tilde{V}_y is strictly increasing in $\tilde{\phi}^1$ and $\tilde{V}_y(\phi_y^*) = \bar{\Pi}$, by definition of ϕ_y^* , so that $\partial V_z(\phi^1, \phi^2)/\partial \phi^2 < 0$ if δ is sufficiently large. Then, as ϕ^2 increases, ϕ^1 must increase for (5) to be satisfied, so that $\phi_z^*(\bar{\phi}^2) > \phi_z^*(0)$. Since the maximal increase in ϕ^2 consistent with (5) is for $\phi^3 = 0$, it follows $\phi_z^*(\bar{\phi}^2) = 1 - \bar{\phi}^2 \in (0, 1)$. The rest of the argument follows as in the proof of Lemma 6. \square

Proof of Proposition 11: (i) Denote by V^i and V^j , respectively, the value function of the firm employing the worker (firm i) and not employing him (firm j) at any given state ϕ . Notice that, if the worker's human capital is general, i.e., fully transferable across firms, both firms share a common belief $\phi = \phi^i = \phi^j$, while, if the worker's human capital is firm-specific, the state is the common belief vector $\phi = (\phi^i, \phi^j)$. Let V^w be the worker's value function. From BV's Lemma 2 (CHECK POSITIVE EXPECTATION) it is immediate that a wage and job assignment strategy for each firm and a worker's acceptance behavior form an MPE in stationary strategies if and only if

$$(1 - \delta)w^i(\phi) + \delta E_{x^i} V^w(\phi) = (1 - \delta)w^j(\phi) + \delta E_{x^j} V^w(\phi) \quad (16)$$

$$(1 - \delta)[x^i(\phi) - w^i(\phi)] + \delta E_{x^i} V^i(\phi) \geq \delta E_{x^j} V^i(\phi) \quad (17)$$

$$\delta E_{x^i} V^j(\phi) \geq (1 - \delta)[x^j(\phi) - w^j(\phi)] + \delta E_{x^j} V^j(\phi) \quad (18)$$

$$x^i(\phi) = \operatorname{argmax}\{(1 - \delta)z^i(\phi^i) + \delta E_{z^i} V^i(\phi), (1 - \delta)y^i(\phi^i) + \delta E_{y^i} V^i(\phi)\} - w^i(\phi) \quad (19)$$

$$x^j(\phi) = \operatorname{argmax}\{(1 - \delta)z^j(\phi^j) + \delta E_{z^j} V^j(\phi), (1 - \delta)y^j(\phi^j) + \delta E_{y^j} V^j(\phi)\} - w^j(\phi). \quad (20)$$

Suppose now that the planner selects an employment policy (i.e., the firm employing the worker at each state) so as to maximize $V^w(\phi) + V^i(\phi) + V^j(\phi)$, say program (P1), and an assignment policy so as to maximize $V^w(\phi) + p^i(\phi)V^i(\phi) + p^j(\phi)V^j(\phi)$, say (P2), where $p^i(\phi)$ and $p^j(\phi)$ indicate, respectively, the equilibrium probability that the worker is employed at firm i and j . Express (19) and (20) as (weak) inequalities. By summing (16), (17) and (18) side by side it is immediate that firm i is selected in equilibrium only if it is selected along some optimal path in the planner's problem (P1). By summing (16), (19) expressed as weak inequality, weighted by $p^i(\phi)$, and (20) expressed as weak inequality, weighted by $p^j(\phi)$, side by side it also follows that the task proposed by each firm is an equilibrium assignment only if it is selected along some optimal path in the

planner's problem (P2), i.e., as before the task the worker is assigned maximizes the surplus of the match. However, the equilibrium task is a solution to program (P1) if and only if the conditions in the statement are satisfied, as it can be easily seen by noticing that

$$\begin{aligned} & (1 - \delta)z^i(\phi^i) + \delta E_{z^i} V^i(\phi) - (1 - \delta)y^i(\phi^i) + \delta E_{y^i} V^i(\phi) \geq 0 \\ \iff & (1 - \delta)z^j(\phi^j) + \delta E_{z^j} V^j(\phi) - (1 - \delta)y^j(\phi^j) + \delta E_{y^j} V^j(\phi) \geq 0 \end{aligned}$$

holds if and only if $\alpha_x^j = \alpha_x^i$ and $\beta_x^j = \beta_x^i$, for $x \in \{z, y\}$, which ensure $E_{z^i} V^i \geq E_{y^i} V^i \iff E_{z^j} V^j \geq E_{y^j} V^j$, and $z^j(\theta) = az^i(\theta) + b$, $y^j(\theta) = ay^i(\theta) + b$, with $a > 0$, at each θ , guaranteed by $z^j = az^i + b$ and $y^j = ay^i + b$, for each $z^i \in \{\bar{z}, \underline{z}\}$, $y^i \in \{\bar{y}, \underline{y}\}$.

