Career Concerns and the Nature of Skills*

Gonzalo Cisternas
MIT Sloan

May 2017

Abstract

I examine how career concerns are shaped by the nature of productive actions taken by workers. A worker's skills follow a Gaussian process with an endogenous component reflecting human-capital accumulation. Effort and skills are substitutes both in the output process (as in Holmström, 1999) and in the skills technology. The focus is on deterministic equilibria by virtue of Gaussian learning. When effort and skills are direct inputs to production and skills are exogenous, effort can be inefficiently high in the beginning of a career. In contrast, when skills are the only input to production but accumulate through past effort choices, the worker always underinvests in skill acquisition. At the center of the discrepancy are two types of ex post errors that arise at interpreting output due to an identification problem faced by the market. The robustness of the underinvestment result is explored via variations in the skill-accumulation technology and in the information structure, and policy implications are discussed.

JEL codes: D82, D83, J24.

1 Introduction

A common force driving many productive decisions by workers is the possibility of being perceived as skilled in a certain area of expertise. While such career concerns are a defining feature of many occupations, new technologies that facilitate the flow of information about potential employees, as well as online marketplaces that enable matches between firms and workers, naturally lead these motives to play an even more predominant role nowadays. Studying how markets reward ability is an relevant topic for two reasons. First, it helps to

*Email: gcistern@mit.edu. A previous version of this paper was circulated under the title "Shock Persistence, Endogenous Skills and Career Concerns". I thank Yuliy Sannikov for advice and guidance, and Alessandro Bonatti, Hector Chade and Bob Gibbons for very helpful discussions. I am also grateful to Benjamin Brooks, Sylvain Chassang, Henry Farber, Leandro Gorno, Johannes Hörner, Jay Lu, Alex Mas, a co-editor and three anonymous referees for very useful comments.

uncover the mechanisms through which workers can build their reputations in labor markets. Second, it helps to evaluate whether market-based incentives induce appropriate productive decisions by workers, thereby contributing to guide economic policy.

In traditional models of career concerns (e.g., Holmström 1982/1999; Dewatripont et al., 1999; Bonatti and Hörner, 2017), reputation is understood as a market's belief concerning a worker's unobserved measure of productive ability modeled as an exogenous type (e.g., "talent," either fixed or stochastically evolving over time); the worker's action (e.g., "effort") is in turn a direct input to production. In many industries, however, workers become more skilled on the job: for example, scientists write better papers as a consequence of doing research; surgeons become more skilled through performing surgeries; and lawyers acquire more experience through repeated litigation. In all these occupations, workers' productive actions can have persistent effects on output via the accumulation of *human capital*.

The systematic analysis of human capital began with the seminal work of Becker (1962, 1975). Despite the importance of this theory in explaining properties of the distribution of earnings (Ben-Porath, 1967; Rosen, 1976) and its implications for economic growth (Lucas, 1988; Mankiw et al., 1992), little attention has been devoted to the determinants of on-the-job skill acquisition in settings where workers behave strategically, such as when they have reputational concerns. Industries such as those previously mentioned are critical components of many developed economies, and a significant amount of skill development takes place on such jobs. Understanding the forces behind reputationally driven skill acquisition is thus of first-order importance, as human capital is arguably the main source of economic growth in modern times (Goldin and Katz, 2008).

Main contribution. This paper contributes to the understanding of incentives in such labor markets by developing a flexible model of career concerns with strategic skill acquisition. The model nests the cases of exogenous productivity studied in Holmström (1999) as special instances within a larger class of Gaussian processes for skills that allow for endogenous trends reflecting learning-by-doing. The main finding relates to how the inability of markets to directly observe workers' productivities interplays with workers' incentives depending on the nature of their productive actions (i.e., effort as a direct input to production versus an input to skills), and thus, on the nature of skills (i.e., exogenous versus endogenous). Specifically, while the inability to observe productivities can generate strong incentives by allowing the output consequences of productive actions of short-term nature to be overrepresented in wages, it can also severely damage workers' incentives by preventing actions that enhance productivity—and thus have long-term impact on output—to be fully priced into wages.

To illustrate the previous finding, Sections 2 and 3 analyze a particular specification of this model: the worker's hidden actions can simultaneously affect current output and the rate of change of skills. At one extreme, there is a *pure career-concerns model*: effort is a direct input to production, and skills evolve exogenously. This model fits situations in which a measure of innate ability is a critical component of the production process (e.g., managerial talent). At the other extreme, there is a *pure human-capital model*: effort affects skills directly, and the latter is the only input to production. This model captures environments in which current effort can have persistent effects on future output through skill acquisition (e.g., experience in litigation). Actions and skills enter additively separably in the output and skill technologies, which renders learning to be Gaussian. Normally distributed beliefs and linear payoffs in turn enable the analysis to focus on deterministic equilibria in which actions depend only on calendar time, as the latter maps one-to-one to the contemporaneous degree of uncertainty about skills.

In this baseline model, output is public, and hence the labor market learns about the worker's skills using the entire time series of realized production. Moreover, competition forces wages to reflect the market's expectation of next-period output at all times, so the wage paid is an affine function of the market belief about the worker's skills. Crucially, the market updates it belief upwards (downwards) if realized output is above (below) the market's expectation. Thus, wages are revised based on the observations of *output surprises*.

To understand how the worker's incentives interact with learning depending on the nature of his skills, suppose the market conjectures that the worker follows a deterministic strategy. Upon observing an output surprise, therefore, the market faces an *identification problem*: how much of the unanticipated output observed is due to a change in skills versus noise in the output signal? This informational limitation leads the market to make two types of ex post errors: first, to incorrectly attach a positive probability that an output surprise driven by noise was instead due to a change in skills; second, if the signal observed was driven by a change in skills, Bayesian learning attaches a non-zero probability that it was due to noise.

In a pure career-concerns setting, effort choices, like noise, influence only current output. Upon observing an output surprise due to an increase in effort, the market attaches a positive probability to such surprise being triggered by a shock to skills, and skills have persistence. Therefore, due to learning, effort has a persistent effect on wages despite its only transitory effect on output. In contrast, in a pure human-capital setting, effort has a persistent effect on output, just as shocks to skills do. However, the identification problem leads the market to always attach a positive probability to noise being the driver of the additional output that results from effort choices. Learning thus leads effort to having an effect on wages that is weaker than the corresponding impact on output.

This dual way in which the identification problem operates can have non-trivial implications. Specifically, when actions have direct, yet short-term, effects on output, and skills are sufficiently autocorrelated and the initial prior sufficiently dispersed, effort is inefficiently high in the beginning of a career; in such cases, beliefs are overly responsive to output surprises, leading effort to have an excessively strong effect on wages. In contrast, when actions have delayed, yet long-lived, effects on output via skill acquisition, effort is inefficiently low throughout a worker's entire career, regardless of the degree of autocorrelation in skills and initial uncertainty. Thus, there is *underinvestment* in skill acquisition.

Dynamics. Regarding the dynamics of incentives, effort choices always decay over time, regardless of the nature of the worker's skills. This unifying result is important for two reasons. First, it confirms that the intuition that the returns to reputation formation can decay over a worker's career purely due to informational considerations is likely to be a more robust phenomenon. Second, it sheds light on the mechanism that drives this intuition. Specifically, while beliefs become less sensitive to output surprises as learning progresses (an effect that dampens incentives), they also acquire more persistence (which strengthens incentives). The fact that the first effect dominates the second simply reflects that the total responsiveness of beliefs over long horizons decays as learning progresses.

In the long run, effort always settles below efficiency due to the combination of impatience and beliefs responding too little to new information; larger degrees of skill autocorrelation nevertheless alleviate this inefficiency by increasing the overall responsiveness of beliefs. Interestingly, as the persistence of skills approximates the discount rate, incentives in pure career-concerns environments become arbitrarily close to efficiency. In fact, the additional persistence of the market's belief that results from higher levels of autocorrelation asymptotically offsets the worker's rate of impatience, leading the sensitivity and persistence of wages to approximately balance each other; but the ratio of the latter fully characterizes incentives when payoffs and learning are linear. The comparative static shows, however, that attaining efficiency requires the worker's utility to become arbitrarily large in absolute value.

Robustness. The existence of deterministic equilibria relies on the linearity of payoffs and the substitutability between effort and skills: if either one is relaxed, the myopic return from affecting output will depend on the worker's reputation, leading to equilibrium actions that will also depend on past performance.[1] Critically, however, both the identification problem and the way in which it hurts incentives in pure human-capital settings, will not disappear.

It is then natural to explore more general skill-accumulation technologies under the umbrella of the previous assumptions. In this line, Section 4 characterizes deterministic equilibria via a *certainty equivalent*: provided the skill-accumulation technology is deterministic

[1] Cisternas (2017) allows for nonlinear payoffs using the signal technology of a pure career-concerns model.

(i.e., the endogenous component of skills is a general function of past effort choices exclusively), equilibria of this sort correspond to solutions to a decision problem with certainty.

The certainty equivalent shows that the equilibrium and socially efficient investment policies found in the baseline model are, more generally, measures of the reputational and social value of a *transitory* increase human capital. While these reputational and social measures of the value of human capital are only of local nature, their uniform ranking suggests that workers will always underinvest in skills. I confirm this intuition for the case of linear technologies; i.e., the rate of growth in skills is a linear function of past effort choices.

Information structures and policy. The previous results offer insights on the role that incomplete information and imperfect monitoring, via learning, can have on labor-market outcomes, and thus have the potential to guide economic policy. Wage subsidies and better monitoring technologies, for instance, alleviate any underprovision of effort by making wages more sensitive to output, and hence, to effort. Moreover, observing skills enables competition to correctly price any change of skills into wages, thereby inducing efficient incentives in pure human-capital settings via Becker's logic. In pure career-concerns settings, however, it destroys the possibility of building a reputation, as beliefs cannot be affected.

In reality, (i) obtaining arbitrarily precise information about skills or (ii) having access to the whole current history of output realizations as in the baseline model is, in many instances, prohibitively costly or simply impossible; potential employers must usually rely on coarser measures of performance to assess workers' skills when the latter are hidden. Prominent examples are online platforms that facilitate the matching between potential employers and workers, such as freelancers, lawyers, and doctors, by offering compressed information about past performance in the form of "scores" or "ratings."[2] In those marketplaces, employer-employee matches are in many instances short-lived, and passing information about performance to future employers occurs only through the platform's ratings process.

Building on recent methods by Hörner and Lambert (2016) on information structures for pure career-concerns models, Section 5 presents an example of how information design can be used as a policy tool to ameliorate the underinvestment result found in pure human-capital settings. This is done in a context where skills are stationary and the demand side operates as short-term employers who have access to a single signal about performance. Intuitively, ratings that exacerbate past performance relative to Bayesian learning as in the baseline model generate beliefs that are less responsive to changes in such ratings, reflecting that belief updating under coarser information structures attempts to correct the overweighing of information. For sufficiently persistent skills, the rating studied takes advantage of the high

[2]Some examples are upwork.com (freelancers), healthgrades.com (doctors), and avvo.com (lawyers).

autocorrelation of skills by discounting past information heavily, thereby making beliefs very responsive to performance. This, in turn, generates incentives arbitrarily close to efficiency.

Related Literature. The reputation literature employing exogenous types is vast. In the context of ex ante symmetric uncertainty, Holmström (1982/1999) analyzes career concerns when effort and skills are substitutes in the production function; complementarities are studied by Dewatripont et. al (1999a,b) in a two-period context and by Bonatti and Hörner (2015) in a dynamic setting with coarse information. Other forms of reputation building, such as mimicking irrational types (Fudenberg and Levine, 1992; Faingold and Sannikov, 2011) and signaling (Spence, 1973; Mailath and Samuelson, 2001), rely on asymmetric information.

This paper also contributes to the literature on strategic models of hidden or unverifiable investments. Regarding agency models, Kahn and Huberman (1988) and Prendergast (1993) find thar up-or-out contracts and promotions can sometimes alleviate hold-up problems when workers can invest in firm-specific human capital and output is non-contractible; I instead explore market-based incentives in settings where competition and human capital of a general nature make such problems mute. More recently, a growing body of literature on stochastic games with hidden investments has analyzed strategic behavior drawing mostly on firms' reputations as the main source of inspiration. Bohren (2016) and Huangfu (2015) study general models in which a long-run player can affect an observable payoff-relevant state that accumulates slowly. Board and Meyer-ter-Vehn (2013) and Dilme (2016) focus on unobservable investments in settings where (i) the state is hidden and can change discretely and (ii) information about it arrives infrequently. From a modeling perspective, the setting I study, like the first two papers, incorporates a payoff-relevant state that changes gradually; in line with the second group, this stock is hidden and inferred from a public signal. From an economic perspective, these papers are concerned with how incentives may vary across different levels of reputation; in contrast, I shut down this channel to isolate how such incentives depend on the degree of uncertainty about current economic conditions exclusively.

The study of non-strategic investments in human capital began with Becker (1962), who introduced the distinction between general and specific skills. In the model I analyze, human capital is general or task-specific (Gibbons and Waldman, 2004) rather than specific to the firm. Ben-Porath (1967) adopts an optimal-control approach with certainty in contexts in which training in general skills drives resources away from production; in contrast, I study skill accumulation resulting from effort directly devoted to production. On the demand side, a large body of work analyzes an employer's incentives to invest in general skills when labor markets are imperfect (e.g., Acemoglu and Pischke, 1999). I instead focus on the supply side, with the imperfection being the unobservability of workers' productivities and actions.

Finally, this paper contributes to the literature using continuous-time methods to analyze learning environments. Jovanovic (1979) analyzes labor-market outcomes in settings where firms and workers learn about the quality of their match and there is turnover. More closely related are strategic models of symmetric uncertainty that exploit the tractability of Holmström's signal-jamming technology: Cisternas (2017) studies learning-driven ratchet effects; Hörner and Lambert (2016) examine optimal (Gaussian) information structures for pure career-concerns models; and Prat and Jovanovic (2014) and DeMarzo and Sannikov (2016) study optimal contracts with learning.

Outline. Section 2 describes the basics of the baseline model, while Section 3 performs the corresponding equilibrium analysis. Section 4 examines general deterministic skill-accumulation technologies. Section 5 studies information design as a policy tool. Appendix A contains all the proofs, while Appendix B discusses the model's prediction on wages.

