

IZA DP No. 5323

Dynamic Incentive Contracts under Parameter Uncertainty

Julien Prat
Boyan Jovanovic

November 2010

Dynamic Incentive Contracts under Parameter Uncertainty

Julien Prat

*IAE (CSIC), Barcelona GSE,
CEPR and IZA*

Boyan Jovanovic

New York University

Discussion Paper No. 5323
November 2010

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Dynamic Incentive Contracts under Parameter Uncertainty*

We analyze a long-term contracting problem involving common uncertainty about a parameter capturing the productivity of the relationship, and featuring a hidden action for the agent. We develop an approach that works for any utility function when the parameter and noise are normally distributed and when the effort and noise affect output additively. We then analytically solve for the optimal contract when the agent has exponential utility. We find that the Pareto frontier shifts out as information about the agent's quality improves. In the standard spot-market setup, by contrast, when the parameter measures the agent's "quality", the Pareto frontier shifts inwards with better information. Commitment is therefore more valuable when quality is known more precisely. Incentives then are easier to provide because the agent has less room to manipulate the beliefs of the principal. Moreover, in contrast to results under one-period commitment, wage volatility declines as experience accumulates.

JEL Classification: D82, D83, E24, J41

Keywords: principal-agent model, optimal contract, learning, private information, reputation, career

Corresponding author:

Julien Prat
Institute of Economic Analysis (IAE-CSIC)
Campus UAB
08193 Bellaterra, Barcelona
Spain
E-mail: julien.prat@iae.csic.es

* We thank Heski Bar-Isaac, Pieter Gautier, Robert Kohn, Albert Marcet, Ennio Stacchetti as well as seminar participants at the IMT Institute in Lucca, the Tinbergen Institute, University of Cergy-Pontoise, University of Essex, Stockholm Institute for International Economic Studies (IIES), and the Minnesota Macro Workshop for comments and suggestions. We acknowledge the support of the Kauffman Foundation and of the Barcelona GSE.

1 Introduction

Agency relationships often preclude complete monitoring so that a principal cannot observe the agent's actions. Other features of the environment, such as the manager's ability, the quality of his match with the firm, or the profitability of the project under management, can also be a source of uncertainty. Many relationships between firms and workers, as well as between lenders and borrowers, are of this general form. Yet, little is known about how parameter and effort uncertainty interact to shape the optimal design of incentive contracts. Does parameter uncertainty reinforce or alleviate moral hazard concerns? Does it render commitment more or less valuable?

This paper provides some answers to these questions by focusing on cases where: (i) the unknown parameter remains constant over time; and (ii) a risk neutral principal and a risk averse agent commit to a long-term contract. Under full-commitment, incentives are designed to reward effort and not ability. Disentangling the two is not always feasible for the principal because they both influence his only source of information, i.e., realized revenues. Signal confusion enables the agent to manipulate the principal's beliefs. If the agent shirks (i.e., provides less effort than recommended), output will be below expectation and the principal will infer that the match productivity is lower than he had thought. The agent, on the other hand, knows that low output was caused not by low productivity but by low effort and so, after shirking, is more optimistic about the value of the unknown parameter than the principal.

Compared to the situation in which all parameters are known, a given indexation of future earnings to performance entails lower punishments for shirkers. By inducing the principal to underestimate the match productivity, a shirker knows that he will benefit in the future from overestimated inferences about his effort and thus higher rewards. In order to prevent such *belief manipulation*, a long-term contract under parameter uncertainty must entail a higher indexation to performance. This raises income volatility, which lowers the welfare of the risk-averse agent. Moreover, if the unknown parameter is constant, belief manipulation is more effective early on in the relationship because posteriors put higher weight on new information. This is why the sensitivity of pay to performance declines over time.¹

These implications stand in sharp contrast to the ones derived in the literature on *career concerns* where the unknown parameter measures the agent's general ability, transferable from job to job. Analyzing this class of problems under spot markets with up-front pay only, Holmström (1999) concludes that incentives are more easily

¹Given a level of lifetime utility that the contract promises him, it is in the agent's interest to bias downward the principal's belief about his ability. This was also a feature in the ratchet effect model of Laffont and Tirole (1988). They consider an environment with adverse selection so that the agent knows the actual productivity. By contrast, asymmetric beliefs do not exist from the outset in our model but arise *endogenously*, and then only off the equilibrium path.

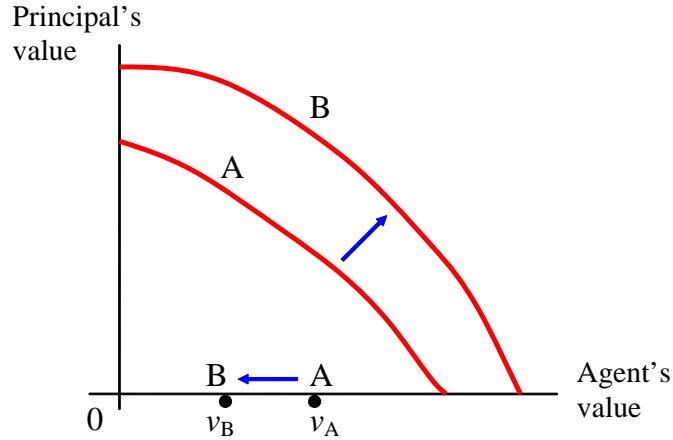


Figure 1: THE EFFECT OF A RISE IN PRIOR PRECISION UNDER SPOT MARKETS AND UNDER COMMITMENT.

provided when the agent's reputation is not established. Agents will generally exert inefficient levels of effort. At first, effort may exceed its first-best level as the agent seeks to build his reputation, but effort diminishes over time, dwindling monotonically to zero. Thus career concerns in competitive markets do not restore correct incentives on the part of agents.² Because of the convexity of the effort-disutility term, as the agent's effort declines, so do his rents. In other words, better information about the agent's quality *reduces* his equilibrium utility.

Figure 1 illustrates the effect that higher precision of information about the agent's quality has on the welfare of the parties. Under spot contracts, competition for the agent's services ensures that the principal earns zero profits, and so we are on the horizontal axis. Starting at a point on the horizontal axis where the agent's value is v_A , a rise in information about the agent's quality leaves the principal's welfare unchanged at zero, but reduces the agent's welfare from v_A to v_B , as illustrated by the arrow pointing to the left on the horizontal axis.

For reasons discussed above, the opposite happens under full commitment. The spot contract is feasible but is generally suboptimal, and therefore the utilities that it generates are strictly inside the Pareto frontier. When we raise precision about the agent's quality, there is less room for belief manipulation and the contract curve

²Holmström assumed that the agent was risk neutral. As shown in Section 6, introducing risk aversion makes little difference to the solution of the problem: Effort still converges monotonically to zero. On the other hand, the contracting problem under risk neutrality becomes trivial: Even one-period contracts with pay for performance can achieve first best. More generally, a contract can attain first-best levels of effort by transferring all uncertainty to the risk neutral agent and effectively selling the project to him.

shifts out, as illustrated by the arrow pointing up and to the right. In contrast to spot markets then, better information *raises* utility and pushes the Pareto frontier out. Consequently, the value of commitment is higher when information about quality is more precise.

Analyzing models with commitment and belief divergence entails the following technical issue: Each deviation drives a permanent wedge between the agent's and the principal's posteriors. As the duration of the relationship increases, the state space is in general unbounded because the entire history of actions matters for evaluating the agent's options off the equilibrium path. Models where the noise is Markovian contain our assumptions about parameter uncertainty as a special case when the persistence becomes infinite and where the initial value is unknown with a common prior attached to it. In that case the unknown parameter is the initial condition of the process. Fernandes and Phelan (2000) or Williams (2008) study such Markovian processes but they assume that the initial value is public knowledge. A recursive approach to the problem would generally need to take beliefs of the agent and beliefs of the principal as separate states. This, broadly speaking, is the approach Fernandes and Phelan (2000) proposed. Unfortunately, it implies that the state space grows with the number of potential deviations and is therefore ill-suited to solving our problem where information persistence extends over several periods and actions are defined over a continuum.

We rely instead on a first-order approach, meaning that we focus on the equilibrium path and establish necessary condition for recommended effort to be optimal. The difficulty with this solution method is that it may identify contracts that are not implementable because the concavity of the agent's objective function is not guaranteed. Sufficient conditions have been established in the static case by Rogerson (1985). Similar results in dynamic environments are not known. One remedy is to numerically check the implementability of the solution, as in Abraham and Pavoni (2008). To the best of our knowledge, the only proof in discrete time is by Kapicka (2006) and is rather specific to the reporting problem analyzed in his paper. Hopenhayn and Jarque (2007) also analyze persistence in a principal-agent model under the assumption that the effort decision occurs solely in the first period, whereas Jarque (2008) assumes that the probability distribution over future output depends positively on a weighted sum of past efforts.

To establish implementability, we cast our problem in continuous time. This allows us to derive a parameter restriction under which recommended effort meets both necessary and sufficient conditions of the agent. The proof relies on the concavity of the agent's Hamiltonian, a strategy that was initially applied by Schättler and Sung (1993) to continuous time contracts without persistent information. Williams (2008, 2009) extends their methodology to incentives contracts with hidden savings or reporting problems with persistent information. Our analysis shares many similarities

with his approach. It differs in that we have to model the learning process and thus need to introduce contract duration as a state. Furthermore, we implement a different proof strategy based on the work of Cvitanić *et al.* (2009).

A burgeoning literature illustrates the advantages of using continuous time methods to analyze dynamic contracts, such as Sannikov (2008) though his model does not feature learning. A series of recent papers on learning and dynamic incentives is even more closely related to our work. Adrian and Westerfield (2009) analyze a dynamic contracting model in which principal and agent disagree about the resolution of uncertainty. They avoid complications linked to private information by assuming that agent's posteriors are common knowledge so that the two parties agree to disagree. Giat *et al.* (2010) extend the model of Holmström and Milgrom (1987) by also allowing initial beliefs to be asymmetric. They focus on contracts specifying a single transfer at the end of the predetermined contracting horizon whereas our setting allows transfers to be made throughout the relationship. Finally, DeMarzo and Sannikov (2008) characterize continuous-time contracts when the agent's quality varies over time and is autocorrelated. On the one hand, our set-up is more specific since we focus on cases where the unknown state remains constant through time and the agent liability is not limited. On the other hand, we introduce risk aversion on the agent's side. Hence, whereas the main insights in DeMarzo and Sannikov (2008) are related to the optimal separation policy, our paper focuses on the incentive-insurance trade-off.

The paper is structured as follows. Section 2 lays out the model's set-up. In section 3, we derive the agent's necessary and sufficient conditions. Then we solve for the optimal contract under exponential utility in Section 4. We propose a closed form solution for the principal's rent and optimal wage schedule. The properties of the optimal contract are discussed in Section 5. Section 6 contrasts the full-commitment with the spot wages solution of Holmström (1999) and the solution under partial commitment of Gibbons and Murphy (1992). Section 7 sums up our findings whereas the proofs of the main Propositions and Corollaries are in Appendix A. We relegate the proofs of some tangential claims to Appendix B, and describe in Appendix C our simulation procedure.

2 The environment

The production process.— Let $\{B_t\}_{t \geq 0}$ be a standard Brownian Motion on a probability space (Ω, \mathcal{F}, P) . The cumulative output Y_t of a match of duration t is observed by both parties and satisfies the stochastic integral equation

$$Y_t = \int_0^t (\eta + a_s) ds + \int_0^t \sigma dB_s . \quad (1)$$

The *time-invariant* productivity is denoted by η whereas $a_t \in [0, 1]$ is the effort provided by the agent. The agent's action thus shifts average output but does not directly affect its volatility.