(ii) The argument closely follows the proof of Theorem 3 in BV. Since the equality in (16) has to hold in each period, it can be extended using any of the alternatives, either i or j . Extend both left-hand side and right-hand side of (16) over the continuation game in which firm j employs the worker forever. Then,

$$(1 - \delta)w^i(\phi) + (1 - \delta)E \left\{ \sum_{s=t+1}^{\infty} \delta^{s-t} w^j(\tilde{\phi}_s) \right\} = (1 - \delta)E \left\{ \sum_{s=t}^{\infty} \delta^{s-t} w^j(\phi_s) \right\}.$$

Note the different decoration of ϕ in the two expressions to emphasize the fact that the associated experimentation paths are different. Solve the above expression for $w^i(\phi)$

$$w^i(\phi) = E \left\{ \sum_{s=t}^{\infty} \delta^{s-t} w^j(\phi_s) \right\} - E \left\{ \sum_{s=t+1}^{\infty} \delta^{s-t} w^j(\tilde{\phi}_s) \right\}. \quad (21)$$

Recall that the condition characterizing a cautious equilibrium is that (18) holds as an equality. Extend (18) in the same way as (16),

$$(1 - \delta)E \left\{ \sum_{s=t+1}^{\infty} \delta^{s-t} [x^j(\tilde{\phi}_s) - w^j(\tilde{\phi}_s)] \right\} = (1 - \delta)E \left\{ \sum_{s=t}^{\infty} \delta^{s-t} [x^j(\phi_s) - w^j(\phi_s)] \right\}$$

or, equivalently,

$$E \left\{ \sum_{s=t}^{\infty} \delta^{s-t} w^j(\phi_s) \right\} = E \left\{ \sum_{s=t}^{\infty} \delta^{s-t} x^j(\phi_s) \right\} - E \left\{ \sum_{s=t+1}^{\infty} \delta^{s-t} [x^j(\tilde{\phi}_s) - w^j(\tilde{\phi}_s)] \right\}.$$

We can then use this expression above to rewrite (21) as

$$\begin{aligned} w^i(\phi) &= E \left\{ \sum_{s=t}^{\infty} \delta^{s-t} x^j(\phi_s) \right\} - E \left\{ \sum_{s=t+1}^{\infty} \delta^{s-t} x^j(\tilde{\phi}_s) \right\} \\ &= x^j(\phi) + E \left\{ \sum_{s=t+1}^{\infty} \delta^{s-t} x^j(\phi_s) \right\} - E \left\{ \sum_{s=t+1}^{\infty} \delta^{s-t} x^j(\tilde{\phi}_s) \right\} = x^j(\phi) \end{aligned}$$

by the linearity of the expectation operator, (CHECK DCT) and the law of iterated expectations. Finally, to show that the wage policy of the non-employing firm is a supermartingale, consider again the cautious equilibrium condition (18), holding as equality. Notice that, conditional on x_t^j being observed in period t , the condition might turn into an inequality

$$\delta(1 - \delta)E[x^j(\varphi(\phi | x^j)) - w^j(\varphi(\phi | x^j))] + \delta^2 E_{x^j, x^j} V^j(\phi) \geq \delta^2 E_{x^i, x^j} V^j(\phi)$$

which can be rewritten, by adding side by side the above and (18), holding as equality (and extended over the continuation game in which firm j employs the worker forever), as

$$\begin{aligned} \delta(1 - \delta)E[x^j(\varphi(\phi | x^j)) - w^j(\varphi(\phi | x^j))] + \delta^2 E_{x^j, x^j} V^j(\phi) + \delta E_{x^i} V^j(\phi) \geq \\ \delta^2 E_{x^i, x^j} V^j(\phi) + (1 - \delta)[x^j(\phi) - w^j(\phi)] + \delta E_{x^j} V^j(\phi) \end{aligned}$$

or, equivalently,

$$\begin{aligned} \delta(1 - \delta)E[x^j(\varphi(\phi | x^j)) - w^j(\varphi(\phi | x^j))] + \delta^2 E_{x^j, x^j} V^j(\phi) \\ + \delta(1 - \delta)E[x^j(\varphi(\phi | x^i)) - w^j(\varphi(\phi | x^i))] + \delta^2 E_{x^i, x^j} V^j(\phi) \geq \\ \delta^2 E_{x^i, x^j} V^j(\phi) + (1 - \delta)[x^j(\phi) - w^j(\phi)] \\ + \delta(1 - \delta)E[x^j(\varphi(\phi | x^j)) - w^j(\varphi(\phi | x^j))] + \delta^2 E_{x^j, x^j} V^j(\phi) \end{aligned}$$