2 Baseline Model

2.1 Output and Human Capital Technologies

A worker and a pool of potential employers (or *market*) interact over an infinite horizon. While working with any employer, the worker produces output continuously over time. Cumulative output $(\xi_t)_{t\geq 0}$ evolves according to the technology

$$d\xi_t = (\lambda a_t + \theta_t)dt + \sigma_\xi dZ_t^\xi,\ t \geq 0, \tag{1}$$

where $(Z_t^\xi)_{t\geq 0}$ is a Brownian motion. In this specification, θ_t represents any aspect of the worker's time-t productive capability that exhibits persistence (and thus behaves as a stock variable); I refer to it as the worker's *skills* at time t. By contrast, a_t denotes the worker's effort at the same time; this (flow) choice variable takes values in an interval $[0, \bar{A}]$, with $\bar{A} > 0$ sufficiently large. The parameter $\sigma_\xi > 0$ measures the volatility of shocks to output, and $\lambda \in [0, 1]$ is a scalar that captures the importance of effort as a direct input to production.[3]

The worker's skills are governed by the dynamic

$$d\theta_t = (\kappa\theta_t + (1-\lambda)a_t)dt + \sigma_\theta dZ_t^\theta,\ t \geq 0, \tag{2}$$

where $(Z_t^\theta)_{t\geq 0}$ is a Brownian motion independent of $(Z_t^\xi)_{t\geq 0}$, and κ is a scalar. The drift

[3]In this continuous-time setting, therefore, the interval $[t, t+dt)$, dt small, can be interpreted as the relevant time unit of production. Realized output over the same "period" is then given by the increment $d\xi_t$; on average, output will change at rate $\lambda a_t + \theta_t$ at time $t \geq 0$.

of (2), $\kappa\theta_t + (1-\lambda)a_t$, reflects that skills are affected by the worker's effort choices. When $\lambda = 1$, skills evolve exogenously and effort is solely a direct input to production, resulting in a *pure career-concerns* model. Instead, when $\lambda = 0$, skills are the only input in (1), and effort is solely an input to skills; I refer to this case as a *pure human-capital* model.

The solution to (2) as a function of past effort choices and productivity shocks is

$$\theta_t = e^{\kappa t}\theta_0 + \int_0^t e^{\kappa(t-s)}(1-\lambda)a_s ds + \int_0^t e^{\kappa(t-s)}\sigma_\theta dZ_s^\theta,\ t \geq 0. \tag{3}$$

Observe first that the second term on the right-hand side of (3) measures the total amount of skills acquired on the job. Importantly, the technology through which the worker becomes endogenously more skilled is given by the linear function $a \mapsto (1-\lambda)a$ in the drift of (2), which is deterministic. General skill-accumulation technologies are studied in Section 4.

Second, observe that, as κ grows, past productivity shocks $(Z_s^\theta : 0 \leq s \leq t)$ have a higher weight on current skills; i.e., shocks to productivity become more persistent.[4] Thus, κ captures the degree of autocorrelation of skills. If $\kappa < 0$, the worker's skills are mean-reverting, as the drift of (2) satisfies

$$\kappa\theta_t + (1-\lambda)a_t = -|\kappa|(\theta_t - \underbrace{(1-\lambda)a_t/|\kappa|}_{\eta_t :=}) = -|\kappa|(\theta_t - \eta_t);$$

i.e., skills tend to move towards the (potentially) time-varying trend $(\eta_t)_{t\geq 0}$ at rate $|\kappa|$. Instead, when $\kappa \geq 0$, productivity shocks have permanent effects on skills.[5]

Remark 1. Holmström's (1982/1999) seminal analysis corresponds to a pure career-concerns model ($\lambda = 1$), in which skills either evolve as a random walk ($\kappa = 0$ and $\sigma_\theta > 0$) or are fixed over time and drawn from a normal distribution ($\kappa = \sigma_\theta = 0$).

Remark 2. Specifications of productivity such as (2) have found support in labor economics. Specifically, Farber and Gibbons (1996) and Kahn and Lange (2014) find that Gaussian models of skills that exhibit autocorrelation fit data on wages better than models in which productivity is fixed over time (or in which productivity is stochastic, but observable; Kahn and Lange, 2014). These papers are concerned with empirically testing the hypothesis that employers learn about workers' productivities, and hence, they do not model incentives.[6]

[4] Alternatively, since $e^{-\kappa s}$ is continuous of bounded variation, $\int_0^t e^{\kappa(t-s)}dZ_s^\theta = Z_t^\theta + \kappa\int_0^t e^{\kappa(t-s)}Z_s^\theta ds$ almost surely, where the last integral is of the Riemann-Stieltjes type (due to $Z_t^\theta \in C([0,t];\mathbb{R})$, $t \geq 0$). See, for instance, Theorem 2.3.7, in Kuo (2006).

[5] The model could be interpreted as one in which skills are fixed over time but their value in a (potentially rapidly) changing world varies stochastically.

[6] See also Altonji and Pierret (2001) for a model of employer learning in which unobserved productivity is constant over time.

2.2 Information and Labor-Market Structures

The output process (1) is publicly observed by all the agents in the economy, while effort choices are directly observed by the worker only. The skill process (2) is hidden to everyone, and all market participants share the prior that its initial value, θ_0, is normally distributed with mean $p_0 \in \mathbb{R}$ and variance $\gamma_0 \in \mathbb{R}_+$ (i.e., there is ex ante symmetric uncertainty). Thus, output is the only source of information about skills to everyone, and it fully determines the market's information structure. Denote by $\mathcal{F}_t$ the market's time-t information (which consists of the observations $\xi^t := (\xi_s : 0 \leq s < t)$ exclusively), and $\mathbb{F} := (\mathcal{F}_t)_{t\geq 0}$ the corresponding public filtration.

Since output is affected by the worker's actions, the market must anticipate the worker's behavior to correctly interpret $(\xi_t)_{t\geq 0}$. The market must then *conjecture* the worker's behavior to learn about skills, as his actions are effectively hidden due to the output technology (1) satisfying the full support assumption with respect to his effort choices. In this line, let $(a_t^*)_{t\geq 0}$ denote the market's conjecture of the (pure) strategy followed by the worker (which, in equilibrium, must be correct), and observe that a_t^* can be a function of time and of the current public history, ξ^t, only.[7] The market's conditional expectation operator under the belief that $(a_t^*)_{t\geq 0}$ is driving (1)-(2) is then denoted by $\mathbb{E}^{a^*}[\cdot|\mathcal{F}_t]$, $t \geq 0$.

The demand side of the labor market is competitive (e.g., the pool of employers consists of a large number of small agents), and no output-contingent contract can be written. As a result, competition forces wages to reflect, at any point in time, the market's expectation of production over the next instant; i.e., the market is *spot*. The time-t wage, w_t, then satisfies

$$w_t := \lim_{h\to 0} \frac{\mathbb{E}^{a^*}[\xi_{t+h} - \xi_t|\mathcal{F}_t]}{h} = \mathbb{E}^{a^*}[\lambda a_t + \theta_t|\mathcal{F}_t] = \lambda a_t^* + p_t^*, \tag{4}$$

where $p_t^* := \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t]$, $t \geq 0$. I sometimes write $w_t[p_t^*, a_t^*]$ to emphasize the dependence of w_t on the pair (p_t^*, a_t^*), and refer to $(w_t)_{t\geq 0}$ as the wage process.[8]

2.3 Strategies, Payoffs, and Equilibrium Concept

A (pure) strategy $(a_t)_{t\geq 0}$ for the worker any $\mathbb{F}$-progressively measurable process that takes values in $[0, \bar{A}]$ and such that (1)–(2) admits a solution.[9] Intuitively, at any time t, the

[7]The formal definition of a strategy is introduced in the next subsection.

[8]Since the marker's behavior is fully summarized by (4), the market plays no strategic role in the subsequent analysis. Also, because the worker internalizes all the expected surplus at all points in time and human capital is equally valued elsewhere, the worker has no strict incentive to switch employers.

[9]Because the action space is bounded, additional integrability conditions usually needed to define strategies (such as $\mathbb{E}\big[\int_0^t a_s^2 ds\big] < \infty$, $t \geq 0$) are automatically satisfied here. The set of pure strategies is non-empty, as any deterministic $(a_t)_{t\geq 0}$ clearly induces a unique solution to (1) where θ_t is given by (3) (i.e., the system

worker takes an action $a_t \in [0, \bar{A}]$ as a function of the current history ξ^t and time, $t \geq 0$.[10]

The worker is risk neutral and discounts flow payoffs according to a discount rate $r > 0$. Exerting effort is costly, which is captured by a strictly increasing and convex function $g : [0, \bar{A}] \to \mathbb{R}$ with $g(0) = g'(0) = 0$. Consequently, given the wage process (4), the total payoff from following the strategy $(a_t)_{t\geq 0}$ is given by

$$\mathbb{E}^a \left[\int_0^\infty e^{-rt}(w_t[p_t^*, a_t^*] - g(a_t))dt \right]. \tag{5}$$

The notation $\mathbb{E}^a[\cdot]$ simply emphasizes that different strategies $(a_t)_{t\geq 0}$ affect the likelihood of different realizations of $(\xi_t)_{t\geq 0}$, which in turn determine the realizations of the wage process.

In the following, I assume that the rate of change of skills satisfies $\kappa \in (-\infty, r)$; this is a necessary and sufficient condition for payoffs to be well defined. The focus is on Nash equilibria of the previous game between the market and the worker:

Definition 1 (Nash Equilibrium). *A pure strategy $(a_t^*)_{t\geq 0}$ is a Nash equilibrium if (i) $(a_t^*)_{t\geq 0}$ maximizes (5) subject to (1)-(2) and $w_t = \lambda a_t^* + p_t^*$, $t \geq 0$, among all pure strategies $(a_t)_{t\geq 0}$; and (ii) $(p_t^*)_{t\geq 0}$, is constructed via Bayes rule using $(a_t^*)_{t\geq 0}$. A Nash equilibrium is deterministic if $(a_t^*)_{t\geq 0}$ is a function of calendar time only.*

In a Nash equilibrium, the worker finds it optimal to follow the market's conjecture of play at all times. Thus, *along the path of play*, (i) the worker's behavior is sequentially rational, and (ii) the market holds a correct belief about his skills at all times.[11]

Deterministic equilibria are of interest because, as shown in the next section, learning is Gaussian; hence, the market's and the worker's posterior beliefs can be summarized by their corresponding (stochastic) posterior means (in the case of the market, by p_t^*, $t \geq 0$) and a common deterministic variance.

2.4 Discussion: On Deterministic Equilibria

More generally, the deterministic equilibria studied here belong to a larger class of *Markov* equilibria in which the worker's behavior depends on the history of the game only through the value that the common belief takes. In fact, since posterior means and time (through

can be solved in two steps; first (2), then (1)). For strategies that depend non-trivially on $(\xi_t)_{t\geq 0}$, existence is usually shown using Girsanov's theorem (in particular, observe that this reduces to showing existence to (2) with θ_t as in (3), as $(a_t)_{t\geq 0}$ depends explicitly on the realizations of $(\xi_t)_{t\geq 0}$, but not of $(Z_t^\theta)_{t\geq 0}$).

[10] For the analysis of equilibrium outcomes (i.e., actions and payoffs), leaving the worker's behavior after deviations unspecified is without loss, as the market can never observe the worker's actions.

[11] Regarding (ii), see Section 3.1 for the law of motion of $(p_t^*)_{t\geq 0}$ and its dependence on $(a_t^*)_{t\geq 0}$

the posterior variance) fully characterize posterior beliefs when learning is Gaussian, deterministic equilibria are simply Markov equilibria in which the dependence on p_t^* is absent.

Deterministic equilibria depend partly on the worker's flow payoff being linear in the public belief (via $w_t = \lambda a_t^* + p_t^*$).[12] If the linearity assumption is relaxed (as it is when the wage is a nonlinear function of p_t^*), incentives will depend on the stochastic history of the game through p_t^*: intuitively, since in this case the myopic return from affecting the public belief vary across different levels of reputation, so does the fully forward-looking value of being perceived as more skilled. This dependence is in fact the object of study of Cisternas (2017), where games of symmetric uncertainty about an exogenous state are analyzed. Importantly, these two papers are complementary. First, in Cisternas (2017), learning is stationary, and hence, beliefs can be identified with posterior means. Reducing beliefs to a one-dimensional state in turn simplifies the task of proving the existence of Markov equilibria when nonlinearities are present. In contrast, the current paper imposes no restriction on the initial degree of uncertainty about skills, but it constrains flow payoffs to be linear. Second, the approaches to performing equilibrium analysis followed in these papers differ, as off-path private beliefs complicate the analysis when nonlinearities are present; in contrast, in this linear model, the same strategy is optimal on and off the path of play.

3 Baseline Model: Equilibrium Analysis

In what follows, I fix a deterministic strategy $(a_t^*)_{t\geq 0}$ as the market's conjecture.

3.1 Learning Dynamics

By standard results in filtering theory, the distribution of θ_t conditional on $\mathcal{F}_t$ is Gaussian, and thus, the market's posterior belief is characterized by its mean and variance.[13] Recall that $p_t^* := \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t]$, and let $\gamma_t := \mathbb{E}^{a^*}[(\theta_t - p_t^*)^2|\mathcal{F}_t]$ denote the variance of the posterior belief at time $t \geq 0$. The Kalman-Bucy filter provides the laws of motion for these states:

[12] To the best of my knowledge, no equilibrium in which behavior explicitly depends on the stochastic history of the game ξ^t has been shown to exist in settings in which (i) the signal technology is additively separable, (ii) uncertainty is symmetric and Gaussian, (iii) and flow payoffs are linear.

[13] See Theorem 11.1 in Liptser and Shiryaev (1977). This follows from the action space being compact and the coefficients in (1)-(2) being constant (so conditions (11.4)-(11.11) in the Theorem are met).

Lemma 1. *The market's posterior mean and variance evolve according to*

$$
\begin{aligned}
dp_t^* &= (\kappa p_t + (1-\lambda)a_t^*)dt + \frac{\gamma_t}{\sigma_\xi^2}[d\xi_t - (\lambda a_t^* + p_t^*)dt],\ t>0,\ p_0 = p^0, \quad \textit{and} && (6)\\
\dot{\gamma}_t &= 2\kappa\gamma_t + \sigma_\theta^2 - \left(\frac{\gamma_t}{\sigma_\xi}\right)^2,\ t>0,\ \gamma_0 \in \mathbb{R}_+, && (7)
\end{aligned}
$$

where $Z_t^{a^*} := \frac{1}{\sigma_\xi}\left(\xi_t - \int_0^t (\lambda a_s^* + p_s^*)ds\right)$ *is a* $\mathbb{F}$*-Brownian motion from the market's perspective.*

Proof: See Theorem 12.1 in Liptser and Shiryaev (1977). □

The last term in (6) shows that changes in the market's posterior mean are driven by the realizations of $\sigma_\xi dZ_t^{a^*} := d\xi_t - \mathbb{E}_t^{a^*}[d\xi_t] = d\xi_t - (\lambda a_t^* + p_t^*)dt$, which is unpredictable from the market's perspective. The sensitivity of beliefs to such output surprises is captured by

$$\beta_t = \frac{\gamma_t}{\sigma_\xi^2},\ t \geq 0,$$

which multiplies unanticipated output in (6). Note that this sensitivity is deterministic because, from (7), the posterior variance $(\gamma_t)_{t\geq 0}$ is a function of time exclusively: when effort and skills are substitutes in the production function, any given change in the worker's action shifts the distribution of output by an amount that is independent of the current value that skills take, thereby leaving the informativeness of the output signal unaffected.