Learning.— No one knows η at the outset, and common priors are normal with mean m_0 and precision h_0 . Posteriors over η depend on Y_t and on cumulative effort $A_t \triangleq \int_0^t a_s ds$. Conditional on (Y_t, A_t, t) , they are also normal with mean

$$\hat{\eta}(Y_t - A_t, t) \triangleq E_t[\eta | Y_t, A_t] = \frac{h_0 m_0 + \sigma^{-2}(Y_t - A_t)}{h_t}, \quad (2)$$

and with precision

$$h_t \triangleq h_0 + \sigma^{-2}t. \quad (3)$$

Focusing on normal priors over the mean of a normally distributed process enables us to summarize all the statistically significant information by just three variables: cumulative output Y , cumulative effort A and elapsed time t . Especially useful for the characterization of optimal contracts is the fact that beliefs depend on the history of a through A alone. Hence it is sufficient to keep track of cumulative effort instead of the whole effort path.

Preferences.—The agent is risk averse and cannot borrow and lend. For all $t \geq 0$ and any given event $\omega \in \Omega$, we define a wage function $w : \mathbb{R}^+ \times \Omega \rightarrow \mathbb{R}$. The agent preferences as of time 0 read

$$\mathcal{U}_0 \triangleq \int_0^\infty e^{-\rho t} U(w_t(\omega), a_t) dt, \quad (4)$$

with $\rho > 0$. Our specification of wages is quite general since they can depend on the entire past and present $\{Y_s; 0 \leq s \leq t\}$ of the output process.

The principal is risk neutral and seeks to maximize output net of wages. His inter-temporal preferences are

$$\pi_0 \triangleq \int_0^\infty e^{-\rho t} (dY_t - w_t(\omega)) dt, \quad (5)$$

where we have imposed a common discount rate for the agent and principal.

Long-term contract.—We assume that the parties are able to commit to a long-term contract that can depend on realized history in an arbitrary way. We follow the usual practice of adding *recommended effort* a^* to the contract definition. Accordingly, since a given output path is a random element of the space Ω , a contract is a mapping $(w, a^*) : \mathbb{R}^+ \times \Omega \rightarrow \mathbb{R} \times [0, 1]$ that associates at each time t a wage-effort pair to any output path. The mapping must be measurable based on information that the principal has, and so, can depend on past output but not on past effort.

Otherwise contracts remain general since they can depend on the entire sample path $\{Y_s; 0 \leq s \leq t\}$ of the output process.³

The principal's beliefs.—The principal assumes that the agent always takes his equilibrium action a_t^* . His beliefs are governed by (2) in which $A = A^*$ and by (3).

The agent's beliefs.—The agent's beliefs incorporate the actual level of effort a which only he knows. Thus his beliefs are governed by (2) in which A and not A^* enters. Let $\mathcal{F}_t^a \triangleq \sigma(Y_s, a_s; 0 \leq s \leq t)$ denote the filtration generated by (Y, a) and $\mathbb{F}^a \triangleq \{\mathcal{F}_t^a\}_{t \geq 0}$ the P -augmentation of this natural filtration. Denote by Z_t the cumulative surprise of someone who believes that Y_t was accompanied by the effort sequence $\{a_s; 0 \leq s \leq t\}$. The filtering theorem of Fujisaki et al. (1972) implies that the *innovation process*

$$dZ_t \triangleq \frac{1}{\sigma} [dY_t - (\hat{\eta}(Y_t - A_t, t) + a_t)dt] \quad (6)$$

is a standard Brownian motion on the probability space $(\Omega, \mathcal{F}^a, P)$.⁴ Moreover, $\hat{\eta}$ is a P -martingale⁵ with decreasing variance:

$$d\hat{\eta}(Y_t - A_t, t) = \frac{\sigma^{-1}}{h_t} dZ_t. \quad (7)$$

The agent is restricted to the class of control processes $\mathcal{A} \triangleq \{a : \mathbb{R}^+ \times \Omega \rightarrow [0, 1]\}$ that are \mathbb{F}^a -predictable.⁶ Given that the principal does not observe actual effort a , the information available to him is restricted to the filtration $\mathcal{F}_t^Y \triangleq \sigma(Y_s; 0 \leq s \leq t)$ generated by Y whose augmentation we denote by $\mathbb{F}^Y \triangleq \{\mathcal{F}_t^Y\}_{t \geq 0}$. An effort path is an equilibrium path when recommended and actual effort do coincide, i.e., if $a_t = a_t^*$ for all (t, ω) .

³Given the diffusion property of the output process, one should think of $\Omega = C([0, T]; \mathbb{R})$ as the space of continuous functions $\omega : [0, T] \rightarrow \mathbb{R}$ and of the process defined in (6) $Z_t(\omega) = \omega(t)$, $0 \leq t \leq T$, as the coordinate mapping process with Wiener measure P on $(\Omega, \mathcal{F}_t^Y)$. Accordingly a contract is a mapping $(w, a^*) : \mathbb{R}^+ \times C([0, T]; \mathbb{R}) \rightarrow \mathbb{R} \times [0, 1]$.

⁴As shown in Section 10.2. of Kallianpur (1980), the linearity of the filtering problem implies that the filtrations generated by the output and innovation processes coincide. More formally, for $\mathcal{F}_t^Z \triangleq \sigma(Z_s; 0 \leq s \leq t)$, we have $\mathcal{F}_t^a = \mathcal{F}_t^Z$.

⁵The equality follows directly from Ito's lemma. Let $X_t \triangleq Y_t - A_t$ denote cumulative output net of cumulative effort so that

$$d\hat{\eta}(X_t, t) = \frac{\partial \hat{\eta}(X_t, t)}{\partial t} dt + \frac{\partial \hat{\eta}(X_t, t)}{\partial X_t} dX_t = -\frac{\sigma^{-2}}{h_t} \hat{\eta}(X_t, t) + \frac{\sigma^{-2}}{h_t} (\hat{\eta}(X_t, t) + \sigma dZ_t) = \frac{\sigma^{-1}}{h_t} dZ_t.$$

⁶A mapping is predictable when it is \mathcal{P} -measurable, with \mathcal{P} denoting the σ -algebra of predictable subsets of the product space $\mathbb{R}^+ \times \Omega$, i.e. the smallest σ -algebra on $\mathbb{R}^+ \times \Omega$ making all left-continuous and adapted processes measurable.

3 Incentive compatibility and implementability

This section focuses on the agent's problem. We derive the necessary conditions for a given action to be optimal and then establish a restriction under which they are also sufficient. We impose a terminal date T on the contracting horizon. Until then, both principal and agent are fully committed to the relationship. The agent's continuation value at time t reads

$$v_t \triangleq \max_{a \in \mathcal{A}} E \left[\int_t^T e^{-\rho(s-t)} U(w(\bar{Y}_s), a_s) ds + e^{-\rho(T-t)} W(\bar{Y}_T) \middle| \mathcal{F}_t^a \right], \quad (8)$$

where the output path is denoted by $\bar{Y}_t \triangleq \{Y_s; 0 \leq s \leq t\}$ and $W(\cdot)$ is the terminal utility which depends on output history.⁷ The agent computes his continuation value by taking a conditional expectation under the filtration \mathcal{F}_t^a which varies with the level of cumulative effort. The principal, however, does not observe actual actions. Thus he needs to keep track of continuation values for any potential level of cumulative effort. We shall simplify the problem by adopting a first order approach: We focus on the continuation value along the equilibrium path and then establish conditions under which our solution is indeed globally optimal.

3.1 Necessary conditions for the agent's problem

The optimization problem (8) cannot be analyzed with standard methods because the objective function depends on the process w_t which is *non-Markovian*. We instead use a martingale approach. Faced with a contract w , the agent controls the distribution of w_t through his choice of effort. Under this interpretation, the agent chooses the probability measure over realizations of w_t . The Radon–Nikodym derivative associated with any effort path is a Markovian process, and so this approach makes our optimization problem treatable with optimal control techniques.⁸

⁷Since we shall let $T \rightarrow \infty$, we have assumed a tractable form for W . It is straightforward to let W also depend on cumulative effort A . Then one would have to adjust the stochastic process p defined in equation (13) as follows

$$p_t = E \left[- \int_t^T e^{-\rho(s-t)} \gamma_s \frac{\sigma^{-2}}{h_s} ds + e^{-\rho(T-t)} W_A(\bar{Y}_T, A_T) \middle| \mathcal{F}_t^a \right].$$

Apart from that, our results hold with few or no changes. The specification of the terminal utility would matter if we were to focus on repeated contracts, with W capturing the agent's outside option and the ability of the principal to reward him at the end of the relationship. We do not consider such generalizations because this paper focuses on the limit situation where both parties are forever committed. Then, as long as standard transversality conditions hold, the specification of the terminal utility is immaterial to the analysis.

⁸A more concise way to formulate the advantages of the martingale approach is to observe that the control is not anymore closed loop but instead open loop with respect to the output process.

The idea of applying this approach to principal-agent models goes back to Mirrlees (1974). Our problem is complicated by the learning mechanism as past efforts affect not only current wages but also future expectations. We show in the Appendix how this difficulty can be handled through an extension of the proof by Cvitanić *et al.* (2009) which leads to the necessary condition stated below.⁹

Proposition 1 *There exists a unique decomposition for the agent's continuation value*

$$dv_t = [\rho v_t - U(w_t, a_t)] dt + \gamma_t \sigma dZ_t, \quad (9)$$

$$v_T = W(Y_T), \quad (10)$$

where γ is a square integrable predictable process. The necessary condition for a^* to be an optimal control reads

$$\left[\gamma_t + E_t \left[- \int_t^T e^{-\rho(s-t)} \gamma_s \frac{\sigma^{-2}}{h_s} ds \right] + U_a(w_t, a^*) \right] (a - a^*) \leq 0, \quad (11)$$

for all $a \in [0, 1]$.

An increase in current effort has two effects: it raises the promised value along the equilibrium path and increases cumulative effort. The first effect is proportional to the process γ which measures the sensitivity of the agent's value to output surprises. The second effect is captured by the expectation term in (11). This term vanishes when η is known, since then $\sigma^{-2}/h_s = 0$ for all $s \geq t$. As a special case of our model, we then get the necessary condition in Sannikov (2008) which says that an optimal control must maximize the expected change in continuation value minus the marginal cost of effort.

Introducing parameter uncertainty leads to the addition of the expected future sensitivities weighted by their precision ratios because they capture the marginal impact of current effort on expected earnings. To see this, observe first that $\partial \hat{\eta}(Y_s - A_s, a) / \partial a_t = -\sigma^{-2}/h_s$ for all $s \geq t$. Hence a marginal increase in a_t lowers date- s posteriors about η by the amount σ^{-2}/h_s . The impact in utils follows multiplying these marginal effects by the expected value of the sensitivity parameter γ .

Analytically, (11) is more convenient when re-written as follows:

$$\left[\frac{\sigma^{-2}}{h_t} p_t + \gamma_t + U_a(w_t, a^*) \right] (a - a^*) \leq 0, \text{ for all } a \in [0, 1], \quad (12)$$

where

$$p_t \triangleq h_t E \left[- \int_t^T e^{-\rho(s-t)} \gamma_s \frac{1}{h_s} ds \mid \mathcal{F}_t^a \right] \quad (13)$$

⁹The necessary condition can also be derived using Williams' (2008, 2009) method based on the stochastic maximum principle.

is a stochastic process capturing the value of private information.