which simplifies to

$$\delta E[x^j(\varphi(\phi | x^i)) - w^j(\varphi(\phi | x^i))] \geq x^j(\phi) - w^j(\phi)$$

implying $w^j(\phi) \geq \delta E w^j(\varphi(\phi | x^i))$ as desired. \square

References

- [1] Aghion, P., P. Bolton, C. Harris and B. Jullien (1991): “Optimal Learning by Experimentation”, *Econometrica* 58(4), 621-654.
- [2] Anderson, P.M., and B.D. Meyer (1994): “The Extent and Consequences of Job Turnover”, *Brookings Papers on Economic Activity: Microeconomics*, Vol. 1994, 177-236.
- [3] Araujo, L., and B.M. Camargo (2002): “Monetary Equilibrium with Decentralized Trade and Learning”, University of Pennsylvania, mimeo.
- [4] Baker, G., M. Gibbs and B. Holmström (1994a): “The Internal Economics of the Firm: Evidence from Personnel Data”, *Quarterly Journal of Economics* 109(4), 881-919.

- [5] Banks, J.S., and R.K. Sundaram (1994): “Switching Costs and The Gittins Index”, *Econometrica* 62(3), 687-694.
- [6] Banks, J.S., and R.K. Sundaram (1992a): “Denumerable-Armed Bandits”, *Econometrica* 60(5), 1071-1096.
- [7] Banks, J.S., and R.K. Sundaram (1992b): “A Class of Bandit Problems Yielding Myopic Optimal Strategies”, *Journal of Applied Probability* 29(3), 625-632.
- [8] Bergemann and Välimäki (1996): “Learning and Strategic Pricing”, *Econometrica* 64(5), 1125-1149.
- [9] Davis, S.J. (1997): “Sorting, Learning, and Mobility When Jobs Have Scarcity Value”, *Carnegie-Rochester Conference Series on Public Policy* 46, 327-337.
- [10] Feller, W. (1965): “*An Introduction to Probability Theory and Its Applications*”, Vol. 1, Wiley, New York, USA.
- [11] Gibbons, R., and M. Waldman (1999a): “Careers in Organizations. Theory and Evidence”, in *Handbook of Labor Economics*, Volume III, O. Ashenfelter and D. Card Eds., North-Holland, Amsterdam, the Netherlands, 2373-2437.
- [12] Gibbons, R., and M. Waldman (1999b): “A Theory of Wage and Promotion Dynamics Inside Firms”, *Quarterly Journal of Economics* 114(4), 1321-1358.
- [13] Harris, M., and Y. Weiss (1984): “Job Matching with Finite Horizon and Risk Aversion”, *Journal of Political Economy* 92(4), 758-779.
- [14] Holmström, B., and P. Milgrom (1991): “Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design”, *Journal of Law, Economics and Organization* 7, 24-52.
- [15] Holmström, B., and J. Tirole (1989): “The Theory of the Firm”, in *Handbook of Industrial Organization*, Volume I, Chapter 2, R. Schmalensee and R.D. Willig Eds., North-Holland, Amsterdam, the Netherlands, 61-133.
- [16] Jovanovic, B., and Y. Nyarko (1997): “Stepping Stone Mobility”, *Carnegie-Rochester Conference Series on Public Policy* 46, 289-325.
- [17] Kakigi, R. (1983): “A Note on Discounted Future Two-Armed Bandits”, *Annals of Statistics* 11(2), 707-711.

- [18] Mas-Colell, A., M.D. Whinston and J.R. Green (1995): “*Microeconomic Theory*”, Oxford University Press, New York, USA.
- [19] Milgrom, P.R., and J. Roberts (1992): “*Economics, Organization, and Management*”, Prentice-Hall, Englewood Cliffs, USA.
- [20] Rhenius, D. (1974): “Incomplete Information in Markovian Decision Models”, *Annals of Statistics* 2(6), 1327-1334.
- [21] Rieder, U. (1975): “Bayesian Dynamic Programming”, *Advances in Applied Probability* 7(2), 330-348.
- [22] Rothschild, M. (1974): “A Two-Armed Bandit Theory of Market Pricing”, *Journal of Economic Theory* 9(2), 185-202.
- [23] Rudin, W. (1976): “*Principles of Mathematical Analysis*”, McGraw-Hill, New York, USA.
- [24] Schäl, M. (1979): “On Dynamic Programming and Statistical Decision Theory”, *Annals of Statistics* 7(2), 432-445.
- [25] Shiryaev, A.N. (1995): “*Probability*”, Springer-Verlag, New York, USA.
- [26] Waldman, M. (1984b): “Job Assignments, Signalling, and Efficiency”, *Rand Journal of Economics* 15(2), 255-267.