Because (6) is linear in p_t^*, it admits a solution that is a linear function of past output realizations:

$$p_t^* = p_0 e^{-\int_0^t (\beta_s - \kappa)ds} + \int_0^t e^{-\int_s^t (\beta_u - \kappa)du}(1-\lambda)a_s^* ds + \int_0^t e^{-\int_s^t (\beta_u - \kappa)du}\beta_s[d\xi_s - a_s^* ds]. \quad (8)$$

In this expression, the last term shows how the market corrects the output signal using the conjecture $(a_t^*)_{t\geq 0}$ in an attempt to construct an unbiased signal of the worker's current skills. In addition, observe that the sensitivity process $(\beta_t)_{t\geq 0}$ also affects the weight that beliefs attach to past output observations, as captured in $e^{-\int_s^t (\beta_u - \kappa)du}$, $t \geq s$.

The worker's belief about his own skills is summarized in dynamics analogous to (6)-(7) but replacing the market's conjecture of equilibrium play, $(a_t^*)_{t\geq 0}$, with the worker's actual strategy, $(a_t)_{t\geq 0}$ (in particular, the corresponding posterior variances coincide). The worker's posterior mean, $(p_t)_{t\geq 0}$, evolves according to

$$dp_t = (\kappa p_t + (1-\lambda)a_t)dt + \beta_t \sigma_\xi dZ_t,\ t \geq 0,$$

where $Z_t := \frac{1}{\sigma_\xi}\big(\xi_t - \int_0^t(\lambda a_t + p_s)ds\big)$ is a Brownian motion from the worker's perspective.[14] Equivalently, from his standpoint, output follows $d\xi_t = (\lambda a_t + p_t)dt + \sigma_\xi dZ_t$, $t \geq 0$.

Since the market's conjecture is deterministic, there are only two channels through which the worker can influence his wage process, $w_t = \lambda a_t^* + p_t^*$, $t \geq 0$: (i) by influencing the market's posterior mean via the channel of his actions *directly* affecting output and (ii) by influencing the market's posterior mean via the channel of his actions *indirectly* affecting output as he becomes endogenously more productive. Specifically, from (8), only the term

$$\int_0^t e^{-\int_s^t(\beta_u-\kappa)du}\beta_s d\xi_s = \int_0^t e^{-\int_s^t(\beta_u-\kappa)du}\beta_s[(\lambda a_s + p_s)dt + \sigma_\xi dZ_s],\ t \geq 0, \tag{9}$$

can be affected by the worker's actions, where in the last term I have used that $d\xi_t = (\lambda a_t + p_t)dt + \sigma_\xi dZ_t$ from the worker's viewpoint. Moreover, because $(Z_t)_{t\geq 0}$ is exogenous, only the component

$$Y_t := \int_0^t e^{-\int_s^t(\beta_u-\kappa)du}\beta_s(\lambda a_s + p_s)ds,\ t \geq 0, \tag{10}$$

of the wage process matters for the worker's incentives (i.e., for his effort choices).

To conclude, observe that the worker's flow payoff is linear in Y_t and the pair $(Y_t, p_t)_{t\geq 0}$ is linear in the worker's action. As a result, deterministic equilibria can be found via using pointwise optimization in the worker's problem. However, to better illustrate how learning and strategic behavior interplay depending on the nature of skills, the next two subsections study changes in $(Y_t, p_t)_{t\geq 0}$ in a heuristic fashion, and separately for the cases of pure career concerns and pure human capital. The formal arguments are direct corollaries of Proposition 5 in Section 4, which contains the main characterization result for deterministic equilibria.

3.2 Pure Career-Concerns Model

Consider now a pure career-concerns model (i.e., $\lambda = 1$) and contemplate a marginal increase in the worker's time-t effort along the path of play over a small period of time $dt > 0$; namely, a deviation $(a_t)_{t\geq 0}$ such that $a_t = a_t^* + \epsilon$ over $[t, t+dt)$, ϵ small, and $a_s = a_s^*$ for $s \neq t$. Denote by $\partial w_s/\partial a_t$ the change in the wage at time $s > t$ resulting from this effort increase. From the previous subsection, $\partial w_s/\partial a_t = \partial p_s^*/\partial a_t = \partial Y_s/\partial a_t$, where $\partial p_s^*/\partial a_t$ and $\partial Y_s/\partial a_t$ have analogous interpretations.

Since the worker's own belief process, $(p_t)_{t\geq 0}$, is exogenous when there is no skill accu-

[14]These results also follow from Theorem 12.1 in Liptser and Shiryaev (1977).

mulation, inspection of $(Y_t)_{t\geq 0}$ in (10) shows that

$$\frac{\partial Y_s}{\partial a_t} = [\beta_t e^{-\int_t^s (\beta_u - \kappa)du}]\epsilon dt, \;\; s \geq t.$$[15]

Intuitively, the deviation leads to an immediate response of the market's belief of size $\beta_t \epsilon dt$: output over $[t, t+dt)$ increases by $\lambda\epsilon dt = \epsilon dt$, whereas p_t^* responds to contemporaneous output with sensitivity β_t in (6). However, as learning progresses, this initial response decays exponentially according to $e^{-\int_t^s (\beta_u - \kappa)du}$, $s > t$.

The value of the additional wage stream consequence of this deviation is thus given by

$$\int_t^\infty e^{-r(s-t)} \frac{\partial w_s}{\partial a_t} ds = \underbrace{\left[\beta_t \int_t^\infty e^{-\int_t^s (r-\kappa+\beta_u)du} ds\right]}_{\mu_t :=} \epsilon dt = \mu_t \epsilon dt, \;\; t \geq 0, \tag{11}$$

and the total change in payoffs given by $[\mu_t - g'(a_t^*)]\epsilon dt$, ϵ small. In a deterministic equilibrium, therefore, $g'(a_t^*) = \mu_t$ must hold at all times; otherwise, the worker could increase his payoff by suitably choosing ϵ.

Importantly, because $(\mu_t)_{t\geq 0}$ is independent of the worker's past actions and output history, changes in effort today have no consequences for future behavior. Thus, the first-order condition $g'(a_t^*) = \mu_t$ is also sufficient:

Proposition 1 (Equilibrium behavior). *In a pure career-concerns model, there is a unique deterministic equilibrium that is characterized by*

$$g'(a_t^*) = \mu_t := \beta_t \int_t^\infty e^{-\int_t^s (r-\kappa+\beta_u)du} ds, \;\; t \geq 0. \tag{12}$$

Proof: See the Appendix. □

As argued earlier, the sensitivity of beliefs to output surprises, $\beta_t = \gamma_t/\sigma_\xi^2$, $t \geq 0$, plays a dual rule in the impulse response of beliefs, and hence, in the overall responsiveness of wages: it measures their initial reaction (β_t outside the integral) and it affects their rate of decay ($(\beta_s)_{s\geq t}$ present in the exponential). Importantly, for large enough γ_0 , $(\gamma_t)_{t\geq 0}$ is decreasing, and hence, there are two opposing forces affecting incentives as learning progresses. First, beliefs become less responsive to output surprises, which puts downward pressure on incentives; I refer to this as the *sensitivity effect*, captured by β_t outside the integral. Second,

[15]Formally, given $a : \mathbb{R}_+ \to [0, \bar{A}]$ measurable, let $D_v Y_s[a] := \lim_{\alpha\to 0} \frac{Y_s[a+\alpha v] - Y_s[a]}{\alpha}$ denote the directional derivative of Y_s at a in the direction $v : \mathbb{R}_+ \to \mathbb{R}$ measurable (which is well defined given that Y_t is a Lebesgue integral). Then, $D_v Y_s[a] = \int_0^s e^{-\int_t^s (\beta_u - \kappa)du} \beta_t v_t dt$ from where it is clear that $\partial Y_s/\partial a_t = D_{\epsilon \mathbb{1}_{[t,t+dt)}} Y_s[a]$. The same notion of directional derivative is behind all other forms of sensitivities $\partial \cdot / \partial a_t$ discussed in the paper.

any change in the public belief has more persistence, which in turn strengthens the worker's incentives to exert effort; I refer to this as the *persistence effect*, driven by $(\beta_s)_{s\geq t}$ in the exponential.

The next result shows that the sensitivity effect always dominates. To this end, let

$$\gamma^* := \sigma_\xi^2\left(\sqrt{\kappa^2 + (\sigma_\theta/\sigma_\xi)^2} + \kappa\right) \tag{13}$$

denote the variance of the market's long-run posterior belief when skills evolve as in (2) (i.e., the unique strictly positive stationary solution of (7)).[16] In particular, note that asymptotically, beliefs react to output surprises according to the sensitivity parameter

$$\beta^* := \frac{\gamma^*}{\sigma_\xi^2} = \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa,$$

and that past output observations are discounted at the constant rate $\beta^* - \kappa$. Whenever convenient, the dependence of γ^* and β^* on κ is explicitly stated.

Proposition 2 (Dynamics and convergence)**.** *Suppose that $\gamma_0 > \gamma^*$. Then, $(a_t^*)_{t\geq 0}$ is strictly decreasing. In the limit,*

$$g'(a_t^*) \to \frac{\beta^*}{r + \beta^* - \kappa} = \frac{\sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa}{\sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + r}, \ \ as \ \ t \to \infty. \tag{14}$$

In particular, long-run incentives are increasing in $\kappa \in (-\infty, r)$.

Proof: See the Appendix. □

The reason behind the decreasing time profile of effort lies in the nature of the learning process. Specifically, since learning is Gaussian, (i) time is a sufficient statistic for the current degree of uncertainty about skills, and (ii) this uncertainty decays monotonically if $\gamma_0 > \gamma^*$. The fact that the sensitivity effect dominates the discounting effect simply reflects that beliefs' overall responsiveness to new information over long horizons decreases as learning progresses, thus diminishing the effect that effort has on wages over time.[17]

The same logic reveals why the worker's long-run incentives become stronger as skills become more autocorrelated. In fact, since the environment becomes more uncertain as productivity shocks gain more persistence (i.e., $\partial\gamma^*(\kappa)/\partial\kappa > 0$), beliefs must become more

[16] Observe from (7) that, when $\gamma_0 > \gamma^*$, $(\gamma_t)_{t\geq 0}$ converges to γ^* monotonically from the right.

[17] As a corollary, the decreasing patterns of incentives found by Holmström (1982/1999) in the constant skills ($\kappa = \sigma_\theta = 0$) and random walk ($\kappa = 0$, $\sigma_\theta > 0$) cases are recovered.

responsive to output surprises overall, an effect that strengthens the worker's incentives—the ratio $\beta^*(\kappa)/(r+\beta^*(\kappa)-\kappa)$, which increases with κ because $\kappa \mapsto \beta^*(\kappa)$ is also increasing, is precisely a measure of the monetary value of such overall responsiveness of beliefs.

The previous discussion suggests that the persistence of shocks to productivity can potentially have important effects on the size of the incentives created by career concerns. Formally, observe that, in this pure career-concerns setting, total surplus satisfies

$$\mathbb{E}\left[\int_0^\infty e^{-rt}(d\xi_t - g(a_t)dt)\right] = \mathbb{E}\left[\int_0^\infty e^{-rt}(a_t - g(a_t))dt\right] + \underbrace{\mathbb{E}\left[\int_0^\infty e^{-rt}\theta_t dt\right]}_{\text{exogenous}}.$$

Thus, the efficient effort policy entails a constant effort level a^e characterized by $g'(a^e) \equiv 1$.

While inspection of (14) shows that long-run incentives are always inefficiently low, the next proposition shows that, given any degree of impatience $r > 0$, high levels of autocorrelation in skills can lead to inefficiently high effort early in a career, provided that there is enough initial uncertainty. To this end, let $a_0^*(\kappa;\gamma_0)$ denote the initial level of effort that arises in equilibrium as a function of κ for a fixed level of initial variance γ_0.

Proposition 3. *Suppose that* $\gamma_0 > \sigma_\xi^2(\sqrt{r^2+\sigma_\theta^2/\sigma_\xi^2}+r)$. *Then,* $a_0^*(\kappa;\gamma_0) > a^e$ *for* κ *sufficiently close to* r.[18]

Proof: See the Appendix. □

To understand the result, recall first that the market always believes that the worker exerts effort according to $(a_t^*)_{t\geq 0}$. Thus, upon observing an output surprise, the market must determine what fraction of such unexpected output is the consequence of a change in skills versus a shock to output. That is, the market faces an *identification problem.*

Importantly, a small deviation from a_t^*, like noise, (i) is unobservable, and (ii) affects only current output. The identification problem thus leads the market to always attach a positive probability to an output surprise consequence of additional effort instead being the outcome of a productivity shock. Productivity shocks, however, unlike noise, have persistent effects on output via the persistence of the skill process. Under sufficiently high initial uncertainty and autocorrelation in skills, the overall response of the market's belief to an output surprise (i.e., the sensitivity and persistence of beliefs combined over long horizons) is large enough

[18] This exercise fixes an initial level of uncertainty γ_0 (beyond a threshold), and then increases the degree of autocorrelation of skills to show that excessively high effort is an equilibrium outcome. An alternative exercise would involve fixing $\kappa \geq 0$ and $\sigma_\theta \geq 0$, and then making γ_0 arbitrarily large. While such an exercise can potentially lead to inefficiently high effort too (as it can be easily verified in a two-period model with a fixed type), it is not pursued here, as the interest is on the impact of the degree of persistence on incentives.

to allow effort to have a persistent effect on wages that is larger than its transitory impact on output, yielding inefficiently high effort.

Remark 3. Holmström (1982/1999) shows that efficiency can be attained when skills evolve as a random walk and the variance of the posterior belief is stationary by letting the worker become patient (i.e., $\kappa = 0$, $\gamma_0 = \sigma_\xi^2\sigma_\theta^2$, and $r \searrow 0$). Instead, (14) shows that, given any discount rate $r > 0$, long-run effort converges to efficiency provided that the skill process exhibits enough persistence (i.e., $r > 0$ is fixed, and $\kappa \nearrow r$). This complementary exercise is useful because it shows that, for the case of deterministic equilibria, efficiency occurs if and only if the worker attains unbounded utility. Specifically, it is easy to see that the worker's equilibrium utility satisfies

$$\mathbb{E}\left[\int_0^\infty e^{-rt}(a_t^* + p_t^* - g(a_t^*))dt\right] = \frac{p_0}{r-\kappa} + \underbrace{\int_0^\infty e^{-rt}(a_t^* - g(a_t^*))dt}_{\text{bounded}},$$

from where utility diverges as $\kappa \to r$.[19] However, there are a plethora of configurations (r, κ) such that, starting from γ_0 close to $\gamma^*(\kappa)$, incentives are arbitrarily close to efficiency and utility still bounded. Finally, observe that efficiency can be approximated only when $\kappa \geq 0$: it is only in this case that beliefs have enough persistence to offset the negative effect that the rate of impatience r has on $\beta^*(\kappa)/(r - \kappa + \beta^*(\kappa))$, thus generating an impact on wages of net present value close to 1.[20]

3.3 Pure Human-Capital Model

Consider now a pure human-capital model, i.e., effort influences the evolution of skills, but not output ($\lambda = 0$).