The reformulated necessary condition (12) involves two stochastic variables, γ_t and p_t . This is a usual result for dynamic contracts with private information.¹⁰ First, we recover the now standard technique of using the promised value to encode past history. A related interpretation can be inferred for p noticing that the incentive constraint implied by (12) is

$$\gamma_t \geq -U_a(w_t, a_t) - \frac{\sigma^{-2}}{h_t} p_t . \quad (14)$$

Given that the agent is risk averse, it is reasonable to conjecture that the principal will minimize the volatility parameter γ . Hence, as long as $a_t^* > 0$, the necessary condition (12) will hold with equality almost everywhere along the equilibrium path. We show below that this indeed holds true when the agent has exponential utility. We therefore replace γ_t by the expression implied for it when (12) binds and, as shown in Appendix B.1., obtain the following solution:

$$p_t = E \left[\int_t^T e^{-\rho(s-t)} U_a(w_s, a_s) ds \middle| \mathcal{F}_t^a \right] < 0 . \quad (15)$$

Intuition behind (15).— The second state variable p is equal to the expected discounted marginal cost of future efforts. Multiplying it by the ratio σ^{-2}/h_t yields the marginal effect of cumulative effort on the continuation value. The intuition for this result can be laid out considering *mimicking strategies*. Fix \bar{Y}_t and lower A_t by $\delta > 0$. Then define a strategy enabling the agent to reproduce the payoffs of an agent with the reference level A_t of past effort. Let a_t^* denote the optimal effort at time t of the reference policy with cumulative effort A_t . By providing $a_t^\delta = a_t^* - \delta\sigma^{-2}/h_t$,¹¹ the agent with cumulative effort $A_t - \delta$ ensures that cumulative output will have the same drift as along the reference path

$$\hat{\eta}(Y_t - (A_t - \delta), t) + a_t^\delta = \frac{h_0 m_0 + \sigma^{-2}(A_t - \delta)}{h_t} + a_t^* - \frac{\sigma^{-2}}{h_t} \delta = \hat{\eta}(Y_t - A_t, t) + a_t^* .$$

Assume now that a similar strategy is employed afterwards, so that $a_s^\delta = a_s^* - (\sigma^{-2}/h_t)\delta$ for all $s \geq t$. Cumulative effort will be $A_s^\delta = A_s^* - [1 + (\sigma^{-2}/h_t)(s-t)]\delta$

¹⁰For example, Werning (2001) shows that in principal-agent problems with hidden savings, one has to introduce both continuation value and expected marginal utility from consumption.

¹¹Such strategies are not feasible when the reference control is at the lower bound, i.e., when $a_t^* = 0$. One should therefore interpret our discussion of mimicking strategies as an heuristic one. The rigorous interpretation being that of the expectation term $E \left[- \int_t^T \gamma_s \frac{\sigma^{-2}}{h_s} ds \middle| \mathcal{F}_t^a \right]$ laid-out in the paragraph above.

leading to the following output drift

$$\begin{aligned} \hat{\eta}(Y_s - A_s^\delta, s) + a_s^\delta &= \frac{h_0 m_0 + \sigma^{-2} (A_s^* - [1 + (\sigma^{-2}/h_t)(s-t)] \delta)}{h_s} + a_s^* - \frac{\sigma^{-2}}{h_t} \delta \\ &= \hat{\eta}(Y_s - A_s^*, s) + a_s^* - \frac{\sigma^{-2}}{h_t h_s} \left[\underbrace{(h_t + \sigma^{-2}(s-t))}_{=h_s} - h_s \right] = \hat{\eta}(Y_s - A_s^*, s) + a_s^* . \end{aligned}$$

As desired, the mimicking strategy reproduces the distribution of Y_s for all $s \geq t$ and the product $-(\sigma^{-2}/h_t)p_t$ measures its expected discounted return in utils.¹² It is positive because it took the agent with cumulative effort A_t more work to produce Y_t , implying that his productivity is likely to be lower. Returns decrease over time as the influence of output on beliefs is lower when η is known more precisely. This suggests that incentives become easier to provide, a result that we will discuss at length in Section 4.

3.2 Sufficient conditions for the agent's problem

First-order conditions rely on the premise that the agent's objective is globally concave. Unfortunately, principal-agent problems do not always fulfill such a requirement. In our case, establishing concavity is complicated by the persistence of private information: As explained in the introduction, deviations from recommended effort drive a permanent wedge between the beliefs of the agent and that of the principal. This is why excluding one shot deviations does not necessary rule out multiple deviations. In order to clarify this distinction we introduce the notion of *implementability* and refer to a control a as implementable if, when assigned the wage function satisfying the local incentive constraint (12) and the promise keeping constraints for v and p , i.e., (9) and (18), the agent finds it optimal to provide effort a .

How to establish implementability for discrete time contracts with persistent information remains an open question.¹³ By contrast, when the model is cast in continuous time, the sufficiency of the necessary conditions and thus the implementability of the control follow from the concavity of the agent's Hamiltonian. This general mathematical result is summarized in Theorem 3.5.2 of Yong and Zhou (1999), and has already been used in principal-agent settings by Schättler and Sung (1993) and more

¹²The correction term σ^{-2}/h_t required to mimic the output distribution remains constant over time because of two countervailing mechanisms. On the one hand, as h_s increases, the impact of past deviations on posteriors decreases over time. On the other hand, the mimicking strategy involves repeated deviations so that the gap between A_s^* and A_s^δ widens over time. When the output distribution is normal, these two opposite forces offset each other.

¹³The difficulties arising in discrete time settings are thoroughly discussed by Abraham and Pavoni (2008). To circumvent them, they propose a numerical procedure verifying *ex-post* the implementability of contracts with hidden effort and savings. See also Kocherlakota (2008) for a discussion of the problem and an analytical example.

recently by Williams (2008). In our case, the agent's Hamiltonian turns out to be concave when the requirements stated in the following proposition are fulfilled.¹⁴

Proposition 2 *A control a is implementable if (11) and*

$$-2U_{aa}(w_t, a_t) \geq e^{\rho t} \xi_t \sigma^2 h_t \quad (16)$$

are true for almost all t , where ξ is the predictable process defined uniquely by

$$E \left[- \int_0^T e^{-\rho s} \gamma_s \frac{\sigma^{-2}}{h_s} ds \middle| \mathcal{F}_t^a \right] - E \left[- \int_0^T e^{-\rho s} \gamma_s \frac{\sigma^{-2}}{h_s} ds \middle| \mathcal{F}_0^a \right] = \int_0^t \xi_s \sigma dZ_s, \text{ for all } t \in [0, T]. \quad (17)$$

According to (15), the process ξ_t is the random fluctuation in the discounted sum of marginal utilities as evaluated from time 0. These restrictions are stronger than required so that a control might violate them and nevertheless be implementable. Moreover, (16) and (17) are stated in terms of γ_t which is endogenous, implying that (16) has to be verified ex-post for any given contract. In some cases, however, one can translate (16) and (17) into a requirement on the parameters of the model. Indeed, when the agent's utility function is as in (20), we shall show that (16) and (17) will hold if (27) holds.

Finally, observe that letting the horizon T go to infinity allows us to discard the terminal condition (10) as long as the transversality condition $\lim_{T \rightarrow \infty} e^{-\rho t} W(\bar{Y}_T)$ is satisfied. Then we can replace the Backward Stochastic Differential Equation¹⁵ (9) by a Stochastic Differential Equation (SDE hereafter) and express the law of motion of the stochastic process p as follows.

Corollary 1 *In the infinite horizon case, p_t (defined in (15)) satisfies*

$$dp_t = \left[p_t \left(\rho + \frac{\sigma^{-2}}{h_t} \right) + \gamma_t \right] dt + \vartheta_t \sigma dZ_t, \quad (18)$$

with

$$\vartheta_t \triangleq e^{\rho t} \sigma^2 h_t \xi_t$$

where ξ_t is defined in (17).

¹⁴The concavity requirement derived in Williams (2008) tends to be violated by his principal-agent problem. Corollary 2 below shows that this is not necessarily the case in our model because implementability is not anymore an issue when parameter precision h_t goes to infinity.

¹⁵A Backward Stochastic Differential Equation is a Stochastic Differential Equation on which a terminal condition has been imposed. In our case, we assumed that the agent's value v_t equals $W(\bar{Y}_t)$ at the end of the contracting horizon, i.e., when $t = T$.

4 Optimal contract under exponential utility

We now show how one can solve for the principal's problem and derive the optimal contract in closed form when attention is restricted to commitment over an infinite horizon and exponential utility functions. The main idea is to simplify the optimization program by eliminating two states: The first one is a component of the sufficient statistics for beliefs, $\hat{\eta}$; and the second one is the value of private information, p . We now describe how each of these is dealt with.

Eliminating $\hat{\eta}$ from the list of states.— According to (5) the principal's problem has an infinite horizon, so that his objective reads¹⁶

$$\begin{aligned} J_t &\triangleq E \left[\int_t^\infty e^{-\rho s} (\hat{\eta}(Y_s - A_s^*, s) + a_s - w_s) ds \middle| \mathcal{F}_t^Y \right] \\ &= \left(\frac{e^{-\rho t}}{\rho} \right) \hat{\eta}(Y_t - A_t^*, t) + E \left[\int_t^\infty e^{-\rho s} (a_s - w_s) ds \middle| \mathcal{F}_t^Y \right]. \end{aligned}$$

The equality follows because the agent is risk neutral and beliefs are a martingale. This implies that one of the two sufficient statistics of beliefs, the mean, can be dispensed with as a state, leaving only precision as the remaining belief state, and since h_t is deterministic, we may index precision by t . This illustrates that incentives are optimally designed to reward effort and not ability.

The principal's optimization problem can therefore be recast as¹⁷

$$j_t \triangleq \max_{\{a, w, \gamma, \vartheta\}} E \left[\int_t^\infty e^{-\rho s} (a_s - w_s) ds \middle| \mathcal{F}_t^Y \right],$$

subject to the two promise-keeping constraints (9) and (18) and subject to the incentive constraint (14). One can assume that the incentive constraint (14) holds with equality almost everywhere because, as shown below, the principal's value function is concave in the promised value v so that he would like to lower the volatility in v as much as possible. Hence we can treat the volatility term $\gamma_t = -U_a(w, a) - \frac{\sigma^{-2}}{h_t} p_t$ as a function of the other controls. Furthermore, (15) implies that the deterministic trend for p is equal to $\rho p - U_a(w, a)$ when (14) binds.

The resulting optimization problem is a standard one since, by (9) and (18), the state variables (v, p) are Markovian. We are therefore justified in using a Hamilton-Jacobi-Bellman (HJB) equation in order to characterize the principal's value func-

¹⁶Profits are discounted from date 0 for analytical convenience.

¹⁷We use a strong formulation for the principal's problem even though we have used a weak formulation to solve for the agent's problem. This change of solution method is usual for principal-agent models. Yet, as discussed in Cvitanic *et al.* (2009), it may lead to measurability issues if the optimal action directly depends on the Brownian motion. In our case, however, a^* turns out to be constant over time so that measurability of the optimal control will not be problematic.

tion.¹⁸ If we had to keep all three states (t, v, p) , the HJB equation would read

$$0 = \max_{\{a, w, \vartheta\}} \left\{ e^{-\rho t} (a - w) + \frac{\partial j}{\partial t} + \frac{\partial j}{\partial v} (\rho v - U(w, a)) + \frac{\partial j}{\partial p} (\rho p - U_a(w, a)) \right. \\ \left. + \frac{\sigma^2}{2} \left[\frac{\partial^2 j}{\partial v^2} \gamma(t, p, w, a)^2 + \frac{\partial^2 j}{\partial p^2} \vartheta^2 + 2 \frac{\partial^2 j}{\partial v \partial p} \gamma(t, p, w, a) \vartheta \right] \right\}. \quad (19)$$

We can, however, reduce the list of states by eliminating p , and this will simplify (19) considerably.