In this context, the main finding of this section pertains to an univocally negative prediction that the identification problem faced by the market has on the incentives to invest in human capital: in contrast to pure career-concerns settings, the worker always underinvests

[19] In fact, since $p_t = p_t^*$ in equilibrium, then, $\mathbb{E}[p_t^*|\mathcal{F}_0] = \mathbb{E}\left[e^{\kappa t}p_0 + \int_0^t \beta_s\sigma_\xi e^{\kappa(t-s)}dZ_s|\mathcal{F}_0\right] = e^{\kappa t}p_0$ from a time-zero perspective. Additionally, it is easy to check that $0 < \beta_t\int_t^\infty e^{-\int_t^s(r+\beta_u-\kappa)du}ds < \beta_0/(r-\kappa+\beta^*) < \beta_0/(r + \sigma_\theta/\sigma_\xi)$, where I used that $\beta^* - \kappa = (\kappa^2 + \sigma_\theta^2/\sigma_\xi^2)^{1/2} > \sigma_\theta/\sigma_\xi$, and so actions are bounded.

[20] Relatedly, the β^* term in the denominator of $\beta^*/(r + \beta^* - \kappa)$ can be seen as a form of *ratcheting*: after one-shot upward deviation from $(a_t^*)_{t\geq 0}$ occurs, the market incorrectly increases its expectation of future output (as it mistakenly believes that the worker is more skilled), and so the worker expects to underperform, thus explaining the extra β in the rate of decay of wages; but if he were to try to avoid underperforming, this would require exerting more effort in the future to account for the market's wrong belief, which in turn dampens his incentives by the same amount. Moreover, in this linear setting, this dynamic cost of exhibiting good performance is uniform across all levels of private beliefs. See Cisternas (2017) for more details.

in skill accumulation, regardless of both the degree of initial uncertainty and autocorrelation of skills. In addition, this incentive problem, as in the previous case, worsens over time.

For illustrational purposes, I proceed heuristically in a way that parallels the analysis of the previous section; i.e., by considering a deviation $(a_t)_{t\geq 0}$ such that $a_t = a_t^* + \epsilon$ over $[t, t+dt)$, ϵ small, and $a_s = a_s^*$ for $s \neq t$. From the worker's perspective, his skills satisfy

$$p_s = p_t e^{\kappa(s-t)} + \int_t^s e^{\kappa(s-u)} a_u du + \int_t^s e^{\kappa(s-u)} \beta_u \sigma_\xi dZ_u, \ s \geq t. \tag{15}$$

Consequently, from the second term in the previous expression, the worker to expects his time-s skills to increase by $\partial p_s/\partial a_t := e^{\kappa(s-t)}\epsilon dt$, $s \geq t$ due to the deviation. Inspection of $(Y_t)_{t\geq 0}$ (see eqn. (10)) then yields that the time-s wage changes by

$$\frac{\partial w_s}{\partial a_t} = \frac{\partial Y_s}{\partial a_t} = \left[\int_t^s e^{-\int_u^s (\beta_v - \kappa)dv} \beta_u e^{\kappa(u-t)} du\right] \epsilon dt = e^{\kappa(s-t)} \left[1 - e^{-\int_t^s \beta_u du}\right] \epsilon dt,$$

as the additional investment at time t affects all intermediate values of skills over $[t, s]$. The term in brackets plays an important role: it reflects that only a fraction $1 - e^{-\int_t^s \beta_u du}$ of the additional time-s output $e^{\kappa(s-t)}\epsilon dt$ that results from becoming more productive actually translates into additional wages.

The total value of the additional wages generated by this investment is thus given by

$$\begin{aligned} \int_t^\infty e^{-r(s-t)} \frac{\partial w_s}{\partial a_t} ds &= \left[\int_t^\infty e^{-r(s-t)} e^{\kappa(s-t)} \left[1 - e^{-\int_t^s \beta_u du}\right] ds\right] \epsilon dt \\ &= \underbrace{\left[\frac{1}{r-\kappa} - \int_t^\infty e^{-\int_t^s (r-\kappa+\beta_u)du} ds\right]}_{\rho_t :=} \epsilon dt, \end{aligned} \tag{16}$$

and observe that $\rho_t > 0$ follows from $\beta_s := \gamma_s/\sigma_\xi^2 > 0$, $s \geq 0$.

Since the total change in payoffs from this investment is, for ϵ small, $[\rho_t - g'(a_t^*)]\epsilon dt$, and ρ_t is independent of both the worker's past performance and actions, the first-order condition $g'(a_t^*) = \rho_t$ fully characterizes the worker's incentives to invest in human capital.

Proposition 4 (Equilibrium investment)**.** *In the pure human capital model, there is a unique deterministic equilibrium* $(a_t^*)_{t\geq 0}$ *that is characterized by*

$$g'(a_t^*) = \rho_t := \frac{1}{r-\kappa} - \int_t^\infty e^{-\int_t^s (r-\kappa+\beta_u)du} ds, \ t \geq 0. \tag{17}$$

Proof: See the Appendix. □

To interpret the result, note that the efficient investment policy is constant and satisfies

$$g'(a^{e,hc}) = \frac{1}{r-\kappa}. \tag{18}$$

In fact, from (3), $\theta_t = e^{\kappa t}\theta_0 + \int_0^t e^{\kappa(t-s)} a_s ds + \int_0^t e^{\kappa(t-s)}\sigma_\theta dZ_s^\theta,\ t \geq 0$, when $\lambda = 0$. Consequently, integration by parts yields that the planner's problem is to maximize

$$\mathbb{E}\left[\int_0^\infty e^{-rt}[d\xi_t - g(a_t)]dt\right] = \mathbb{E}\left[\int_0^\infty e^{-rt}\left[\frac{1}{r-\kappa}a_t - g(a_t)\right]dt\right] + \underbrace{\frac{\mathbb{E}[\theta_0]}{r-\kappa}}_{\text{exogenous}}; \tag{19}$$

(18) then follows from pointwise optimization in (19). Intuitively, a marginal increase in effort at $t \geq 0$ impacts future skills according to $e^{\kappa(s-t)}$, $s \geq t$, and hence, the additional stream of output created has a value $1/(r-\kappa)$. The following result holds.

Corollary 1 (Underinvestment)**.** *In the pure human-capital model, $a_t^* < a^{e,hc}$, for all $t \geq 0$.*

Proof: Follows directly from (17) and (18). □

At the center of this underinvestment result is the second type of ex post error that can arise as a result of the identification problem faced by the market.

Specifically, recall that, in a pure career-concerns setting, the worker's incentives can be inefficiently high due to the combination of (i) the market attaching a positive probability to output surprises consequence of extra effort instead being driven by changes in skills, and (ii) effort having a short-lived impact on output. In a pure human-capital setting, however, effort choices influence skills directly, and thus have persistent effects on output; i.e., they act as productivity shocks. However, due to the noise present in the output process, the market now attaches a positive probability to the output consequences of such investments in skills instead being the outcome of exogenous shocks to the output process. The economic consequence of this statistical result is that, as argued earlier, only $1 - e^{-\int_t^s \beta_u du}$ of the extra time-s output created by a time-t investment is priced into the time-s wage, $s \geq t$. More generally, from (17), only a fraction

$$1 - (r-\kappa)\int_t^\infty e^{-\int_t^s (r-\kappa+\beta_u)du} ds \in (0,1)$$

of the total output generated by a time-t investment maps into wages. Thus, while the identification problem can generate strong incentives by permitting productive actions of a short-term nature to be perceived as productivity shocks, it can damage incentives by preventing actions that have a long-term impact on output to be fully priced into wages.[21]

Regarding the dynamics of investment, recall that $(\beta_t)_{t\geq 0}$ is strictly decreasing when there is enough initial uncertainty. Thus, from (17), the underinvestment problem worsens over time as information accumulates:

Corollary 2 (Dynamics and convergence)**.** *Suppose that $\gamma_0 > \gamma^*$. Then, $(a_t^*)_{t\geq 0}$ is strictly decreasing. In the limit,*

$$g'(a_t^*) \to \frac{1}{r-\kappa} - \frac{1}{r-\beta^*+\kappa} = \frac{1}{r-\kappa} - \frac{1}{r+\sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2}}, \quad \textit{as} \ \ t \to \infty. \tag{20}$$

In particular, long-run incentives increase in $\kappa \in (-\infty, r)$.

The previous corollary is the investment analog of Proposition 2, i.e., the worker's incentives to invest in human capital are always stronger when the environment is more uncertain. In particular, the investments undertaken decay over time if uncertainty about the worker monotonically drops, and in the long-run if skills are less autocorrelated. As in the pure career-concerns case, this is the outcome of the market's belief becoming less responsive overall as the sensitivity $(\beta_t)_{t\geq 0}$ shrinks, which reduces the returns from exerting effort.[22]

To complete the analysis of this section, I make two observations. First, while the incentives to invest in skills are inefficiently low from a social perspective, they can be arbitrarily large relative to the incentives to exert effort when skills are exogenous (compare (14) and (20) as $\kappa \nearrow r$). This is because actions have a long-term impact on output when they affect skills, and this impact increases with the persistence of the skill process.[23]

Second, due to the linearity and additive separability of the technologies and flow payoffs, it is easy to verify that, in the combined model $\lambda \in (0,1)$, equilibrium incentives, $(a_t^{*,\lambda})_{t\geq 0}$, and efficiency, $(a_t^{e,\lambda})_{t\geq 0}$, are characterized by the linear combinations

$$g'(a_t^{*,\lambda}) = \lambda\mu_t + (1-\lambda)\rho_t \ \text{ and } \ g'(a_t^{e,\lambda}) = \lambda + (1-\lambda)\frac{1}{r-\kappa},$$

respectively.[24] Thus, for specific configurations of parameters (i.e., large enough γ_0, λ and

[21] In independent and parallel work, Kahn and Lange (2014) perform a numerical impulse-response analysis of the Kalman filter that leads them to suggest potential underinvestments in human capital. Because in their model worker behavior is not modeled, no explanatory mechanism is offered.

[22] Observe that $1/(r-\kappa) - 1/(r-\beta^*+\kappa) = [1/(r-\kappa)][\beta^*/(r-\kappa+\beta^*)]$ and so both pure models are isomorphic in the long run: for any unit of effort exerted, only $\beta^*/(r-\kappa+\beta^*) < 1$ of the output created translates into wages.

[23] Convergence to efficiency as $\kappa \nearrow r$ also occurs in this pure human-capital setting, now in terms of the fraction of total cost of effort exerted (as the efficiency benchmark diverges). Specifically, in the long run, $\lim_{\kappa\to r} g(a^*)/g(a^{e,hc}) = \lim_{\kappa\to r} g'(a^*)/g'(a^{e,hc}) = \lim_{\kappa\to r}[1/(r-\kappa) - 1/(r+\beta^*(\kappa)-\kappa)]/[1/(r-\kappa)] = 1$.

[24] These expressions follow directly from (22) in Proposition 5 and (23), respectively, in the next section.

κ), the combined model can contribute to generating incentives closer to efficiency early in a worker's career by averaging the excessive provision of effort as a production input, with the low provision of effort as an input to skills. Over time, however, inefficiently low effort provision prevails.

The next two sections are devoted to further exploring the robustness of deterministic equilibria and the associated underinvestment finding by studying (i) more general skill-accumulation technologies, and (ii) the information and labor market structures.

4 Deterministic Skill-Accumulation Technologies

The baseline model shows that deterministic equilibria can exist even though the environment being stochastic. Key to the analysis are that (i) the wage is linear in the market's belief, (ii) actions and skills are substitutes in (1) and (2), and that (iii) current effort directly affects the rate of change in contemporaneous skills.

It is easy to see that relaxing (i) or (ii) makes the instantaneous impact on effort on wages to depend non-trivially on the worker's reputation, and thus equilibrium behavior will depend non-trivially on past performance (i.e., on the stochastic history of the game). Importantly, since the identification problem faced by the market will not disappear in either case, the market will continue excessively punishing output surprises that result from unobserved changes in skills. Put differently, the mechanism that drives the underinvestment problem will continue negatively affecting the worker's incentives.

It is then natural to study the robustness of the findings of the previous section with respect to (iii). More generally, consider the skill process

$$d\theta_t = (\kappa\theta_t + \eta_t(a^t))dt + \sigma_\theta dZ_t^\theta,\ t \geq 0, \tag{21}$$

where $\eta_t(a^t) \in \mathbb{R}$ encodes how the past history of effort choices, $a^t := (a_s : 0 \leq s \leq t)$, and time itself, affect the current rate of skill change. The only economic restriction on such a technology is that it is *deterministic*. To this end, let $\mathcal{A} := \{x : \mathbb{R}_+ \to [0, \bar{A}] : x \text{ is measurable}\}$.

Definition 2. *The skill accumulation technology* $(\eta_t(\cdot))_{t\geq 0}$ *is said to be deterministic if there is* $\eta : \mathbb{R}_+ \times \mathcal{A} \to \mathbb{R}$ *such that (i)* $\eta_0(\cdot) \equiv \eta_0 \in \mathbb{R}$*, (ii)* $\eta_t(x) = \eta_t(\hat{x})$ *if* $x^t = \hat{x}^t$*,* $t > 0$*, and (iii) for all* $x \in \mathcal{A}$*, the function* $t \mapsto \eta_t(x) = \eta_t(x^t) \in L^1(\mathbb{R}_+)$ *(under the Lebesgue measure).*[25]

[25]Observe that (ii) simply states that $\eta_t(x)$ cannot depend on future values $\{x_s : s > t\}$, whereas (iii) ensures that payoffs are finite. Importantly, given a public strategy $(a_t)_{t\geq 0}$ and a realization $\tilde{\xi}$ of the public signal, $t \mapsto a_t(\tilde{\xi}) \in \mathcal{A}$ by definition of progressive measurability; thus, the set of feasible strategies

I assume that η_0 is publicly observed, but subsequent values of the trend are hidden from all market participants. Because the technology is deterministic, however, once a path of effort $(a_t)_{t\geq 0}$ is fixed, the entire continuation trend $(\eta_s(a^s))_{s\geq t}$ is known by the worker at time $t \geq 0$. Similarly, in the case of the market, provided it anticipates the worker's behavior.

The interpretation of the model is that η_0 represents an observable measure of skills upon entering the market (e.g., education). Worker heterogeneity thus comes from idiosyncratic realizations of the productivity shocks, $(Z_t^\theta)_{t\geq 0}$.[26] In the mean-reverting case ($\kappa < 0$), for instance, there is a high probability that any realization of the skills process fluctuates around the endogenous trend, $t \mapsto \eta_t(a^t)/|\kappa|$. Skilled workers are those who experience positive productivity shocks that make skills deviate more persistently from its trend.

4.1 Characterization of Deterministic Equilibria

Consider the model defined by the output and skill processes (1) and (21), with η deterministic. Further, recall that in the baseline model, the functions

$$\mu_t = \beta_t \int_t^\infty e^{-\int_t^s (r+\beta_u-\kappa)du} ds \;\text{ and }\; \rho_t = \frac{1}{r-\kappa} - \int_t^\infty e^{-\int_t^s (r+\beta_u-\kappa)du} ds,\; t \geq 0,$$

characterized incentives in the pure career concerns ($\lambda = 1$) and pure human capital ($\lambda = 0$) model, respectively. Specifically, μ_t captured the reputational value of a marginal increase in effort at time t when effort is a direct production input; ρ_t instead captured the reputational value of a marginal investment in skills at t when effort enters directly into the skills process.

The next result presents a characterization of deterministic equilibria as solutions to a simplified dynamic optimization problem with certainty; that is, a *certainty equivalent* for the worker's stochastic optimization problem given a deterministic conjecture. The functions $(\mu_t)_{t\geq 0}$ and $(\rho_t)_{t\geq 0}$ play a key role.

Proposition 5 (Certainty Equivalent for Deterministic Equilibria)**.** $a^* : \mathbb{R}_+ \to [0, \bar{A}]$ *measurable is a deterministic equilibrium if and only if it is the solution to*

$$\max_{a\in\mathcal{A}} \int_0^\infty e^{-rt} \left[\lambda \mu_t a_t + \rho_t \eta_t(a^t) - g(a_t)\right] dt. \tag{22}$$

Proof: See the Appendix. □

remains unchanged. Finally, the definition could be further generalized to $\eta_t = F(t, x^t, \eta^{t-})$, some $F : \mathbb{R}_+ \times \mathcal{A} \times L^1(\mathbb{R}) \to \mathbb{R}$, but this requires additional conditions that ensure that a solution $(\eta_t)_{t\geq 0}$ exists at the expense of no economic insights.

[26] In addition, specific parameters of the skill accumulation technology could vary across workers as well, but this type of heterogeneity is assumed to be observed by the market.

The previous result states that when the skill-accumulation technology is deterministic, the task of finding equilibria that depend exclusively on calendar time is not impossible; i.e. equilibria of this sort can arise despite the environment being stochastic. Moreover, from the certainty equivalent (22), the existence and uniqueness of deterministic equilibria is subject to technical conditions on η only.

To understand (22), note that, when $\eta_t(a^t) = (1-\lambda)a_t$ as in the baseline model, $g'(a_t^*) = \lambda\mu_t + (1-\lambda)\rho_t$ follows from applying pointwise optimization to (22); in this case, incentives at different points in time are connected only through the way in which uncertainty evolves. Under more general skill-accumulation technologies, however, the worker must also take into account how effort today affects his future wages via the channel of influencing η. Incentives thus become linked via the marginal impact of effort on the trend.

Remark 4. The specific technical conditions under which (22) has a solution will depend on the particular structure imposed on $(\eta_t)_{t\geq 0}$. Nevertheless, it is important to stress the generality of this result: it can accommodate (i) complementarities between η_t and a_t, (ii) multitasking environments, as well as (iii) problems of optimal control (e.g., settings where $(\eta_t)_{t\geq 0}$ is the solution to an ordinary differential equation, $\dot{\eta}_t = f(t, a_t, \eta_t)$).

4.2 The Reputational and Social Values of Human Capital

From (22) it is clear that, more generally, ρ_t captures the value of a *transitory* increase in the trend at time t.[27] Thus, ρ_t is a local measure of the reputational value of an additional unit of human capital. In contrast, it is easy to verify that the planner's solution is given by

$$\max_{(a_t)_{t\geq 0}} \int_0^\infty e^{-rt}\left[\lambda a_t + \frac{1}{r-\kappa}\eta_t(a^t) - g(a_t)\right] dt. \tag{23}$$

Consider now a pure human capital environment, i.e., $\lambda = 0$. In this case, inspection of (22)-(23) shows that the (local) social value of human capital, $1/(r-\kappa)$, is always above its contemporaneous reputational counterpart, ρ_t, $t \geq 0$.[28] The fact that this ranking is uniform across time suggests that, for technologies that are increasing in a^t (in a pathwise sense), the worker will acquire an inefficiently low stock of skills. This is because the sensitivity of the worker's flow payoff to changes in the trend is smaller in the equilibrium problem than in the planner's counterpart.

[27] Namely, the value of the additional wages resulting from a strategy $(\hat{a}_t)_{t\geq 0}$ such that $\eta_t(\hat{a}^t) = \eta_t(a^t)+1$ and that $\eta_s(\hat{a}^s) = \eta_s(a^s)$ for $s \neq t$.

[28] The interpretation of this result is exactly as in the baseline model, now considering the identification problem that is created by a transitory increase in the human capital trend.

Policies that increase $\rho_t > 0$ (thereby reducing the wedge $1/(r-\kappa)-\rho_t > 0$) will generate higher levels of human capital. In particular, a wage subsidy $w_t \mapsto (1+\alpha)w_t = (1+\alpha)p_t^*$, $\alpha > 0$, leads to $(1+\alpha)\rho_t > \rho_t$, $t \geq 0$. Similarly, better monitoring technologies, captured in a reduction in $\sigma_\xi > 0$ in (1), will increase $(\rho_t)_{t\geq 0}$ by making beliefs more sensitive to output surprises. This is transparent in the case of long-run incentives, where both

$$\beta^* = \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa \text{ and } \rho^* = \frac{1}{r-\kappa} - \frac{1}{r + \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2}}$$

increase as $\sigma_\xi > 0$ falls.

To conclude this subsection, I confirm the underinvestment intuition suggested by (22)-(23) for the case of linear technologies.

Definition 3. *A skill accumulation technology is linear if there are functions $h : \mathbb{R}_+ \to \mathbb{R}$ and $K : \mathbb{R}_+^2 \to \mathbb{R}_+$, $t \geq 0$, such that*

$$\eta_t(a^t) = h(t) + \int_0^t K(t,s)a_s ds, \; t \geq 0, \tag{24}$$

and where $\int_t^\infty K(s,t)ds < \infty$, for all $t \geq 0$.

In this specification, $h(\cdot)$ captures any aspect of skills growth that varies with experience but that is unrelated to the job itself, while $K(t,s)$ captures the impact of time-s effort on the trend at time $t \geq s$.[29] The following result follows directly from (22)-(23):

Proposition 6. *If skills accumulate linearly, equilibrium investment, $(a_t^*)_{t\geq 0}$, and socially efficient investment $(a_t^e)_{t\geq 0}$, satisfy*

$$g'(a_t^*) = \int_t^\infty e^{-r(s-t)}K(s,t)\rho_s ds \;\; \textit{and} \;\; g'(a_t^e) = \int_t^\infty e^{-r(s-t)}\frac{K(s,t)}{r-\kappa}ds \; t \geq 0,$$

respectively. Hence, there is underinvestment throughout the worker's entire career.

Proof: Direct from integration by parts in (22)-(23). □

Finally, observe that the baseline model is recovered when $h(t) \equiv 0$ via a sequence $(K_n(\cdot,\cdot))_{n\in\mathbb{N}}$ such that for each $t \geq 0$, $\lim_{n\to\infty} K_n(\cdot,t) = \lim_{n\to\infty} K_n(t,\cdot) = \delta_t(\cdot)$, where $\delta_t(\cdot)$ is the

[29] As an example of a linear technology, consider a stationary exponential kernel $h(t) = \eta_0 e^{-\delta t}$ and $K(t,s) = e^{-\delta(t-s)}\alpha$, for $\delta, \alpha \in \mathbb{R}_+$; i.e., $(\eta_t)_{t\geq 0}$ is the solution to the ODE $\dot{\eta}_t = \alpha a_t - \delta\eta_t$, $t \geq 0$. Non-stationary kernels can be interpreted as capturing time-varying returns to experience, i.e., non-increasing returns to experience arise when $K(t+h,t) \geq K(t'+h,t')$ for all $t' > t$ and $h > 0$.

Dirac measure at $t \geq 0$. The reader interested in the properties of the distribution of wages generated by the model can skip directly to Appendix B.

5 Information Structures and Policy

A strong informational assumption of the baseline model is that the demand side of the labor market has access to the complete time series of the worker's performance. While this is a reasonable approximation in industries such as academia (where the specifics of the output generated and its chronological appearance are public), it is less so in other occupations. At the other extreme are marketplaces in which employers enter a market platform to purchase short-term services from a worker and exit it once the relationship ends. For instance, many programmers, developers, accountants, and even doctors and lawyers operate as independent contractors (or "freelancers") who sell short-term labor to clients in online marketplaces (e.g., upwork.com, healthgrades.com and avvo.com). In those transactions, potential employers usually rely on coarser measures of past performance that are provided by the market platform.

The purpose of this subsection is to apply recent methods by Hörner and Lambert (2016) on effort-enhacing information structures for pure career concern models as a policy tool to alleviate the underinvestment problem in pure human capital settings. Specifically, employing a specific type of information structure, I show how a simple adaptation of their analysis to the case of human capital accumulation yields incentives that are arbitrarily close to efficiency when skills exhibit enough persistence; interestingly, the corresponding rating system used exhibits virtually no persistence.

5.1 Labor Market and Information

Consider a market in which a worker faces a sequence of short-term employers. Heuristically, at every instant t, a short-lived agent enters a marketplace demanding the worker's service for the period $[t, t+dt)$. Once production $d\xi_t$ is realized, the employer exits the market, and the same sequence of events repeats over $[t+dt, t+2dt)$, now with a new employer, etc.

Upon entering the market, the time-t employer observes a one-dimensional statistic of the worker's past performance $Y_t \in \mathbb{R}$ published by the platform. For simplicity, the worker has all the bargaining power, and hence, the time-t wage takes the form $\mathbb{E}^{a^*}[\theta_t|Y_t]$ (i.e., potential employers do not monitor changes in $(Y_t)_{t\geq 0}$). The specification for $(Y_t)_{t\geq 0}$ under study is

$$Y_t = \int_{-\infty}^{t} e^{-\phi(\beta^*+\kappa)(t-s)}[d\xi_s - A_s^* ds],\ t \geq 0, \tag{25}$$

where $\phi > 0$, $A_t^* := \mathbb{1}_{t\geq 0}\int_0^t e^{\kappa(t-s)}a_s^* ds$ is the total contribution of past (equilibrium) effort choices to time-t skills, and $\beta^* := \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa$ is the long-run sensitivity of beliefs of the baseline model. In particular, when $\phi < 1$, $(Y_t)_{t\geq 0}$ discounts past performance less heavily than the market's belief does asymptotically in the baseline model.[30]

The interpretation of (25) is one where the worker enters the market platform at time zero, which determines the starting point of the window over which the incentives for skill acquisition are studied. Output observations before time zero are then interpreted as measures of performance prior to entering the marketplace that were collected by the platform. The unbounded domain for the time variable (jointly with Assumption 1 below) is required to obtain a stationary investment policy, as in the long run analysis of the baseline model.

5.2 Stationary Exogenous Component of Skills

Recall from (3) that, in a pure human-capital setting, the worker's skills satisfy

$$\theta_t = \underbrace{\int_0^t e^{\kappa(t-s)}a_s ds}_{A_t:=} + \underbrace{e^{\kappa t}\theta_0 + \int_0^t e^{\kappa(t-s)}\sigma_\theta dZ_s^\theta}_{L_t:=} = A_t + L_t,\ t\geq 0. \tag{26}$$

In this expression, A_t is the endogenous component of the worker's time-t skills, which depends only on the realized history $(\xi_s : 0 \leq s \leq t)$. Instead, L_t is the exogenous component of the worker's time-t skills, which depends only on the realizations of $(Z_s^\theta : 0 \leq s \leq t)$. In particular, when $\kappa < 0$, $(L_t)_{t\geq 0}$ follows an Ornstein-Uhlenbeck process $dL_t = \kappa L_t dt + \sigma_\theta dZ_t^\theta$, $t > 0$, $L_0 = \theta_0$, with rate of mean reversion $|\kappa| > 0$. As Hörner and Lambert (2016) note, the environment becomes fully stationary when the following holds.

Assumption 1 (Stationarity of L)**.** *$(L_t)_{t\in\mathbb{R}}$ is a stationary centered (i.e., zero mean) Ornstein-Uhlenbeck process with covariance* $\text{Cov}[L_t, L_s] = \frac{\sigma_\theta^2}{2|\kappa|}e^{\kappa|t-s|}$*, $\kappa < 0$, defined on $\mathbb{R}$.*[31]

Under this assumption, (26) holds for all $t \geq 0$ under the additional requirement that $\theta_0 = L_0$ is independent of $(Z_t^\theta)_{t>0}$ and distributed according to $\mathcal{N}(0, \sigma_\theta^2/2|\kappa|)$. In particular, the efficiency benchmark is also trivially given by $g'(a^e) = 1/(r-\kappa)$.

[30] The rating (25) is the natural human-capital counterpart of the "exponential smoothing" rating introduced by Hörner and Lambert (2016). While restricted to stationary pure career-concerns settings, their analysis is more general, as they solve for the optimal information structure. From this perspective, the value of the application is threefold. First, it shows how their methods can be translated to pure-human capital environments. Second, it shows how information design can be used as a tool for economic policy. Third, it shows that the structure of ratings can depend nontrivially on the degree of persistence of skills, an exercise not explored in their analysis.

[31] Given $(B_t)_{t\geq 0}$ a standard Brownian motion, $(L_t)_{t\in\mathbb{R}}$ is constructed via $L_t := \frac{\sigma_\theta}{\sqrt{2|\kappa|}}e^{-|\kappa|t/2}B_{e^{|\kappa|t}}$, $t \in \mathbb{R}$.

5.3 Optimal $\phi > 0$ and asymptotic efficiency as $\kappa \nearrow 0$

Observe that given a deterministic conjecture $(a_t^*)_{t\geq 0}$, the process $(Y_t)_{t\geq 0}$ is centered Gaussian from the market's perspective. Thus, $\mathbb{E}^{a^*}[\theta_t|Y_t] = A_t^* + \mu_t$ where

$$\mu_t := \mathbb{E}^{a^*}[L_t|Y_t] = \frac{\text{Cov}[L_t, Y_t]}{\text{Var}[Y_t]} Y_t = \frac{\text{Cov}[L_t, Y_t]}{\text{Var}[Y_t]} \int_{s\leq t} e^{-\phi(\beta^*+\kappa)(t-s)}[d\xi_s - A_s^* ds].$$

as (Y_t, L_t) is normally distributed with zero mean. Moreover, since $(L_t, Y_t)_{t\geq 0}$ is stationary, $\text{Cov}[L_t, Y_t]/\text{Var}[Y_t]$ turns out to be constant over time. Let

$$\alpha(\phi) := \frac{\text{Cov}[L_t, Y_t]}{\text{Var}[Y_t]}.$$

where the dependence on $\phi > 0$ has been made explicit.

Proposition 7. *The equilibrium investment policy is constant and characterized by*

$$g'(a^*(\phi)) = \frac{\alpha(\phi)}{(r-\kappa)(r+\phi(\beta^*-\kappa))}, \quad \textit{where} \quad \alpha(\phi) = \frac{\phi\sigma_\theta^2}{\sigma_\xi^2(\beta^*-\kappa-\phi\kappa)}. \tag{27}$$

The optimal $\phi > 0$ is given by $\phi^*(r, \kappa) = \sqrt{r/|\kappa|}$. Moreover, $g'(a^*(\phi^*(r, \kappa))) \to 1/r$ as $\kappa \to 0$, i.e., incentives become asymptotically efficient as $\kappa \to 0$.