Eliminating p from the list of states.— In order to dispense with p as a state, we assume the following utility function:¹⁹

$$U(w, a) = -\exp(-\theta(w - \lambda a)), \text{ with } \lambda \in (0, 1), \quad (20)$$

for $a \in [0, 1]$. Imposing $\lambda < 1$ ensures that the first-best action is $a = 1$ because the marginal utility of an additional unit of output exceeds the marginal cost of effort regardless of η .²⁰ The utility is defined even for negative consumption which in equilibrium occurs with positive probability.

When $U(a, w)$ is given by (20), the problem greatly simplifies because $U_a(w, a) = \theta \lambda U(w, a)$. Then (8) and (15) imply that

$$p_t = \theta \lambda v_t.$$

The proportionality of v and p means that keeping track of one of the two states is sufficient.²¹ This further reduces the dimensionality of the problem and allows us to rewrite the HJB equation (19) as

$$0 = \max_{\{a, w\}} \left\{ e^{-\rho t} (a - w) + \frac{\partial j}{\partial t} + \frac{\partial j}{\partial v} (\rho v - U(w, a)) + \left(\frac{\sigma^2}{2} \right) \frac{\partial^2 j}{\partial v^2} \gamma(t, v, w, a)^2 \right\}. \quad (21)$$

Given that effort levels lie in a compact set, the recommended action satisfies

$$e^{-\rho t} - \frac{\partial j}{\partial v} U_a(w, a) + \sigma^2 \frac{\partial^2 j}{\partial v^2} \gamma(t, v, w, a) \frac{\partial \gamma(t, v, w, a)}{\partial a} \geq 0,$$

¹⁸Appendix B.2 shows that the HJB equations defined below can be extended to include $\hat{\eta}$ and would still be satisfied.

¹⁹Even though the full characterization of the contract will hold only for utilities of the form (20), the optimality conditions derived in Section 3 hold independently of this parametric restriction. One of its implications is that there is no wealth effect on leisure because $U_w/U_a = -\lambda^{-1}$, a constant independent of w .

²⁰Accordingly, one could interpret our model as resulting from a situation where the agent is able to divert cash flows $1 - a$ at the rate λ . As in DeMarzo and Sannikov (2009), setting λ below one ensures that cash diversion entails linear losses. Our problems differ because DeMarzo and Sannikov (2009) focus on risk neutral agents whereas we introduce risk aversion by taking a concave transformation of the agent's income net of his opportunity cost λa .

²¹To the best of our knowledge, this simplification of the principal's problem with private information and exponential utility was first noticed by Williams (2008).

whereas wages take value over the real line and so fulfill the optimality condition

$$-e^{-\rho t} - \frac{\partial j}{\partial v} U_w(w, a) + \sigma^2 \frac{\partial^2 j}{\partial v^2} \gamma(t, v, w, a) \frac{\partial \gamma(t, v, w, a)}{\partial w} = 0 .$$

Using once again the fact that the Incentive Constraint (14) holds with equality, we obtain $\partial \gamma / \partial w = -\lambda \partial \gamma / \partial a > -\partial \gamma / \partial a$, which implies in turn that, when the optimality condition for wages binds, the one for effort is not tight. It follows that optimal effort is constant and set equal to the upper-bound $a = 1$. Fixing the agent's action to its first best level allows us to solve for the value function by guess-and-verify.

Proposition 3 *Assume that U is as specified in (20). Then the recommended effort is set equal to the first best level $a^* = 1$ and the principal's value function is of the form*

$$j(t, v) = \frac{e^{-\rho t}}{\rho} \left[j_0(t) + \frac{\ln(-v)}{\theta} \right] , \quad (22)$$

The function $j_0(t)$ is the unique solution of the first order ODE

$$j_0'(t) - \rho j_0(t) = -\rho \left(1 - \lambda + \frac{\ln(-k_t)}{\theta} \right) + \frac{\theta (\sigma \lambda)^2}{2} \left(\frac{1}{\sigma^4 h_t^2} - k_t^2 \right) , \quad (23)$$

with boundary condition $\lim_{t \rightarrow \infty} j_0'(t) = 0$ and k_t being given by the negative root of the quadratic equation

$$k_t^2 (\sigma \lambda \theta)^2 - k_t \left(1 + \frac{1}{h_t} (\lambda \theta)^2 \right) - \rho = 0 . \quad (24)$$

The optimal wage is

$$w_t^*(v) = -\frac{\ln(k_t v)}{\theta} + \lambda , \quad (25)$$

and the optimal volatility reads

$$\gamma_t^*(v) = \Gamma_t v \triangleq \lambda \theta \left(k_t - \frac{\sigma^{-2}}{h_t} \right) v . \quad (26)$$

To establish the implementability of the first best action, remember that our parametrization of $U(w, a)$ is such that $p_t = \theta \lambda v_t$. Consequently, the volatility terms γ_t^* and ϑ_t^* must also remain proportional. Reinserting $\vartheta_t^* = \theta \lambda \gamma_t^*$ into (16) and using the explicit solution (26) for γ_t^* yields the following requirement.

Corollary 2 *First best effort is implementable (i.e., meets conditions (11) and (16)) when*

$$\rho \sigma^2 > \frac{1}{h_0} + 2 (\lambda \theta)^2 \frac{1}{h_0^2} . \quad (27)$$

Since precision h_t is increasing with time, the condition then holds at all subsequent dates as $h_t > h_0$.

The sufficient condition (27) is more likely to hold when: Both parties are impatient, output noise is high, the marginal cost of effort λ is low, the coefficient of absolute risk aversion θ is small, or parameter precision h_0 is high. Indeed, (27) always holds in the limit case without parameter uncertainty ($h_0 = \infty$) because multiple deviations are not anymore a concern.

We shall henceforth assume that our parameters satisfy (27). The condition is sufficient and not necessary, however, and our comparative statics results hold independently of it, which suggests that they are robust over a wider region of the parameter space.

5 Characterization of the optimal contract

The optimal wage process described in (25) has a declining volatility, as well as a drift converging to a negative limit. The first property appears to be quite general, and should hold for any utility function. The second property is specific to the parametrization in (20). The following arguments will suggest that if we could solve the problem for a utility function for which the inverse marginal utility of income ($1/U'(w)$) is concave in w , the drift would converge to a positive limit.

5.1 Wage dynamics

The mechanism driving wage volatility is the decrease in the ability of the agent to manipulate beliefs as they become more precise over time. It enables the principal to sustain first best effort with less variance and to trade lowers wages in exchange of more stable income. This channel is easily derived from the analytical expression (25) for wages.

Corollary 3 *For any given promised value v , the optimal wage $w_t^*(v)$ is a decreasing function of beliefs precision and thus time.*

Corollary 3 does not directly apply to income dynamics because the promised value v evolves over time. To obtain the law of motion of v , we reinsert the optimal volatility $\gamma_t^*(v)$ defined in (26) into the SDE (9)

$$dv_t = v_t [(\rho + k_t) dt + \Gamma_t \sigma dZ_t] . \quad (28)$$

Since k_t is the negative root of (24), the drift of the promised value can be positive or negative. Its sign indicates how earnings are allocated over time: When the drift is positive, wages are back loaded, meaning that the expected average wage is above its current level. Conversely, when the trend is negative, payments are front loaded. Given that k_t is decreasing over time,²² the principal resorts more intensively to

²²See the proof of Corollary 3.

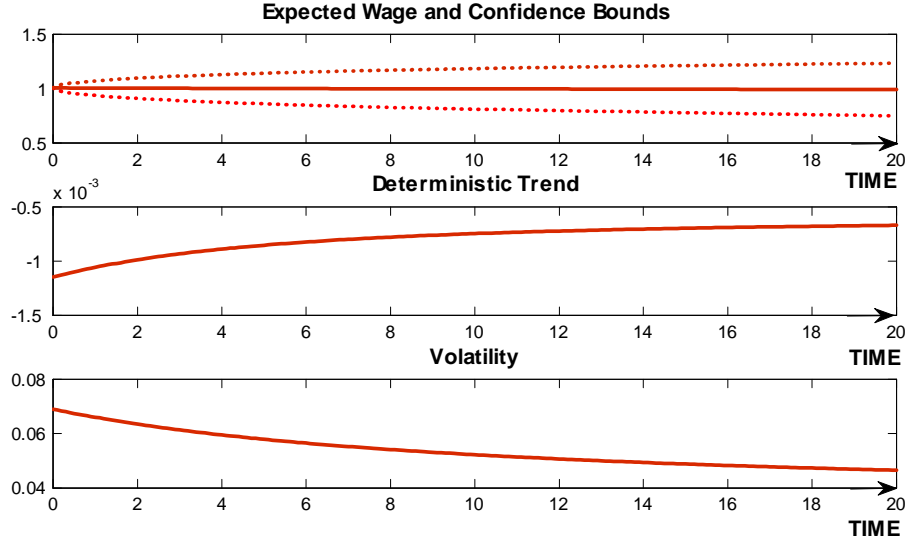


Figure 2: WAGE DYNAMICS AS A FUNCTION OF CONTRACT DURATION

back loading when parameter uncertainty is higher. Payments are deferred because incentives can be provided at a cheaper cost in the future through higher income stabilization.

Accordingly, income dynamics result from the interaction of the following three mechanisms: (i) For a constant promised value, wages decrease over time, as stated in Corollary 3; (ii) Back loading weakens over time, raising current income; (iii) Wages are driven downwards by the agent's immiserization. Of the three channels, only the first two are specific to the learning process whereas the third one remains relevant when belief precision is infinite. Deriving the law of motion of wages allows one to analytically identify each mechanism. The optimal wage at time t as a function of the promised value v is given by

$$w_t^* = - \left(\frac{1}{\theta} \right) [\ln(-k_t) + \ln(-v_t) + \lambda] ,$$

so that its law of motion reads

$$dw_t^* = - \left(\frac{1}{\theta} \right) \left[\left(\frac{1}{k_t} \right) dk_t + d \ln(-v_t) \right] . \quad (29)$$

Reinserting from (28) into (29) and applying Ito’s lemma to the logarithmic transformation of v yields the “reduced form” for wage growth

$$dw_t^* = \frac{1}{\theta} \left(\underbrace{-\frac{dk_t/dt}{k_t}}_{\text{Income Stabilization}} + \underbrace{\frac{(\theta\lambda)^2}{2} \left(\frac{\sigma^{-1}}{h_t}\right)^2}_{\text{Back Loading}} - \underbrace{\frac{(\sigma\theta\lambda)^2}{2} k_t^2}_{\text{Immiserization}} \right) dt + \frac{\Gamma_t}{\theta} \sigma dZ_t . \quad (30)$$

The first two terms in the expression for the trend are due to parameter uncertainty and they vanish when belief precision h_t is infinite. The trend and volatility terms in (30) are both deterministic, and are plotted in the second and third panels of Figure 2. The assumed parameter values are shown in Table 1. They will be used as baseline numbers for all the simulations reported below.

ρ	σ^2	θ	λ	h_0
0.1	0.5	1	0.5	20.48
TABLE 1				

The value $h_0 = 20.48$ is the smallest precision that satisfies the second-order condition (27) given the assumed values of the other parameters. The middle panel of Figure 2 shows that the trend is increasing over time. Hence, parameter uncertainty reinforces the immiserization process because the back loading channel is dominated by the income stabilization channel. This is not a general result, however, as other parameter constellations yield decreasing or even hump-shaped profiles for the deterministic trend.