Proof: See the Appendix. □

First, observe that if $\phi = 1$, $\alpha = \beta^*$ follows from $(\beta^*)^2 + \sigma_\theta^2/\sigma_\xi^2 + 2\kappa\beta^* = 0$, and thus

$$g'(a^*(1)) = \frac{1}{r-\kappa} - \frac{1}{r+\beta^*-\kappa};$$

i.e., the incentives that arise in the long run of the baseline model are recovered (cf. Corollary 2), and thus there is underinvestment. As Hörner and Lambert note, when $(Y_t)_{t\geq 0}$ attaches the same weight to past observations as the long-run belief that conditions on the whole time series of output, the rating essentially discloses the belief that arises when output is public.

Interesting insights appear when studying the optimal persistence of $(Y_t)_{t\geq 0}$ as a function of the autocorrelation of skills. Specifically, using that $\beta^* = \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa$,

$$(\beta^* - \kappa)\phi^*(r, \kappa) = \left[r\left(|\kappa| + \frac{\sigma_\theta^2}{|\kappa|\sigma_\xi^2} \right) \right]^{1/2},$$

is non-monotone in the rate of mean reversion $|\kappa| > 0$. In particular, $(\beta^* - \kappa)\phi^*(r, \kappa) \to$

$+\infty$ as $\kappa \to 0$, and so the rating has no persistence in the limit. Moreover, using that $-\kappa\phi^*(r,\kappa) = \sqrt{r|\kappa|}$,

$$\alpha(\phi^*(r,\kappa)) = \frac{\phi^*(r,\kappa)\sigma_\theta^2}{\sigma_\xi^2(\beta^* - \kappa - \kappa\phi^*(r,\kappa))} \to +\infty \text{ as } \kappa \to 0,$$

so beliefs become infinitely sensitive to changes in the rating. Importantly, by continuity of both $a^*(\phi^*(r,\kappa))$ and $a^e(\kappa)$ at $\kappa = 0$, the limit $g'(a^*(\phi^*(r,\kappa))) \to 1/r$ as $\kappa \to 0$ in Proposition 7 implies that $|a^*(\phi^*(r,\kappa)) - a^e(\kappa)| < \epsilon$ for $|\kappa|$ small; i.e., for any given $r > 0$, sufficiently high persistence in the $\kappa < 0$ range leads the information structure $(Y_t)_{t\geq 0}$ to generate incentives arbitrarily close to the corresponding efficiency level.

The intuition for the result lies in the following trade-off: a statistic that overweighs past signals (i.e., ϕ is low) relative to Bayesian updating using the whole history of observations, by covarying less with current skills, generates beliefs that are less sensitive to changes in the statistic, and the latter effect hurts incentives ($\alpha(\phi) \to 0$ as $\phi \to 0$ in (27)). When $\kappa \in \mathbb{R}_-$ is close to 0, therefore, the optimal information structure within this class sacrifices persistence in $(Y_t)_{t\geq 0}$ in exchange for a large sensitivity in the market's belief, and this is optimal because skills are already very persistent. The result is a sensitivity effect $\alpha(\phi^*(r,\kappa))$ and a persistence effect $\phi^*(r,\kappa)(\beta^* - \kappa)$ that perfectly balance each other in the limit. Unlike in the baseline model, crucially, approaching efficient incentives does not require $r - \kappa$ to be arbitrarily small.

6 Concluding Remarks

This paper has analyzed the interplay between career concerns and the nature of skills, uncovering a reputationally driven mechanism that can lead to inefficiently low investments in human capital. In light of this finding, and given the role that human capital plays in economic growth, examining policies that incentivize skill acquisition acquires relevance. Promising areas include the optimal design of information structures and compensation schemes that account for actions having a long-term impact on output variables via skill acquisition.[32]

Regarding the assumptions of the model, the additive separability of the production function plays an important role by eliminating incentives for strategically affecting the speed of everyone's learning (i.e., experimentation). Importantly, this modeling device is convenient if (i) the goal is to understand how reputational incentives vary over workers'

[32]Preliminary work in the area of optimal contracting includes Kwon (2006) and Prat (2015) in specific environments, and the more general approach of Sannikov (2014).

career due to informational considerations, and (ii) one is willing to assume that uncertainty about a worker's skills, to a first-order approximation, decays with tenure in the labor market.

Second, while the Gaussian structure adds considerable tractability, the main forces found in this paper are likely to arise under other specifications of the randomness present in the environment. In particular, the incentives for effort provision (either as a direct production or skills input) are weaker when beliefs become less sensitive to new information, as the impact that effort has on wages is then diminished. Similarly, the underinvestment result is the consequence of a more fundamental identification problem that will be present under other distributional assumptions.

Finally, the assumption that the demand side of the labor market is perfectly competitive leads to wages that are linear in the market's belief, which is critical for the existence of deterministic equilibria. In particular, (i) labor supply, understood as effort provision, is completely inelastic at all points in time, and (ii) the covariance structure of wages is unaffected by the workers' equilibrium actions. As is well known, however, labor market imperfections can lead to wage compression (Acemoglu and Pischke, 1999), which can be modeled as wages that are a nonlinear function of the market's expectation of next-period production. Such nonlinearity would lead to incentives that explicitly depend on the posterior mean of the skill distribution. As a result, controlling for tenure, the supply of labor would depend non-trivially on wages, and the covariance structure of the latter would be endogenous. This and other questions are left for future research.

References

[1] Acemoglu, D. and J-S. Pischke (1999): "The Structure of Wages and Investment in General Training," *Journal of Political Economy*, 107, 539-572.

[2] Altonji, J. and C. Pierret (2001): "Employer Learning and Statistical Discrimination," *The Quarterly Journal of Economics*, 116 (1), 313-350.

[3] Baker, G. M. Gibbs and B. Holmstrom (1994). "The Wage Policy of a Firm". *The Quarterly Journal of Economics*, 109(4), 881-919.

[4] Becker, G. (1962): "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy*, 70(5), 9-49.

[5] ——. (1975): *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education.* New York: Columbia University Press.

[6] Ben-Porath, Y. (1967): "The Production of Human Capital and the Life Cycle of Earnings". *Journal of Political Economy, Vol. 75, No. 4, Part 1, 352-365.*

[7] Billingsley, P. (1995): *Probability and Measure, Third Edition.* John Wiley and Sons.

[8] Board, S., and M. Meyer-ter-Vehn (2013): "Reputation for Quality," *Econometrica*, 81(6), 2381-2462.

[9] Bohren, A. (2016): "Using Persistence to Generate Incentives in a Dynamic Moral Hazard Problem," mimeo, University of Pennsylvania.

[10] Bonatti, A., and J. Hörner (2017): "Career Concerns with Exponential Learning," *Theoretical Economics*, forthcoming.

[11] DeMarzo, P. and Y. Sannikov (2016): "Learning, Termination and Payout Policy in Dynamic Incentive Contracts," *Review of Economic Studies*, forthcoming.

[12] Cisternas, G. (2017): "Two-Sided Learning and the Ratchet Principle," *Review of Economic Studies*, forthcoming.

[13] Dewatripont, M., I. Jewitt, and J. Tirole, (1999a): "The Economics of Career Concerns, Part I: Comparing Information Structures". *Review of Economic Studies, Vol. 66, No. 1, 199-217.*

[14] ——. (1999b). "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies". *Review of Economic Studies, Vol. 66, No. 1, 199-217.*

[15] Dilme, F. (2016): "Reputation Building through Costly Adjustment," mimeo, University of Bonn.

[16] Faingold, E., and Y. Sannikov (2011): "Reputation in Continuous-Time Games." *Econometrica*, 79(3), 773-876.

[17] Farber, H. and R. Gibbons (1996): "Learning and Wage Dynamics". *The Quarterly Journal of Economics, Vol. 111, No. 4, 1007-1047.*

[18] Fudenberg, D. and D. Levine (1992): "Maintaining a Reputation when Strategies are Imperfectly Observed," *Econometrica, 59, 561-579.*

[19] Gibbons, R. and M. Waldman (2004): "Task-Specific Human Capital". *The American Economic Review, Vol. 94, No. 2, 203-207.*

[20] Goldin, C. and L. Katz (2008): *The Race Between Education and Technology.* Harvard University Press.

[21] Holmström, B. (1982): "Managerial Incentive Problems: A Dynamic Perspective". In *Essays in Economics and Management in Honor of Lars Wahlbeck,* Helsinki: Swedish School of Economics.

[22] ——. (1999): "Managerial Incentive Problems: A Dynamic Perspective". *The Review of Economic Studies, 66, 169-182.*

[23] Hörner, J. and N. Lambert (2016): "Motivational Ratings," Cowles Foundation Discussion Paper.

[24] Huangfu, B. (2016): "Stochastic Reputation Cycles," mimeo, University of Rochester.

[25] Jovanovic, B. (1979). "Job Matching and the Theory of Turnover". *The Journal of Political Economy*, 87(5), 972-990.

[26] Kahn, C. and G. Huberman (1988): "Two-Sided Uncertainty and "Up-or-Out" Contracts," *Journal of Labor Economics, 6(4), 423-424.*

[27] Kahn, L. and F. Lange (2014). "Employer Learning, Productivity and the Earnings Distribution: Evidence from Performance Measures". IZA Discussion Papers 5054, Institute for the Study of Labor (IZA).

[28] Kwon, I. (2006): "Incentives, Wages, and Promotions: Theory and Evidence," *RAND Journal of Economics, 37(1), 100-120.*

[29] Kuo, H-H. (2000). *Introduction to Stochastic Integration*. Springer.

[30] Liptser, R. and A. Shiryaev (1977): *Statistics of Random Processes I and II.* New York: Springer-Verlag.

[31] Lucas, R. (1988): "On the Mechanics of Economic Development," *Journal of Monetary Economics, 22, 3-42.*

[32] Mailath, G. and L. Samuelson (2001): "Who Wants a Good Reputation?," *Review of Economic Studies, 68, 415-442.*

[33] Mankiw, N., D. Romer and D. Weil (1992): "A Contribution to the Empirics of Economic Growth," *The Quarterly Journal of Economics, 107(2), 407-437.*

[34] Neal, D. and S. Rosen (2000): "Theories of the distribution of earnings," in Atkinson A.B. and Bourguignon, F. (eds) *Handbook of Income Distribution.* Amsterdam: North Holland Publishing Company.

[35] Oksendal, B. (2007): *Stochastic Differential Equations, Sixth Edition.* Springer.

[36] Prat, J. (2015): "Dynamic Contracts and Learning by Doing," *Mathematics and Financial Economics, 9(3), 169-193.*

[37] Prat, J. and B. Jovanovic (2014): "Dynamic Contracts when the Agent's Quality is Unknown," *Theoretical Economics*, 9(3), 865-914.

[38] Prendergast, C. (1993): "The Role of Promotion in Inducing Specific Human Capital Acquisition," *The Quarterly Journal of Economics, 108(1), 523-534.*

[39] Rosen, S. (1976): "A Theory of Life Earnings," *Journal of Political Economy, Vol. 84, No. 4, 545-567.*

[40] Sannikov, Y. (2014): "Moral Hazard and Long-Run Incentives," Working paper, Princeton University.

[41] Spence, M. (1973): "Job Market Signaling," *The Quarterly Journal of Economics, 87(3), 355-374.*

Appendix

Proofs of Propositions 1 and 4: They follow directly from point-wise optimization in (22) when $\lambda = 1$ and $\lambda = 0$, respectively. □

Proof of Proposition 2: To show that $\mu_t = \beta_t \int_t^\infty e^{-\int_t^s (r+\beta_t-\kappa)du} ds$ decreases over time, observe first that

$$\frac{d\log(\mu_t)}{dt} = \frac{\dot{\gamma}_t}{\gamma_t} + r + \frac{\gamma_t}{\sigma_\xi^2} - \kappa - \frac{1}{\int_t^\infty e^{-\int_t^s (r+\gamma_u/\sigma_\xi^2-\kappa)du} ds}.$$

If $\gamma_t > \gamma^*$, then,

$$\int_t^\infty e^{-\int_t^s (r+\gamma_u/\sigma_\xi^2-\kappa)du} ds < \frac{1}{r+\gamma^*/\sigma_\xi^2-\kappa} \Rightarrow \frac{d\log(\mu_t)}{dt} < \frac{\dot{\gamma}_t}{\gamma_t} + \frac{\gamma_t}{\sigma_\xi^2} - \frac{\gamma^*}{\sigma_\xi^2}.$$

Finally, from the ODE that governs γ_t, $\dot{\gamma}_t/\gamma_t + \gamma_t/\sigma_\xi^2 = 2\kappa + \sigma_\theta^2/\gamma_t$, so

$$\frac{d\log(\mu_t)}{dt} < 2\kappa + \frac{\sigma_\theta^2}{\gamma_t} - \frac{\gamma^*}{\sigma_\xi^2} < 2\kappa + \frac{\sigma_\theta^2}{\gamma^*} - \frac{\gamma^*}{\sigma_\xi^2} = 0$$

by definition of γ^*. The result follows from $\gamma_t > \gamma^*$ being equivalent to $\dot{\gamma}_t < 0$.

Finally, note that when $\kappa \in (-\infty, r)$,

$$\operatorname{sign}\frac{\partial}{\partial\kappa}\left(\frac{\sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2}+\kappa}{\sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2}+r}\right) = \operatorname{sign}\left[\left(\sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2}+\kappa\right)\left(\sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2}+r-\kappa\right)\right] > 0,$$

i.e., the long run level of effort increases with κ. □

Proof of Proposition 3: Note that since $\sigma_\theta > 0$, $\gamma^*(\kappa)$ defined in (13) is strictly positive for all $\kappa \in \mathbb{R}$. Let $(\nu_t)_{t\geq 0}$ denote the solution to the ordinary differential equation $\dot{\nu}_t = 2r\nu_t + \sigma_\theta^2 - \nu_t^2/\sigma_\xi^2$ with initial condition $\nu_0 = \gamma_0 > \gamma^*(r) := \sigma_\xi^2\left(\sqrt{r^2+\sigma_\theta^2/\sigma_\xi^2}+r\right)$. It is easy to see that $(\nu_t)_{t\geq 0}$ is strictly decreasing, and that it converges to $\gamma^*(r)$.

Define the function

$$\tilde{\mu}_t = \frac{\nu_t}{\sigma_\xi^2}\int_t^\infty e^{-\int_t^s \nu_u/\sigma_\xi^2 du} ds,\ t \geq 0,$$

It is straightforward to see that the Proof of Proposition 2 still goes through if $\kappa = r$. Thus, $(\tilde{\mu}_t)_{t\geq 0}$ is strictly decreasing. Also, because $(\nu_t)_{t\geq 0}$ converges, $(\tilde{\mu}_t)_{t\geq 0}$ converges to 1. As a result, $\tilde{\mu}_t > 1$ for all $t > 0$. It remains to show that, for a fixed $\gamma_0 > \gamma^*(r)$, equilibrium effort, $(a_t^*(\kappa;\gamma_0))_{t\geq 0}$, converges pointwise to $(\tilde{\mu}_t)_{t\geq 0}$ from below when $\kappa \nearrow r$.