The third term in the trend in (30) captures the agent’s immiserization which is specific to the utility function (20). It follows from the inverse Euler equation that can be established in the infinite-precision limit using Ito’s lemma

$$d\left(\frac{1}{\partial U/\partial w_t}\right) = -\frac{\lambda\sigma}{v} dZ_t , \text{ when } \frac{\sigma^{-2}}{h_t} = 0 .$$

Under (20), $(\partial U/\partial w)^{-1} = \exp(\theta[w - \lambda])/\theta$ is convex in w , hence the immiserization. However, if utility were $U(c) = c^{1-\phi}/(1-\phi)$ and $\phi < 1$, $(\partial U/\partial w)^{-1} = c^\phi$ would be concave and the inverse Euler equation would imply that wages exhibit a positive trend. In the knife-edge case $\phi = 1$, the utility would be logarithmic and wages would follow a martingale.

The top panel of Figure 2 plots the mean wage and the one-standard-deviation bands for the parameter values in Table 1. The stochastic term σdZ is the output surprise defined in (6), which means that the solution w_t^* to the stochastic difference equation is a normally distributed random variable, and that the distribution of wages at date t is the frequency distribution of wages among age- t workers with abilities randomly drawn from $\eta \sim N(0, h_0^{-1})$. By normality, the bands are equidistant from the mean, hence, symmetric.

Now, from (48) we find that k_t has a strictly negative limit so that

$$|k_t| \rightarrow \frac{1}{2} \left(\sqrt{\left(\frac{1}{(\sigma\lambda\theta)^2}\right)^2 + \frac{4\rho}{(\sigma\lambda\theta)^2} - \frac{1}{(\sigma\lambda\theta)^2}} \right) > 0 ;$$

implying that the volatility of the wage increments does not die off

$$\left| \frac{\Gamma_t}{\theta} \sigma \right| = \left| \lambda \sigma \left(k_t - \frac{\sigma^{-2}}{h_t} \right) \right| \rightarrow \lambda \sigma |k_\infty| > 0 .$$

Since these increments are independent, the cross-section variance of wages converges to infinity. We sum up our findings in the Corollary below, whereas Figure 2 illustrates them.

Corollary 4 *The volatility of the wage increments is decreasing to a positive limit so that the cross-sectional variance of wages grows without bound. Provided that the sufficient condition (16) is satisfied, wages exhibit a negative trend.*

5.2 Value of Commitment

Instead of focusing on wage dynamics within a given match, we can use the model to compare the value of commitment across different environments. As discussed in the Introduction and in Section 6 below, the total surplus is decreasing in prior precision when wages are set through spot contracts. To the contrary, when parties are able to commit, the surplus is higher when priors are more accurate.

Corollary 5 *The principal's expected lifetime profit as a function of the value v promised to the agent is increasing in the prior precision h_0 .*

The intuition for this result directly follows from Corollary 3: An increase in the precision with which the productivity of the match is known enables the principal to stabilize further the agent's income. As contracts get closer to the second best, the principal can deliver the promised value v at a lower expected cost.

Figure 3 plots the agent's value as a function of the prior variance $1/h_0$ and of the marginal cost of effort parameter λ , holding the principal's value constant at zero. The vertical line labeled "sufficient condition" identifies the maximal prior variance $1/h_0$ and λ above which implementability holds surely. The other parameters are as given in Table 1. In particular, (27) (which involves both λ and h) holds to the left of the solid black line. For the parameter values used in the plot, (27) reads $\frac{1}{20} > \frac{1}{h_0} + 2 \left(\frac{\lambda}{h_0} \right)^2$, and so the maximal λ as a function of h is given by

$$\lambda = \sqrt{\frac{h_0}{2} \left(\frac{h_0}{20} - 1 \right)} . \quad (31)$$

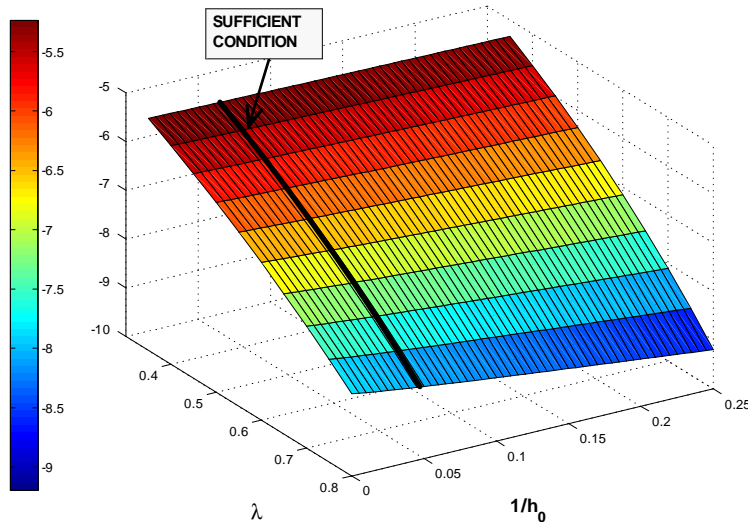


Figure 3: AGENT'S VALUE AS A FUNCTION OF $1/h_0$ AND λ

The RHS of this equation is positive only if $h_0 \geq 20$. In other words, (27) can be met only if $1/h_0 < 5\%$, and then more easily if λ is low enough. Once $h_0 \geq 21.83$, however, the RHS of (31) exceeds unity, and (27) then holds for all $\lambda \in [0, 1]$.

As stated in Corollary 5, the agent's value is decreasing in the prior variance $1/h_0$. Figure 3 also illustrates how an increase in λ lowers the surplus because it intensifies the moral hazard problem, thus making it more costly for the principal to deliver a given utility level.

Williams (2009) proves qualitatively similar results in a reporting problem with persistent income shocks: Efficiency losses due to private information increase with the persistence of the endowment and, parallel to our result that the principal back loads payments more when h_t is lower, Williams also finds that persistence of shocks leads to a tendency to back load payments that is absent in reporting problems with i.i.d. shocks.

6 Limited *vs.* full commitment

Our full-commitment solution applies equally to various interpretations for η ; it can denote the agent's general ability fully transferable across matches, it can denote a match-specific productivity, or any combination of the two. Which interpretation one adopts can affect the solution only via the positioning of the initial point on the Pareto frontier. Under limited commitment, however, the form of the participation constraints will depend on the interpretation of η .

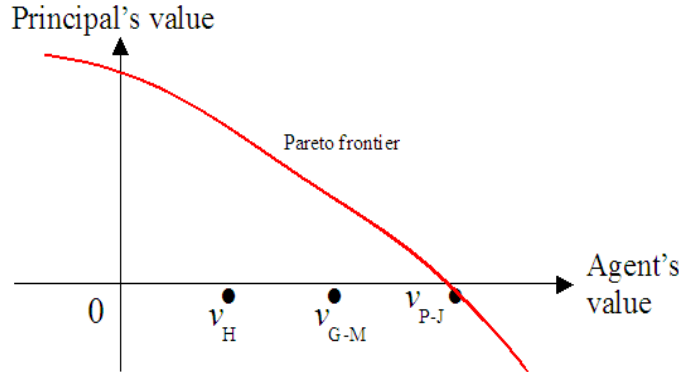


Figure 4: LIFETIME RENTS OF THE AGENTS IN THE THREE MODELS.

We now wish to relate our model to the literature on reputations that typically adopts the interpretation that η is general ability. We shall focus on two models of reputations. The first, considered by Holmström (1999; “H” hereafter), assumes spot-market wages that may reflect the worker’s history but cannot reflect current output. The second, considered by Gibbons and Murphy (1992; “G-M” hereafter) allows wages to respond linearly to performances during the period at hand. Since G-M have some form of partial commitment, the G-M agents receive a higher utility after every history than the H agents.

In both H and G-M the principal is assumed to be risk neutral, and they impose zero expected profits for the principal after every history and at each date. The agents’ utility functions, however, differ from our assumed form in (4) and (20): H assumes that agents are risk-neutral and have time-additive utility, whereas G-M assumes that agents are risk averse but that their utility is not time separable. To make our analysis of commitment comparable to their analyses of limited commitment, we shall derive the equilibria of H and G-M in our environment, i.e., for the case where the agent has lifetime utility (4) and period utility (20). This is the only change we make to H and G-M.

In our model (“P-J” hereafter), the principal has full commitment and his profits will not be zero at an arbitrary date. To compare our solution to H and G-M, it is natural to impose zero expected lifetime profits on the principal at the outset. Thus we shall assume that at date zero, the agent gets all the rents from the relationship. If we maintain the same belief about η across the three models, and if we use v_H , v_{G-M} , and v_{P-J} to denote the agent’s lifetime utility, then they are related as shown in Figure 4.

The relation depicted in Figure 4 exists only at the outset, when under commitment risk-neutral firms compete for the agent. Of course, here we are discussing three separate economies each with its own distinct contracting arrangement, and

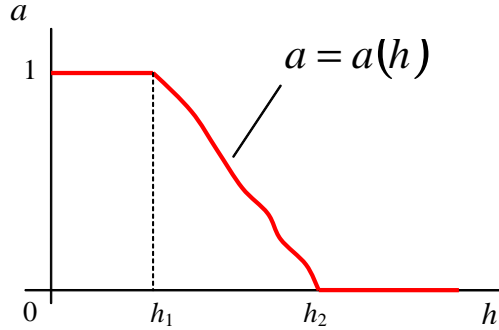


Figure 5: EFFORT AS A FUNCTION OF PRECISION IN SPOT MARKETS.

not a single economy in which lifetime contracts and spot contracts could coexist. We now wish to transport this intuition to the behavior of wages.

6.1 Ex-ante payments

In this section we show that the equilibrium behavior of wages and effort under risk aversion is essentially the same as in H: Reputational concerns are the only reason why the agent exerts any effort, and when information about η accumulates and as these concerns disappear, his effort converges to zero, just as in the risk-neutral case. Of itself this is not surprising. Rather, the result is useful because it enables us to isolate the role that full commitment plays in generating economic outcomes for the parties to the contract.

Employers cannot commit to paying wages that depend on performance, and competition among employers bids wages up to expected output. Denoting as before equilibrium actions by an asterisk, expected productivity reads:

$$w_t = \hat{\eta}(Y_t - A_t^*, t) + a_t^* . \quad (32)$$

In H, equilibrium action entails a strictly declining deterministic sequence for a_t^* . Effort is sustained by the market's imprecise knowledge of η and the agent's attempts to raise the market's expectation. With our utility function and a spot market, the sequence a_t also decreases, eventually reaching zero and remaining there, as drawn in Figure 5 and described in Proposition 4.

Proposition 4 (i) *The equilibrium effort path a_t is deterministic, and it depends on t only through $h_t = h_0 + \sigma^{-2}t$, as drawn in Figure 5.*

(ii) *There exist two numbers h_1 and h_2 satisfying $0 \leq h_1 \leq h_2$ such that (A) $a(h) = 1$ for $h \leq h_1$; (B) $a(h)$ is strictly decreasing for $h \in (h_1, h_2)$; and (C) $a(h) = 0$ for*

$h \geq h_2$.
 (iii)(A) If

$$\lambda < \int_0^\infty e^{-\rho\tau} \left[\left(\frac{\sigma^{-2}}{h_0 + \tau\sigma^{-2}} \right) \exp \left(\frac{\theta^2}{2} \left(\frac{\sigma^{-2}}{h_0 + \tau\sigma^{-2}} \right)^2 (\tau\sigma^2 + h_0^{-1}) \right) \right] d\tau . \quad (33)$$

then $h_2 > 0$. (B) $h_1 < h_2$. Moreover, if

$$\frac{\partial U(m_0, 1)}{\partial a} + \int_0^\infty e^{-\rho t} \frac{\partial}{\partial Y} E_0 [U(\hat{\eta}(Y_t - A_t, s), 1)] ds > 0 , \quad (34)$$

then (C) $h_1 > 0$, i.e., an initial horizontal segment at $a = 1$ exists.