It is easy to verify that the solution to the ODE (7) is given by

$$\gamma_t(\kappa;\gamma_0) = \frac{b - cae^{qt}}{1 - ce^{qt}}$$

where

$$a := \sigma_\xi^2\left(\kappa + \sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2}\right) > 0 \ \text{ and } \ b := \sigma_\xi^2\left(\kappa - \sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2}\right) < 0$$

(i.e., a and $-b$ are the roots of $2\kappa x + \sigma_\theta^2 - x^2/\sigma_\xi^2$), and where

$$q := \frac{a-b}{\sigma_\xi^2} \ \text{ and } \ c := \frac{\gamma_0 - b}{\gamma_0 - a}.$$

(cf. Example 6.2.13 in Oksendal, 2005).

By definition of $(\nu_t)_{t\geq 0}$, $\gamma_t(r;\gamma_0) = \nu_t$. Also, observe that since $\gamma^*(\kappa)$ is increasing in κ, we have that $\gamma_0 > \gamma^*(r)$ implies that $t \mapsto \gamma_t(\kappa;\gamma_0)$ is strictly decreasing for all $\kappa \in (-\infty, r)$.

It is clear that $\gamma_t(\kappa;\gamma_0) \to \nu_t$ when $\kappa \to r$, as c, a, b and q are all continuous in κ at $\kappa = r$. Moreover, $\gamma_t(\kappa;\gamma_0) \in [0, \gamma_0]$ for all $\kappa \in \mathbb{R}$ and $t \geq 0$. By the Dominated Convergence Theorem (DCT), for every fixed $s \geq t$,

$$\int_t^s (r + \gamma_u(\kappa;\gamma_0)/\sigma_\xi^2 - \kappa)du \to \int_t^s (r + \nu_u/\sigma_\xi^2 - r)du$$

when $\kappa \to r$, and hence, by continuity,

$$f_\kappa(s; t, \gamma_0) := e^{-\int_t^s (r+\gamma_u(\kappa;\gamma_0)/\sigma_\xi^2 - \kappa)du} \to e^{-\int_t^s \nu_u/\sigma_\xi^2 du}, \text{ as } \kappa \to r.$$

Now, observe that since $r \geq \kappa$ and $t \mapsto \gamma_t(\kappa;\gamma_0)$ is strictly decreasing, we have that

$$f_\kappa(s; t, \gamma_0) \leq e^{-\int_t^s \gamma_u(\kappa;\gamma_0)/\sigma_\xi^2 du} \leq e^{-\gamma^*(\kappa;\gamma_0)(s-t)/\sigma_\xi^2}, \; s \geq t.$$

As a result, given $\overline{\kappa} < r$ we deduce that for all $\kappa \in [\overline{\kappa}, r]$ and $s \geq t$,

$$f_\kappa(s; t, \gamma_0) \leq e^{-\gamma^*(\overline{\kappa};\gamma_0)(s-t)/\sigma_\xi^2} := g_{\overline{\kappa}}(s; t, \gamma_0).$$

as $\gamma^*(\kappa;\gamma_0)$ is increasing in κ. Moreover, $g_{\overline{\kappa}}(s; t, \gamma_0) \in L^1([t, \infty); \mathbb{R})$ due to $\gamma^*(\overline{\kappa};\gamma_0) > 0$. The DCT allows us then to conclude that

$$\int_t^\infty f_\kappa(s; t, \gamma_0)ds \to \int_t^\infty e^{-\int_t^s \nu_u/\sigma_\xi^2 du}ds.$$

The result then follows from the continuity of $\kappa \mapsto \gamma_t(\kappa;\gamma_0)$ at $\kappa = r$ noticing that

$$a_t^*(\kappa;\gamma_0) = \frac{\gamma_t^*(\kappa;\gamma_0)}{\sigma_\xi^2} \int_t^\infty f_\kappa(s; t, \gamma_0)ds.$$

□

Proof of Proposition 5: Suppose that the market conjectures that the worker will follow a deterministic strategy $a^* := (a_t^*)_{t\geq 0}$. Since wages take the form $w_t = \lambda a_t^* + p_t^*$, only $(p_t^*)_{t\geq 0}$ matters for incentives.

The solution to (6) as a function of the public history is

$$p_t^* = e^{-\int_0^t (\beta_s - \kappa)ds} p_0 + \int_0^t e^{-\int_s^t (\beta_u - \kappa)du} [\eta_s(a^{*s})ds + \beta_s(d\xi_s - \lambda a_s^* ds)], \tag{A.1}$$

where $\beta_t = \gamma_t/\sigma_\xi^2$ for all $t \geq 0$. Also, because $(\eta_t(\cdot))_{t\geq 0}$ is deterministic, the conjectured trajectory $(\eta_t(a^{*t}))$, is fixed at time zero and unaffected by the worker's effort choice. Thus, the only component of the public belief that can be affected by the worker's actions is

$$Y_t := \int_0^t e^{-\int_s^t (\beta_u - \kappa)du} \beta_s d\xi_s = \int_0^t e^{-\int_s^t (\beta_u - \kappa)du} \beta_s [(p_s + \lambda a_s)ds + \sigma_\xi dZ_s], \tag{A.2}$$

where in the last equality I used that $d\xi_t = (p_t + a_t)dt + \sigma_\xi dZ_t$ from the worker's perspective.

Similarly, it is easy to see that the worker's belief about his own skills satisfies

$$p_t = e^{-\int_0^t \kappa ds} p_0 + \int_0^t e^{\kappa(t-s)du}[\eta_s(a^s)ds + \beta_s \sigma_\xi dZ_s], \tag{A.3}$$

where $Z := (Z_t)_{t\geq 0}$ is a Brownian motion from the worker's standpoint. Inserting this into (A.2) yields

$$Y_t = \int_0^t e^{-\int_s^t (\beta_u - \kappa)du} \beta_s \left[e^{\kappa s} p_0 + \int_0^s e^{\kappa(s-u)}(\eta_u(a^u)du + \beta_u \sigma_\xi dZ_u) + \lambda a_s ds + \sigma_\xi dZ_s \right].$$

Since the first term in Y_t is unaffected by the effort decision, the worker's optimization problem is reduced to maximizing

$$\mathbb{E}\left[\int_0^\infty e^{-rt} \left(\int_0^t e^{-\int_s^t (\beta_u - \kappa)du} \beta_s \left\{ \int_0^s e^{\kappa(s-u)}(\eta_u(a^u)du + \beta_u \sigma_\xi dZ_u) + \lambda a_s ds + \sigma_\xi dZ_s \right\} - g(a_t) \right) dt \right].$$

The following lemma is a direct consequence of $r > \kappa$, and its proof deferred to the end:

Lemma 2 (Stochastic integrals vanish). *When $r > \kappa$,*

$$\begin{aligned} J &:= \mathbb{E}\left[\int_0^\infty e^{-rt} \left\{ \int_0^t e^{-\int_s^t (\beta_u - \kappa)du} \beta_s \left(\int_0^s e^{\kappa(s-u)} \beta_u \sigma_\xi dZ_u \right) ds \right\} dt \right] \\ &= \mathbb{E}\left[\int_0^\infty e^{-rt} \left(\int_0^t e^{-\int_s^t (\beta_u - \kappa)du} \beta_s dZ_s \right) dt \right] = 0. \end{aligned}$$

□

With this in hand, the problem of the worker is reduced to

$$\max_{a \in \mathcal{A}} \mathbb{E}\left[\int_0^\infty e^{-rt} \left(\int_0^t e^{-\int_s^t (\beta_u - \kappa)du} \beta_s \left\{ \int_0^s e^{\kappa(s-u)} \eta_u(a^u)du + a_s ds \right\} - g(a_t) \right) dt \right].$$

Using integration by parts

$$\int_0^t e^{-\int_s^t (\beta_u - \kappa)du} \beta_s \int_0^s e^{\kappa(s-u)} \eta_u(a^u)du = e^{\kappa t} \int_0^t e^{-\kappa s} \eta_s(a^s)ds - e^{-\int_0^t (\beta_s - \kappa)ds} \int_0^t e^{\int_0^s (\beta_u - \kappa)du} \eta_s(a^s)ds.$$

Thus, the worker's objective function has 3 integrals of the form (up to multiplicative constants)

$$I := \int_0^\infty e^{-rt}\left[e^{-\int_0^t \tau_s ds}\int_0^t e^{\int_0^s \tau_u du}\nu_s\right]dt$$

where $\tau_t = \beta_t - \kappa$ or $-\kappa$ and $\nu_t = \lambda\beta_t a_t$ or $\eta_t(a^t)$. Since in any case $r+\tau > 0$, a direct application of Fubini's theorem implies that

$$I = \int_0^\infty e^{\int_0^t \tau_t dt}\nu_t \int_t^\infty \exp^{-\int_0^s (r+\tau_u)du} dsdt = \int_0^\infty e^{-rt}\nu_t \int_t^\infty e^{-\int_t^s (r+\tau_u)du} dsdt$$

Defining $\rho_t := 1/(r-\kappa) - \mu_t/\beta_t$ where $\mu_t := \beta_t \int_t^\infty e^{-\int_t^s (r+\beta_u-\kappa)du}ds$, we conclude that the objective becomes

$$\mathbb{E}\left[\int_0^\infty e^{-rt}\left(\lambda\mu_t a_t + \rho_t\eta_t(a^t) - g(a_t)\right)dt\right]. \tag{A.4}$$

To prove the statement of the Proposition, suppose first that $(a_t^*)_{t\geq 0}$ is a deterministic equilibrium. Then, $(a_t^*)_{t\geq 0}$ must maximize (A.4). But this implies that $(a_t^*)_{t\geq 0}$ must solve (22), i.e. $\max_{a\in\mathcal{A}} \int_0^\infty e^{-rt}\left(\lambda\mu_t a_t + \rho_t\eta_t(a^t) - g(a_t)\right)dt$, because otherwise the worker could improve his payoff.

To prove the converse, suppose that $(a_t^*)_{t\geq 0}$ is not a deterministic equilibrium; in particular, when the market forms beliefs using this strategy, there is a deviation $(\hat{a}_t)_{t\geq 0}$ that is progressively measurable with respect to $(\xi_t)_{t\geq 0}$, such that

$$\mathbb{E}^{\hat{a}}\left[\int_0^\infty e^{-rt}\left(\lambda\mu_t\hat{a}_t + \rho_t\eta_t(\hat{a}^t) - g(\hat{a}_t)\right)dt\right] > \int_0^\infty e^{-rt}\left(\lambda\mu_t a_t^* + \rho_t\eta_t(a^{*t}) - g(a_t^*)\right)dt$$

where $\mathbb{E}^{\hat{a}}[\cdot]$ reflects that the expectation is taken across all possible realizations of $(\xi_t)_{t\geq 0}$ according to the measure induced by $(\hat{a}_t)_{t\geq 0}$ on the set of continuous functions from $\mathbb{R}_+$ to $\mathbb{R}$. But in this case, there must exist a set of non-zero measure of paths of $(\xi_t)_{t\geq 0}$ over which

$$\int_0^\infty e^{-rt}\left(\lambda\mu_t\hat{a}_t(\xi^t) + \rho_t\eta_t(\hat{a}^t(\xi^t)) - g(\hat{a}_t(\xi^t))\right)dt > \int_0^\infty e^{-rt}\left(\lambda\mu_t a_t^* + \rho_t\eta_t(a^{*t}) - g(a_t^*)\right)dt.$$

Fixing any path $\{\tilde{\xi}_t : t\geq 0\}$ in that set, let $t\mapsto \tilde{a}_t = \hat{a}_t(\tilde{\xi}^t)$, which is a (measurable) function of time only. Thus, $(a_t^*)_{t\geq 0}$ does not solve (22), from where the conclusion follows.

To conclude, the proof of Lemma 2. Rearranging terms in J,

$$J = \mathbb{E}\left[\int_0^\infty e^{-(r-\kappa)t}\left\{\int_0^t e^{-\int_s^t \beta_u du}\beta_s\left(\int_0^s e^{-\kappa u}\beta_u\sigma_\xi dZ_u\right)ds\right\}dt\right].$$

Now,

$$\underbrace{\int_0^t e^{-\int_s^t \beta_u du}\beta_s \mathbb{E}\left[\left|\int_0^s e^{-\kappa u}\beta_u \sigma_\xi dZ_u\right|\right] ds}_{\hat{J}_t} \leq \int_0^t e^{-\int_s^t \beta_u du}\beta_s \left(\mathbb{E}\left[\left(\int_0^s e^{-\kappa u}\beta_u \sigma_\xi dZ_u\right)^2\right]\right)^{1/2} ds$$

$$= \int_0^t e^{-\int_s^t \beta_u du}\beta_s \left(\int_0^s e^{-2\kappa u}\beta_u^2 \sigma_\xi^2 du\right)^{1/2} ds,$$

where the last equality follows from $\int_0^s e^{-\kappa u}\beta_u\sigma_\xi dZ_u \sim \mathcal{N}(0, \int_0^s e^{-2\kappa u}(\beta_u\sigma_\xi)^2 du)$. Using that $(\beta_t)_{t\geq 0}$ is bounded, there is $C > 0$ such that

$$\hat{J}_t \leq C\int_0^t (1 - e^{-2\kappa s})^{1/2} ds \leq Ct < +\infty. \tag{A.5}$$

By Fubini's theorem (Theorem 18.3 in Billingsley, 1995),

$$\mathbb{E}\left[\int_0^t e^{-\int_s^t \beta_u du}\beta_s \left(\int_0^s e^{-\kappa u}\beta_u\sigma_\xi dZ_u\right) ds\right] = \int_0^t e^{-\int_s^t \beta_u du}\beta_s \underbrace{\mathbb{E}\left[\int_0^s e^{-\kappa u}\beta_u \sigma_\xi dZ_u\right]}_{=0} ds = 0. \tag{A.6}$$

Thus, it remains to show that the first two integrals in J can be interchanged, i.e., that

$$J = \int_0^\infty \mathbb{E}\left[e^{-(r-\kappa)t}\int_0^t e^{-\int_s^t \beta_u du}\beta_s \left(\int_0^s e^{-\kappa u}\beta_u\sigma_\xi dZ_u\right) ds\right] dt, \tag{A.7}$$

as $J = 0$ would then follow from (A.6). To this end, observe that

$$\begin{aligned}
K_0 &:= \int_0^\infty \mathbb{E}\left[\left|e^{-(r-\kappa)t}\int_0^t e^{-\int_s^t \beta_u du}\beta_s\left(\int_0^s e^{-\kappa u}\beta_u\sigma_\xi dZ_u\right) ds\right|\right] dt \\
&\leq \int_0^\infty \mathbb{E}\left[e^{-(r-\kappa)t}\int_0^t e^{-\int_s^t \beta_u du}\beta_s\left|\int_0^s e^{-\kappa u}\beta_u\sigma_\xi dZ_u\right| ds\right] dt \\
&= \int_0^\infty e^{-(r-\kappa)t}\underbrace{\int_0^t e^{-\int_s^t \beta_u du}\beta_s \mathbb{E}\left[\left|\int_0^s e^{-\kappa u}\beta_u\sigma_\xi dZ_u\right| ds\right]}_{=J_t} dt,
\end{aligned} \tag{A.8}$$

where the last equality follows from Tonelli's theorem. Using that $\hat{J}_t \leq Ct$ and $r > \kappa$, it follows that $K_0 < \infty$, and hence, from Fubini's theorem, (A.7) holds. Finally, that the other expectation is also zero follows from identical arguments. This concludes the proof. □