The following properties are of note:

1. Since a depends on t only through the effect that t has on h , lowering the initial precision of the prior (i.e. decreasing h_0) raises the time T at which the agent stops providing effort. In (a, t) space, the entire effort path shifts to the right.
2. Since a_t is deterministic, wage volatility is declining with experience because the volatility of $\hat{\eta}$ is declining with t . Of course, conditional on $\hat{\eta}$ and h , the wage is not random.
3. Since first-best effort is equal to the upper bound of unity, effort cannot ever exceed its first-best level. In terms of welfare this is the only difference from H.

Recall that the equilibrium wage is $\hat{\eta} + a_t^*$. If we normalize the mean of η to zero (as we shall do throughout this section), the average equilibrium wage is

$$\bar{w}_t^H = a_t^* , \quad (35)$$

with the sequence of a_t^* depicted in Figure 5. The efficient level $a = 1$ is implementable only early on, and wages reflect that fact. We shall compare P-J with H in a simulation of both using the parameter values in Table 1. Figure 6 reports the simulation result. We choose $h_0 = 20.48$ so that the contract under commitment is implementable. But this level of precision is too high to generate a reputational concern in the H model that would be sufficient to sustain first-best effort $a = 1$. Indeed, by period 3, effort has already reached zero. Therefore \bar{w}_t^H starts out below unity and itself reaches zero by period 3. Wages in both models are normally distributed at each date and Figure 6 shows that most agents would receive higher wages under commitment than they would in the spot market, but that commitment entails more wage dispersion.

The distribution of lifetime utilities.—While Figure 6 shows the distribution of wages across agents at each date, it does not accurately represent the distribution of lifetime gains that full commitment offers. That would be the distribution of the

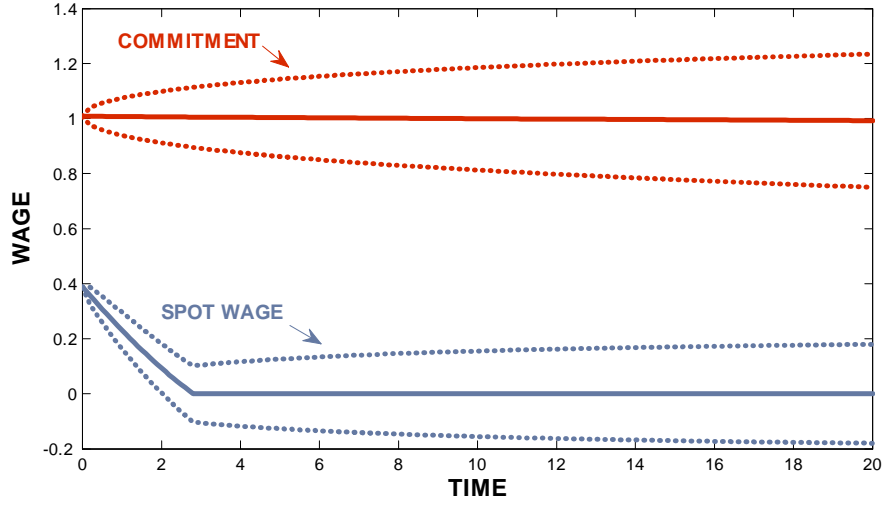


Figure 6: MEAN WAGE AND STANDARD DEVIATION BANDS OF COMMITMENT WAGES AND SPOT WAGES

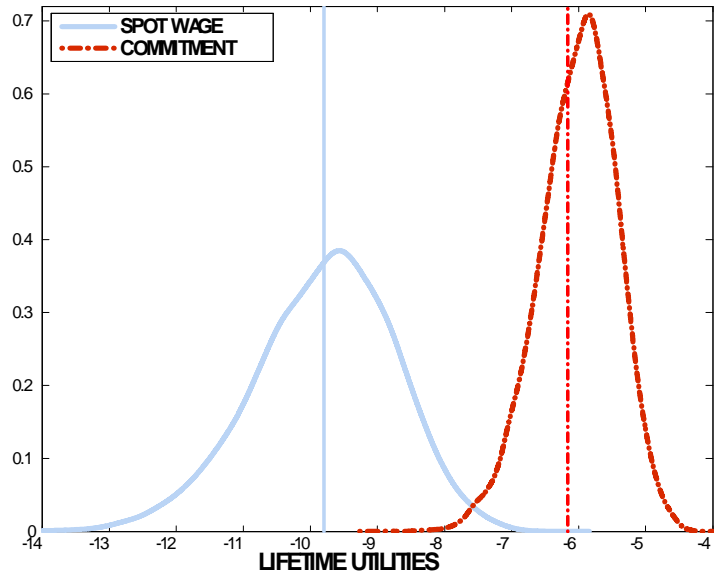


Figure 7: DISTRIBUTION OF LIFETIME UTILITIES UNDER SPOT WAGES AND COMMITMENT

random variable \mathcal{U}_0 defined in (4), which we report in Figure 7.²³ Commitment largely dominates spot agreements: On average, it raises the agent’s utility by 37%,²⁴ which is equivalent to an increase of 89.2% in wages across first best allocations.²⁵ Even though utilities derived from contracts exhibit more dispersion, they dominate from a stochastic point of view. While wages themselves are normally distributed, utilities are nonlinear and bounded above. This is why the resulting distributions of \mathcal{U}_0 are skewed to the left, and more so for commitment. As a result, the means (represented by the vertical lines) are to the left of the modes.

6.2 Ex-post linear payments

Between the extremes of the no-commitment model H on the one hand and the full-commitment model P-J on the other, there is the partial-commitment model G-M in which a contract lasts one period: Wage are paid at the end of each period and can depend linearly on output that period. The market is otherwise still a spot market as there is no contracting for more than one period. Expected profits must still equal zero, but the set of contracts is richer than in H, as it also includes piece rate compensations. The G-M solution therefore provides the agent with a higher expected utility than the H solution, but a lower lifetime utility than our full-commitment solution.

We use discrete time to explain the G-M results for our utility function. A remark about eq. (1). First, if a_t is continuous, (1) can be thought of as the limit of the following discrete time process when the interval length Δ converges to zero

$$Y_t^\Delta \triangleq \sum_{i=1}^{t/\Delta} \left((\eta + a_i) \Delta + \sigma \varepsilon_i \sqrt{\Delta} \right),$$

where $a_i = a_{\Delta i}$, and where ε_i is an i.i.d. shock with unit variance.²⁶

Assume then, that output is $y_t = a_t + \eta + \varepsilon_t$. Since wages are restricted to be linear in output, we can denote the one-period wage function by $w_t = b_{0,t} + b_{1,t}y_t$. We now simulate our solution together with the piece-rate spot-market solution along with

²³The distribution of lifetime utilities in both models is obtained through Monte Carlo simulations. We simulate 10000 sample paths and compute the resulting kernel densities. The accuracy of the procedure is confirmed by the approximation error of less than one percent between the simulated and theoretical average utility in the commitment scenario.

²⁴The welfare gain is obtained dividing the difference between the expected utilities $E_0[\mathcal{U}_0]$ with and without commitment by the expected utility when wages are set on the spot market, i.e., $(-9.792 + 6.162) / -9.792 = 0.37$.

²⁵We first derive the wage such that $U(w, 1) / r = E_0[\mathcal{U}_0]$, which yields $w^{Com} = 0.984$ under commitment and $w^{Spot} = 0.52$ under spot market. The compensating variation follows taking the difference between the two wages and dividing it by w^{Spot} , i.e., $(0.984 - 0.52) / 0.52 = 0.8923$.

²⁶In (1) only mean output depends on a , not its variance. For if, instead, σ also depended on a , say as $\sigma(a)$, the principal could perfectly infer $\sigma(a)$ and hence a , from the observed quadratic variation of Y as $\Delta \rightarrow 0$, for then the signal-noise ratio becomes unbounded.

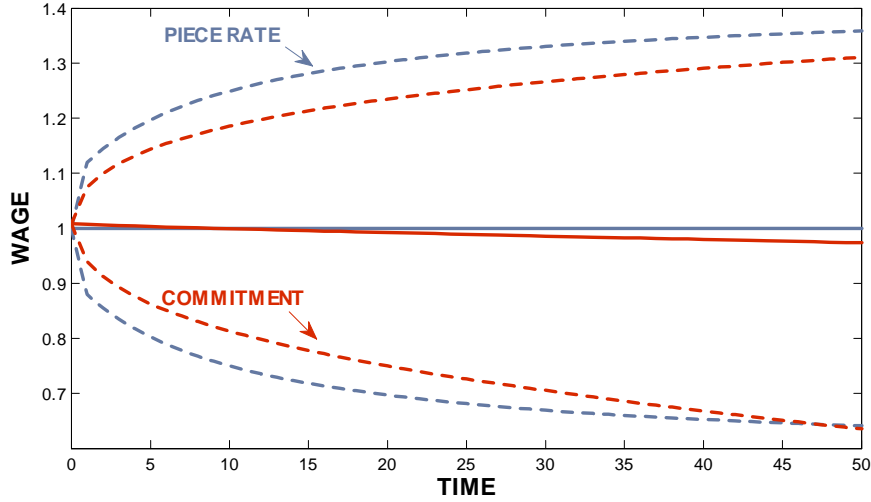


Figure 8: MEAN WAGE AND STANDARD DEVIATION BANDS OF COMMITMENT WAGES AND PIECE-RATE WAGES

the one-standard-deviation bands for the two models. The piece-rate contracts can implement the efficient level of effort. This will be achieved if λ is small enough and as long as the zero-expected-profit constraint holds. We know that $E[w] = E[y] = 1 + \hat{\eta}$, and therefore for the mean agent $\bar{w}_t^{GM} = 1$, for all t , ensuring that piece-rate contracts are linear with zero profit on a period-by-period basis. Observe that in the commitment solution we impose a zero expected lifetime value on the principal, whereas in the spot-market solution the expected profit is zero in each period.

At each t , wages maximize the agent's lifetime utility subject to non-negativity of profits

$$E[w_t] = b_{0,t} + b_{1,t}(\hat{\eta}_t + a_t^*) \leq \hat{\eta}_t + a_t^* , \quad (36)$$

and subject to incentive compatibility (see (11) and its simplification in (56) and (58)). Details are in Appendix C.

Equation (59) reports the standard deviation of the piece-rate wage to be

$$\sqrt{h_0^{-1} - h_{t-1}^{-1} + \sigma^2 b_{1,t}^2} \xrightarrow{t \rightarrow \infty} \sqrt{h_0^{-1} + \lambda^2/2} = \sqrt{(20.48)^{-1} + 1/8} = 0.4169 .$$

It is bounded because level shocks are assumed to be i.i.d.. To the contrary, the standard deviation of wages under commitment diverges to infinity since, as noted in Corollary 4, the variance of its increments does not die off. Yet, for most of the periods reported in Figure 8, the cross-sectional variance of commitment wages is smaller than that of piece-rate wages. It takes around 50 periods for the standard-deviation bands under commitment to become wider than the piece-rate bands. Surprisingly, the

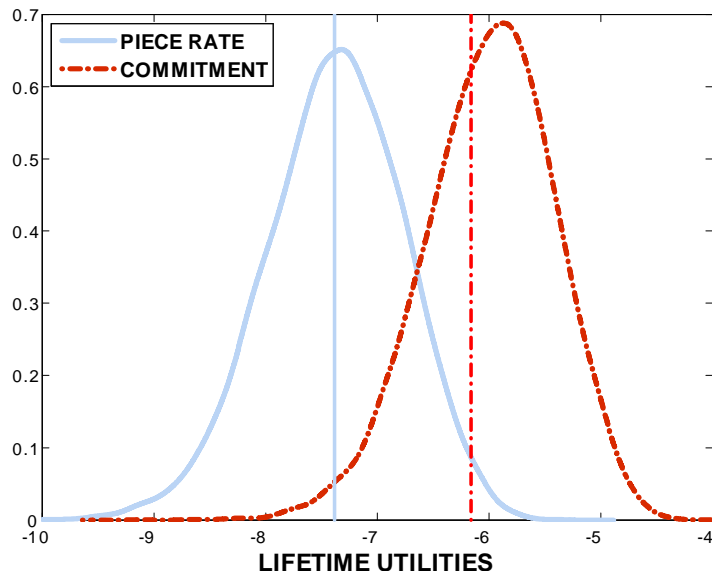


Figure 9: DISTRIBUTION OF LIFETIME UTILITIES UNDER PIECE RATES AND COMMITMENT

standard deviations of piece-rate wages take a long time to stabilize. In Appendix C, Figure 10, we show that this is because $b_{1,t}$ is slow to converge to λ .

We note, however, that the component $\sigma^2 b_{1,t}^2$ represents the contribution of transitory wages, which means that increments to wages are more variable in the piece-rate case, more persistent under commitment. This is why one can easily be misled by Figure 8 into believing that there are not much gains to commitment. Yet Figure 8 does not inform us about the cross-sectional distributions of lifetime utilities which turn out to be quite different across models, as can be seen from Figure 9. Even though the gap is smaller than with spot wages (the distribution on the right is for \mathcal{U}_0 , the same, of course, as the distribution in Figure 7), commitment still offers a noticeably higher expected lifetime utility $E_0[\mathcal{U}_0]$: Long-term contracts raise the agent’s utility by 18.8%, a gain that is equivalent to a compensating variation of 25.9% in wages across first best allocations.²⁷ Furthermore, stochastic dominance continues to hold so that not only the average worker but most workers do benefit from contracting.

6.3 Other remarks on limited commitment

η as a match-specific ability.—If, instead of denoting general ability, η were match specific, then neither the optimal contract nor the Pareto frontier would change under

²⁷For an explanation of how these statistics are derived, see footnotes 24 and 25.

full commitment. By contrast, spot-markets would work poorly. The agent now has no reputational concern, and receives lower lifetime utility. With up-front wages as in H, all reputational concerns disappear; implying that effort would remain constant at zero. The wage would equal $E_t[\eta]$ at *all* dates and so the median agent would receive a wage of zero. On the other hand, linear piece rates, i.e., contracts of the G-M type, would sustain first-best effort, but with a contract that does not change over time:

$$b_{0,t} = 1 - \lambda \quad \text{and} \quad b_{1,t} = \lambda$$

for all t . The mean wage would remain constant $\bar{w}_t^{GM} = 1$ at the level where the principal breaks even. The value of commitment is then even larger than in the case where ability is transferable and again increasing in the precision of the match-quality parameter.

Participation constraints and equilibrium.—In this section we have described three separate economies, each with its own contracting protocol determined by the agents' ability to commit. The P-J solution is for a contract that would yield the principal zero expected profit at the outset, but after some histories his expected profit will fall below zero. Similarly, the agent's continuation value may fall below v_{P-J} and even below v_H . An extension would add participation constraints as Rudanko (2010) and Lustig *et al.* (2007) have done for multi-agent environments without learning. In partial equilibrium settings without learning there are more papers with limited commitment. Closely related to ours is the principal-agent model of Sannikov (2008) which, under some adjustments to the parametric form of the utility function, is encompassed in our framework as $h_0 \rightarrow \infty$, i.e., when posteriors have converged to the true value of η . More precisely, Sannikov considers a utility function that is (i) defined over the positive real line; (ii) bounded from below; and (iii) separable in income and effort. By contrast, our utility function (20) is not bounded from below and, as a result, we do not have a low retirement point. Observe, however, that our characterization of the agent's necessary condition (11) does not depend on the parametric assumption (20) and so coincides with Sannikov's when $h_0 = \infty$.

7 Conclusion

We have solved a contracting problem involving parameter uncertainty and described a mechanism by which uncertainty about the environment worsens the incentive insurance trade-off. We developed an approach that works for any utility function when the parameter and noise are normally distributed. We found that the agent faces two opposite effects when considering a downward deviation from recommended effort. On the one hand, he will be punished by a lower promised value because of the decrease in observable output. On the other hand, he will benefit from higher expectations than the principal about the unknown productivity of the match. This

second channel that we label belief manipulation is specific to problems under parameter uncertainty. The extent to which it influences incentive provisions depends on the remaining length of the relationship. This is why it is not relevant in markets based on spot agreements.

Although the prospect of belief manipulation reduces the gains from commitment, it does not eliminate them altogether. We found, in particular, that the Pareto frontier shifts out when information about quality improves, and this we contrasted to spot markets where, at least when ability is transferable, the Pareto frontier shifts inwards. Therefore incentives are easier to provide and commitment is more valuable when the agent's quality is known more precisely. In further contrast to results under partial commitment, wage volatility under full commitment declines with experience.

Thus, we have shown that parameter uncertainty makes it harder to reward effort under full commitment, in direct contrast to its tendency to stimulate effort in spot markets. However, spot and full commitment settings are both highly stylized depictions of how markets operate in reality. We therefore believe that the most promising task would be to combine the two environments in a model with limited commitment so as to evaluate how the two incentive channels interact.

References

- [1] Abraham, Arpad and Nicola Pavoni. "Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending: A Recursive Formulation." *Review of Economic Dynamics* 11 (2008): 781-803.
- [2] Adrian, Tobias and Mark Westerfield. "Disagreement and Learning in a Dynamic Contracting Model." *Review of Financial Studies* 22(10) (2009): 3839-71.
- [3] Chernoff, Herman. "Optimal Stochastic Control." *Sankhya*, Ser. A (1968): 221-252.
- [4] Cvitanić Jakša, Wan, Xuhu and Jianfeng Zhang. "Optimal Compensation with Hidden Action and Lump-Sum Payment in a Continuous-Time Model". *Applied Mathematics and Optimization* 59 (2009): 99-146.
- [5] DeMarzo, Peter M. and Yuliy Sannikov. "Learning in Dynamic Incentive Contracts." Stanford University, Unpublished manuscript (2008).
- [6] Fernandes, Ana and Christopher Phelan. "A Recursive Formulation for Repeated Agency with History Dependence." *Journal of Economic Theory* 91 (2000): 223-247.
- [7] Fujisaki, Masatoshi, Kallianpur, Gopinath and Hiroshi Kunita. "Stochastic Differential Equation for the Non Linear Filtering Problem." *Osaka Journal of Mathematics* 9 (1972):19-40.

- [8] Giat, Yahel, Hackman, Steve and Ajay Subramanian. “Investment under Uncertainty, Heterogeneous Beliefs, and Agency Conflicts.” *Review of Financial Studies* 23(4) (2010): 1360-1404.
- [9] Gibbons, Robert and Kevin J. Murphy. “Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence.” *Journal of Political Economy* 100(3) (1992): 468-505.
- [10] Holmström, Bengt. “Managerial Incentive Problems: A Dynamic Perspective.” *Review of Economic Studies* 66 (1999): 169-182.
- [11] Holmström, Bengt and Paul Milgrom. “Aggregation and Linearity in the Provision of Intertemporal Incentives.” *Econometrica* 55 (1987): 303-328.
- [12] Hopenhayn, Hugo and Arantxa Jarque. “Moral Hazard and Persistence.” FRB Richmond, Working Paper 07-07 (2007).
- [13] Jarque, Arantxa. “Repeated Moral Hazard with Effort Persistence.” FRB Richmond, Working Paper 08-04 (2008).
- [14] Jovanovic, Boyan. “Job Matching and the Theory of Turnover.” *Journal of Political Economy*, 87(5) pt. 2 (1979): 972-90.
- [15] Kallianpur, Gopinath. *Stochastic Filtering Theory*. Springer-Verlag, New York, (1980).
- [16] Kapicka, Marek. “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach.” University of California, Santa Barbara, Working paper, (2006).
- [17] Kocherlakota, Narayana. “Figuring out the Impact of Hidden Savings on Optimal Unemployment Insurance.” *Review of Economic Dynamics* 7(3) (2004): 541-554.
- [18] Laffont, Jean-Jacques and Jean Tirole. “The Dynamics of Incentive Contracts.” *Econometrica* 56, No. 5 (Sep., 1988): 1153-1175.
- [19] Lustig, Hanno, Chad Syverson and Stijn Van Nieuwerburgh, “IT, Corporate Payouts, and the Growing Inequality in Managerial Compensation.” University of Chicago, (2007).
- [20] Mirrlees, James. “Notes on Welfare Economics, Information and Uncertainty.” In: M. Balch, D. McFadden and S.-Y. Wu, Editors, *Essays on Economic Behavior under Uncertainty*. North-Holland, Amsterdam (1974).
- [21] Rogerson, William. “The First Order Approach to Principal- Agent Problems.” *Econometrica*, 53(6) (1985): 1357-1367.

- [22] Rudanko, Leena. “Labor Market Dynamics under Long-Term Wage Contracting.” *Journal of Monetary Economics*, 56(2) (2010): 170-183.
- [23] Sannikov, Yuliy. “A Continuous-Time Version of the Principal-Agent Problem.” *Review of Economic Studies*, 75(3) (2008): 957-984.
- [24] Schättler, Heinz and Jaeyoung Sung. “The First-Order Approach to the Continuous-Time Principal-Agent Problem with Exponential Utility.” *Journal of Economic Theory* 61 (1993): 331-371.
- [25] Werning, Iván. “Moral Hazard with Unobserved Endowments: A Recursive Approach.” University of Chicago, Unpublished manuscript (2001).
- [26] Williams, Noah. “On Dynamic Principal-Agent Models in Continuous Time.” UW Madison, Unpublished manuscript (2008).
- [27] Williams, Noah. “Persistent Private Information.” UW Madison, Unpublished manuscript (2009).
- [28] Yong, Jongmin and Xun Yu Zhou. *Stochastic Controls*. New York: Springer-Verlag (1999).

Appendix A: Proofs of propositions and corollaries

Proof. Proposition 1: Consider the Brownian motion Z^0 under some probability space with probability measure Q , and $\mathbb{F}^{Z^0} \triangleq \left\{ \mathcal{F}_t^{Z^0} \right\}_{0 \leq t \leq T}$ the suitably augmented filtration generated by Z^0 . Let

$$Y_t = \int_0^t \sigma dZ_s^0 ,$$

so that Y_t is also a Brownian motion under Q . Given that expected output is linear in cumulative output,²⁸ the exponential local martingale

$$\Lambda_{t,\tau}^a \triangleq \exp \left(\int_t^\tau \left(\frac{\hat{\eta}(Y_s - A_s, s) + a_s}{\sigma} \right) dZ_s^0 - \frac{1}{2} \int_t^\tau \left| \frac{\hat{\eta}(Y_s - A_s, s) + a_s}{\sigma} \right|^2 ds \right) , \quad t \leq \tau \leq T ,$$

is a martingale, i.e. $E_t [\Lambda_{t,T}^a] = 1$. Hence Girsanov theorem holds and ensures that

$$Z_t^a \triangleq Z_t^0 - \int_0^t \left(\frac{\hat{\eta}(Y_s - A_s, s) + a_s}{\sigma} \right) ds$$

is a Brownian motion under the new probability measure $dP^a/dP \triangleq \Lambda_{0,T}^a$. Given that both measures are equivalent, the triple (Y, Z^a, Q^a) is a weak solution of the SDE

$$Y_t = \int_0^t (\hat{\eta}(Y_s - A_s, s) + a_s) ds + \int_0^t \sigma dZ_s^a .$$

Adopting a weak formulation allows us to view the choice of control a as determining the probability measure Q^a . In order to define the agent's optimization problem, let $R^a(t)$ denote the reward from time t onwards so that

$$R^a(t) \triangleq e^{\rho t} \left[\int_t^T U(s, \bar{Y}_s, a_s) ds + W(T, \bar{Y}_T) \right] ,$$

where, with a slight abuse of notation, $U(s, \bar{Y}_s, a_s) \triangleq e^{-\rho s} U(w(\bar{Y}_s), a_s)$ and $W(T, \bar{Y}_T) \triangleq e^{-\rho T} W(\bar{Y}_T)$ are utilities at time t discounted from time 0. The agent's objective is to find an admissible control process that maximizes the expected reward $E^a [R^a(0)]$ over all admissible controls $a \in \mathcal{A}$. In other words, the agent solves the following problem

$$v_t = \sup_{a \in \mathcal{A}} V^a(t) \triangleq \sup_{a \in \mathcal{A}} E_t^a [R^a(t)] , \quad \text{for all } 0 \leq t \leq T .$$

²⁸More formally, the martingale property holds true because

$$|\hat{\eta}(Y_t - A_t, t) + a_t| \leq K (1 + \|Z^0\|_t) , \quad \text{for all } t \in [0, T] ,$$

with $K = \frac{\sigma^{-1}}{h_0} + 1$ and $\|Z^0\|_t \triangleq \max_{0 \leq s \leq t} |Z^0(s)|$.