Proof of Proposition 7: Observe first that

$$\mathrm{Cov}[L_t, Y_t] = \int_{s\leq t} e^{-\phi(\beta^*-\kappa)(t-s)}\mathrm{Cov}[L_t, L_s]ds.$$

But $\mathrm{Cov}[L_t, L_s] = -\sigma_\theta^2 e^{\kappa|t-s|}/2\kappa$ for a stationary O-U process (recall that $\kappa < 0$), which yields

$$\mathrm{Cov}[L_t, Y_t] = -\frac{\sigma_\theta^2}{2\kappa[\phi(\beta^*-\kappa)-\kappa]}.$$

Now, by independence

$$\mathrm{Var}[Y_t] = \underbrace{\mathrm{Var}\Big[\int_{s\leq t} e^{-\phi(\beta^*-\kappa)(t-s)}L_s ds\Big]}_{\mathbb{V}:=} + \underbrace{\mathrm{Var}\Big[\int_{s\leq t} e^{-\phi(\beta^*-\kappa)(t-s)}\sigma_\xi dZ_t^\xi\Big]}_{=\frac{\sigma_\xi^2}{2\phi(\beta^*-\kappa)}}$$

However,

$$\begin{aligned}
\mathbb{V} &= \int_{s\leq t} e^{-\phi(\beta^*-\kappa)(t-s)} \int_{u\leq t} e^{-\phi(\beta^*-\kappa)(t-u)}\mathrm{Cov}[L_s, L_u]duds \\
&= -\frac{\sigma_\theta^2}{2\kappa}\int_0^\infty e^{-\phi(\beta^*-\kappa)t}\left[\int_0^t e^{-\phi(\beta^*-\kappa)s}e^{\kappa(t-s)}ds + \int_t^\infty e^{-\phi(\beta^*-\kappa)s}e^{\kappa(s-t)}ds\right]dt \\
&= -\frac{\sigma_\theta^2}{2\kappa}\left[\frac{1}{[\phi(\beta^*-\kappa)+\kappa][\phi(\beta^*-\kappa)-\kappa]} - \frac{1}{2\phi(\beta^*-\kappa)[\phi(\beta^*-\kappa)+\kappa]}\right. \\
&\quad \left. + \frac{1}{2\phi(\beta^*-\kappa)[\phi(\beta^*-\kappa)-\kappa]}\right] \\
&= -\frac{\sigma_\theta^2}{2\kappa}\frac{2\phi(\beta^*-\kappa) - [\phi(\beta^*-\kappa)-\kappa] + [\phi(\beta^*-\kappa)+\kappa]}{2\phi(\beta^*-\kappa)[\phi(\beta^*-\kappa)-\kappa][\phi(\beta^*-\kappa)+\kappa]} \\
&= -\frac{\sigma_\theta^2}{2\kappa}\frac{2[\phi(\beta^*-\kappa)+\kappa)}{2\phi(\beta^*-\kappa)[\phi(\beta^*-\kappa)-\kappa][\phi(\beta^*-\kappa)+\kappa]} \\
&= -\frac{\sigma_\theta^2}{2\kappa\phi(\beta^*-\kappa)[\phi(\beta^*-\kappa)-\kappa]}
\end{aligned}$$

So,

$$\mathrm{Var}[Y_t] = \frac{1}{2\phi(\beta^*-\kappa)}\left[\sigma_\xi^2 - \frac{\sigma_\theta^2}{\kappa[\phi(\beta^*-\kappa)-\kappa]}\right]$$

Using that $\sigma_\theta^2 = \sigma_\xi^2[(\beta^* - \kappa)^2 - \kappa^2]$,

$$\begin{aligned}
\mathrm{Var}[Y_t] &= \frac{\sigma_\xi^2}{2\phi(\beta^* - \kappa)}\left[\frac{\kappa\phi(\beta^* - \kappa) - \kappa^2 - (\beta^* - \kappa)^2 + \kappa^2}{\kappa[\phi(\beta^* - \kappa) - \kappa]}\right] = \frac{\sigma_\xi^2}{2\kappa\phi}\frac{\phi\kappa - \beta^* + \kappa}{\phi(\beta^* - \kappa) - \kappa} \\
\Rightarrow \alpha &= \frac{\mathrm{Cov}[Y_t, L_t]}{\mathrm{Var}[Y_t]} = \frac{\sigma_\theta^2 \phi}{\sigma_\xi^2[\beta^* - \kappa - \phi\kappa]} > 0.
\end{aligned}$$

Regarding incentives, note that, from the perspective of the worker,

$$\begin{aligned}
\mathbb{E}\left[\int_0^\infty e^{-rt}\mu_t dt\right] &= \mathbb{E}\left[\int_0^\infty e^{-rt}\alpha\left(\int_{s\le t} e^{-\phi(\beta^*-\kappa)(t-s)}[d\xi_s - A_s^* ds]\right)dt\right] \\
&= \mathbb{E}\left[\int_0^\infty e^{-rt}\alpha\left(\int_{s\le t} e^{-\phi(\beta^*-\kappa)(t-s)}[A_s + \mathbb{E}[L_s|\mathcal{F}_s] - A_s^*]ds\right)dt\right]
\end{aligned}$$

where $(\mathbb{E}[L_s|\mathcal{F}_s])_{s\ge 0}$ is exogenous and $\mathcal{F}_t$ is, as in the baseline model, the σ-algebra generated by $(\xi_s)_{s\le t}$. Thus, from the point of view of incentives, the only relevant part to consider is

$$I := \mathbb{E}\left[\int_0^\infty e^{-rt}\alpha\left(\int_0^t e^{-\phi(\beta^*-\kappa)(t-s)}\left\{\int_0^s e^{\kappa(s-u)}a_u du\right\}ds\right)dt\right] \tag{A.9}$$

as the worker's investment problem starts at time zero (which modifies $\int_{s\le t}$ to $\int_0^t$). Now,

$$\begin{aligned}
I &= \mathbb{E}\left[\int_0^\infty e^{-[r+\phi(\beta^*-\kappa)]t}\alpha\left(\int_0^t e^{[\phi(\beta^*-\kappa)+\kappa]s}\left\{\int_0^s e^{-\kappa u}a_u du\right\}ds\right)dt\right] \\
&= \mathbb{E}\left[\int_0^\infty e^{-[r+\phi(\beta^*-\kappa)]t}\frac{\alpha}{\phi(\beta^*-\kappa)+\kappa}\left(e^{[\phi(\beta^*-\kappa)+\kappa]t}\int_0^t e^{-\kappa s}a_s ds - \int_0^t e^{\phi(\beta^*-\kappa)s}a_s ds\right)dt\right] \\
&= \mathbb{E}\left[\int_0^\infty e^{-rt}\frac{\alpha a_s}{\phi(\beta^*-\kappa)+\kappa}\left(\int_t^\infty e^{-(r-\kappa)(s-t)}ds - \int_0^t e^{-[r+\phi(\beta^*-\kappa)](s-t)}ds\right)dt\right] \\
&= \mathbb{E}\left[\int_0^\infty e^{-rt}\alpha a_s \frac{1}{(r-\kappa)(r+\phi(\beta^*-\kappa))}dt\right]
\end{aligned}$$

Which implies that the optimal investment policy is stationary according to

$$g'(a^*) = \frac{\alpha}{(r-\kappa)(r+\phi(\beta^*-\kappa))}$$

To find the optimal ϕ^*, observe that $\alpha = \alpha(\phi) = \sigma_\theta^2 \phi / \sigma_\xi^2 [\beta^* - \kappa - \phi\kappa]$ yields

$$\begin{aligned} \frac{d}{d\phi}\left(\frac{\alpha}{(r-\kappa)(r+\phi(\beta^*-\kappa))}\right) &= 0 \\ \Leftrightarrow \sigma_\theta^2[\beta^*-\kappa-\phi\kappa][r+\phi(\beta^*-\kappa)] - \sigma_\theta^2\phi\{(\beta^*-\kappa)[\beta^*-\kappa-\kappa\phi] - \kappa[r+\phi(\beta^*-\kappa)]\} &= 0 \\ \Leftrightarrow (\beta^*-\kappa)[r+\phi(\beta^*-\kappa)] - \phi(\beta^*-\kappa)[\beta^*-\kappa-\kappa\phi] &= 0 \\ \Rightarrow \phi^* &= \sqrt{\frac{r}{|\kappa|}} \end{aligned}$$

This is clearly a maximum, as

$$\frac{d}{d\phi}\left(\frac{\alpha}{(r-\kappa)(r+\phi(\beta^*-\kappa))}\right) = \frac{(\beta^*-\kappa)[r+\phi^2\kappa]}{[(r-\kappa)(r+\phi(\beta^*-\kappa))]^2} > 0 \ \text{ iff } \ \phi < \phi^*.$$

Finally,

$$\begin{aligned} g'(a^*(\kappa)) &= \frac{\alpha}{(r-\kappa)(r+\phi^*(\beta-\kappa))} \\ &= \frac{1}{(r-\kappa)}\frac{\sigma_\theta^2}{\sigma_\xi^2(\beta^*-\kappa-\phi^*\kappa)}\frac{\phi^*}{(r+\phi^*(\beta^*-\kappa))} \\ &= \frac{1}{(r-\kappa)}\frac{\sigma_\theta^2}{\sigma_\xi^2(\sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2}+\sqrt{r|\kappa|})}\frac{\sqrt{r}}{(r\sqrt{|\kappa|}+\sqrt{r}\sqrt{\kappa^2+\sigma_\theta^2/\sigma_\xi^2})} \nearrow \frac{1}{r}, \end{aligned}$$

as $\kappa \nearrow 0$, concluding the proof. □

B Appendix B: Wages

This Appendix establishes the properties of the distribution of wages generated by the general model of Section 4.1. In particular, human capital accumulation in conjunction with learning are able to replicate two well documented empirical regularities of the time series of wages. First, average wages (or earnings) tend to be increasing and concave over the life cycle (e.g., Becker, 1975, chapter VII). Second, the dispersion of wages within cohorts tends to increase with experience (Neal and Rosen, 2000).[33]

[33] Two models that deliver profiles of earnings with properties that are similar to the ones previously mentioned are Ben-Porath (1964) and Jovanovic (1979). In the first model, productivity is observable and the worker invests in skills. As in the model I examine, average wages increase with experience due to human capital accumulating over time. Wage dispersion that increases with cohort experience instead arises from heterogeneity in the skill-accumulation technology. In Jovanovic's model, skills are exogenous and hidden. In addition, he allows for turnover, and hence experience and tenure do not coincide. As in the current model, wage dispersion that increases with cohort experience results from learning. The model instead predicts that average wages increase with *tenure*, as matches that survive longer are precisely the highly productive ones.

Proposition 8 (Distribution of wages)**.** *In the general model defined by (1), (21) and* $\lambda \in [0,1]$*, the wage process is Gaussian with mean and covariance*

$$\mathbb{E}[w_t] = \lambda a_t^* + e^{\kappa t} p_0 + \int_0^t e^{\kappa(t-s)} \eta_s(a^{*s}) ds, \; t \geq 0, \; \textit{and} \tag{B.1}$$

$$\mathrm{Cov}(w_t, w_s) = e^{\kappa(t+s)} \int_0^s e^{-2\kappa u} \frac{\gamma_u^2}{\sigma_\xi^2} du, \; t \geq s. \tag{B.2}$$

In particular, $\mathrm{Var}(w_t)$ *is increasing over time.*

Because the model under study allows for learning, equilibrium wages are perfectly aligned with the second empirical regularity: since the market gradually distinguishes between workers who have suffered good and bad shocks, wages becomes more dispersed over time as workers are rewarded based on their past performance.

To see how the model can generate the first regularity, consider the baseline specification $\eta_t(a^t) = (1-\lambda)a_t$, in its pure human capital version ($\lambda = 0$), with skills evolving as a martingale in the absence of effort ($\kappa = 0$). If information accumulates over time, therefore,

$$\frac{d\mathbb{E}[w_t]}{dt} = a_t^* > 0 \;\text{ and }\; \frac{d^2\mathbb{E}[w_t]}{dt^2} = \frac{da_t^*}{dt} < 0.$$

By continuity, the same conclusion holds in a neighborhood of $(\lambda, \kappa) = (0, 0)$.

Proof of Proposition 8: The expressions for $\mathbb{E}[w_t]$ and $\mathrm{Cov}(w_t, w_s)$ follow directly from: (i) $w_t = \lambda a_t^* + p_t^*$; (ii) $(p_t^*)_{t\geq 0}$ satisfying

$$p_t^* = e^{\kappa t} p_0 + \int_0^t e^{\kappa(t-s)} \eta_s(a^{*s}) ds + \int_0^t e^{\kappa(t-s)} \frac{\gamma_s}{\sigma_\xi} dZ_s, \; t \geq 0,$$

from a time-zero perspective; (iii) $\mathbb{E}\left[\int_0^t e^{\kappa(t-s)}\gamma_s/\sigma_\xi dZ_s\right] = 0$, $t \geq 0$, in equilibrium; and from (iv)

$$\begin{aligned} \mathrm{Cov}\left(\int_0^t e^{\kappa(t-u)} \frac{\gamma_u}{\sigma_\xi} dZ_u, \int_0^s e^{\kappa(s-u)} \frac{\gamma_u}{\sigma_\xi} dZ_u\right) &= e^{\kappa t} e^{\kappa s} \mathbb{V}\left(\int_0^s e^{-\kappa u} \frac{\gamma_u}{\sigma_\xi} dZ_u\right) \\ &= e^{\kappa(t+s)} \int_0^s \left(e^{-\kappa s} \frac{\gamma_u}{\sigma_\xi}\right)^2 du, \; t \geq s, \end{aligned}$$

Moreover, once controlling for tenure, there are no returns to experience. Recently, Bonatti and Hörner (2017) have also made progress in this respect by showing that complementarities and coarse information structures can lead to wages that are single peaked, a feature documented by Baker, Gibbs and Holmström (1994) who analyzed data on wages coming from a single firm.

where the first equality is due to the independent increments of $(Z_t^*)_{t\geq 0}$, and the second follows from the stochastic integral being of the Ito type.

To conclude, it is immediate that $t \mapsto \mathbb{V}[w_t]$ is increasing when $\kappa \geq 0$. If $\kappa < 0$, using that $(\gamma_t)_{t\geq 0}$ is decreasing we deduce

$$\frac{d\text{Var}[w_t]}{dt} = 2\kappa e^{2\kappa t} \int_0^t e^{-2\kappa s} \frac{\gamma_s^2}{\sigma_\xi^2} ds + \frac{\gamma_t^2}{\sigma_\xi^2} > 2\kappa e^{2\kappa t} \frac{\gamma_t^2}{\sigma_\xi^2} \int_0^t e^{-2\kappa s} ds + \frac{\gamma_t^2}{\sigma_\xi^2} = \frac{e^{2\kappa t}\gamma_t^2}{\sigma_\xi^2} > 0.$$

□