The objective function can be recast as

$$V^a(t) = E_t^a [R^a(t)] = E_t [\Lambda_{t,T}^a R^a(t)] , \quad (37)$$

where the operator $E^a[\cdot]$ and $E[\cdot]$ are expectation under the probability measure Q^a and Q , respectively. One can see from (37) that varying a is indeed equivalent to changing the probability measure. The key advantage of the weak formulation is that, under the reference measure Q , the output process does not depend on a . Hence, we can treat it as fixed which enables us to solve our problem in spite of its non-Markovian structure.

Our derivation of the necessary conditions builds on the variational argument in Cvitanić *et al.* (2009). Define the control perturbation

$$a^\varepsilon \triangleq a + \varepsilon \Delta a ,$$

such that there exists an $\varepsilon_0 > 0$ for which any $\varepsilon \in [0, \varepsilon_0)$ satisfy $|a^\varepsilon|^4$, $|U^{a^\varepsilon}|^4$, $|U_a^{a^\varepsilon}|^4$, $|\Lambda_{t,\tau}^{a^\varepsilon}|^4$, $(\mathcal{U}_{t,\tau}^{a^\varepsilon})^2$ and $(\partial_a \mathcal{U}_{t,\tau}^{a^\varepsilon})^2$ being uniformly integrable in $L^1(Q)$ where

$$\mathcal{U}_{t,\tau}^a \triangleq \int_t^\tau U(s, \bar{Y}_s, a_s) ds .$$

We introduce the following shorthand notations for "variations"

$$\nabla \mathcal{U}_{t,\tau}^a \triangleq \int_t^\tau U_a(s, \bar{Y}_s, a_s) \Delta a_s ds , \quad (38)$$

$$\nabla A_t \triangleq \int_0^t \Delta a_s ds , \quad (39)$$

$$\begin{aligned} \nabla \Lambda_{t,\tau}^a &\triangleq \Lambda_{t,\tau}^a \left(\frac{1}{\sigma} \right) \left[\int_t^\tau \left(-\frac{\sigma^{-2}}{h_s} \nabla A_s + \Delta a_s \right) dZ_s^0 - \int_t^\tau (\hat{\eta}_s + a_s) \left(-\frac{\sigma^{-2}}{h_s} \nabla A_s + \Delta a_s \right) ds \right] \\ &= \Lambda_{t,\tau}^a \left(\frac{1}{\sigma} \right) \int_t^\tau \left(-\frac{\sigma^{-2}}{h_s} \nabla A_s + \Delta a_s \right) dZ_s^a . \end{aligned} \quad (40)$$

Step 1: We first characterize the variations of the agent's objective with respect to ε

$$\begin{aligned} \frac{V^{a^\varepsilon}(t) - V^a(t)}{\varepsilon} &= E [\Lambda_{t,T}^{a^\varepsilon} R^{a^\varepsilon}(t) - \Lambda_{t,T}^a R^a(t)] \\ &= E \left[\left(\frac{\Lambda_{t,T}^{a^\varepsilon} - \Lambda_{t,T}^a}{\varepsilon} \right) R^{a^\varepsilon}(t) + \Lambda_{t,T}^a \left(\frac{R^{a^\varepsilon}(t) - R^a(t)}{\varepsilon} \right) \right] \\ &= E \left[\nabla \Lambda_{t,T}^{a^\varepsilon} R^{a^\varepsilon}(t) + \Lambda_{t,T}^a \left(\frac{R^{a^\varepsilon}(t) - R^a(t)}{\varepsilon} \right) \right] . \end{aligned}$$

To obtain the limit of the first term as ε goes to zero, observe that

$$\nabla \Lambda_{t,T}^{a^\varepsilon} R^{a^\varepsilon}(t) - \nabla \Lambda_{t,T}^a R^a(t) = [\nabla \Lambda_{t,T}^{a^\varepsilon} - \nabla \Lambda_{t,T}^a] R^a(t) + \nabla \Lambda_{t,T}^{a^\varepsilon} [R^{a^\varepsilon}(t) - R^a(t)] .$$

As shown in Cvitanić *et al.* (2009), for any $\varepsilon \in [0, \varepsilon_0)$, this expression is integrable uniformly with respect to ε and so

$$\lim_{\varepsilon \rightarrow 0} E [\nabla \Lambda_{t,T}^{a^\varepsilon} R^{a^\varepsilon}(t)] = E [\nabla \Lambda_{t,T}^a R^a(t)] .$$

The limit of the second term reads

$$\lim_{\varepsilon \rightarrow 0} \frac{R^{a^\varepsilon}(t) - R^a(t)}{\varepsilon} = e^{\rho t} \nabla \mathcal{U}_{t,T}^a .$$

Due to the uniform integrability of $\Lambda_{t,T}^a (R^{a^\varepsilon}(t) - R^a(t)) / \varepsilon$, the expectation is also well defined. Combining the two expressions above, we finally obtain

$$\lim_{\varepsilon \rightarrow 0} \frac{V^{a^\varepsilon}(t) - V^a(t)}{\varepsilon} = E [\nabla \Lambda_{t,T}^a R^a(t) + \Lambda_{t,T}^a e^{\rho t} \nabla \mathcal{U}_{t,T}^a] \triangleq \nabla V^a(t) . \quad (41)$$

Step 2: We are now in a position to derive the necessary condition. Consider total earnings as of date 0

$$I^a(t) \triangleq E_t^a \left[\int_0^T U(s, \bar{Y}_s, a_s) ds + W(T, \bar{Y}_T) \right] = \int_0^t U(s, \bar{Y}_s, a_s) ds + e^{-\rho t} V^a(t) . \quad (42)$$

By definition, it is a Q^a -martingale. According to the extended Martingale Representation Theorem²⁹ of Fujisaki et al. (1972), all square integrable Q^a -martingales are stochastic integrals of $\{Z_t^a\}$ and there exists a unique process ζ in $L^2(Q^a)$ such that

$$I^a(T) = I^a(t) + \int_t^T \zeta_s \sigma dZ_s^a . \quad (43)$$

This decomposition allows us to solve for $\nabla V^a(t)$. Reinserting (38), (39) and (40) into (41) yields³⁰

$$\begin{aligned} \nabla V^a(t) &= E_t \left[\Lambda_{t,T}^a R^a(t) \sigma^{-1} \int_t^T \left(-\frac{\sigma^{-2}}{h_s} \nabla A_s + \Delta a_s \right) dZ_s^a + \Lambda_{t,T}^a e^{\rho t} \left(\int_t^T U_a \Delta a_s ds \right) \right] \\ &= e^{\rho t} E_t^a \left[I^a(T) \sigma^{-1} \int_t^T \left(-\frac{\sigma^{-2}}{h_s} \nabla A_s + \Delta a_s \right) dZ_s^a + \int_t^T U_a \Delta a_s ds \right] . \end{aligned}$$

²⁹We cannot directly apply the standard Martingale Representation theorem because we are considering weak solutions, so that $\{Z_t^a\}$ does not necessarily generate $\{\mathcal{F}_t^Y\}$.

³⁰The additional expectation term vanishes because both $\left(\frac{h_\varepsilon}{h_s}\right) \nabla A_s$ and Δa_s are bounded and so

$$\left(\int_0^t U(\tau, \bar{Y}_\tau, a_\tau) d\tau \right) E_t^a \left[\int_t^T \left(-\left(\frac{h_\varepsilon}{h_s}\right) \nabla A_s + \Delta a_s \right) dZ_s^a \right] = 0 .$$

where subscripts denote derivatives and arguments are omitted for brevity. Given the law of motion (43), applying Ito's rule to the first term yields

$$d \left(I^a(\tau) \int_t^\tau \left(-\frac{\sigma^{-2}}{h_s} \nabla A_s + \Delta a_s \right) dZ_s^a \right) = \left[\zeta_\tau \sigma \left(-\left(\frac{\sigma^{-2}}{h_\tau} \right) \nabla A_\tau + \Delta a_\tau \right) \right] d\tau \\ + \left[\zeta_\tau \sigma \int_t^\tau \left(-\frac{\sigma^{-2}}{h_s} \nabla A_s + \Delta a_s \right) dZ_s^a + I_t^a(\tau) \left(-\left(\frac{\sigma^{-2}}{h_\tau} \right) \nabla A_\tau + \Delta a_\tau \right) \right] dZ_\tau^a .$$

Hence $\nabla V^a(t)$ can be represented as

$$e^{-\rho t} \nabla V^a(t) = E_t^a \left[\int_t^T \Gamma_s^1 ds + \int_t^T \Gamma_s^2 dZ_s^a \right] ,$$

where

$$\Gamma_s^1 \triangleq \zeta_s \left[-\frac{\sigma^{-2}}{h_s} \int_0^s \Delta a_\tau d\tau + \Delta a_s \right] + U_a(s, \bar{Y}_s, a_s) \Delta a_s , \\ \Gamma_s^2 \triangleq \zeta_s \left[\int_t^s \left(-\frac{\sigma^{-2}}{h_\tau} \int_0^\tau \Delta a_r dr + \Delta a_\tau \right) dZ_\tau^a \right] + I_t^a(s) \left(-\frac{\sigma^{-2}}{h_s} \int_0^s \Delta a_\tau d\tau + \Delta a_s \right) .$$

Given that Γ_s^2 is square integrable,³¹ we have

$$E_t^a \left[\int_t^T \Gamma_s^2 dZ_s^a \right] = 0 .$$

As for the deterministic term, collecting the effect of each perturbation Δa_s yields

$$e^{-\rho t} \nabla V^a(t) = E_t^a \left[\int_t^T \left(-\int_s^T \zeta_\tau \left(\frac{\sigma^{-2}}{h_\tau} \right) d\tau + \zeta_s + U_a(s, \bar{Y}_s, a_s) \right) \Delta a_s ds \right] .$$

Finally, noticing that Δa_s was arbitrary leads to

$$\left(E_t^a \left[-\int_t^T \zeta_s \frac{\sigma^{-2}}{h_s} ds \right] + \zeta_t + U_a(t, \bar{Y}_t, a_t^*) \right) (a_t - a_t^*) \leq 0 . \quad (44)$$

Step 3: We now rewrite our solution as a function of the promised value v_t . Differentiating (42) with respect to time yields

$$e^{-\rho t} dv_t - \rho e^{-\rho t} v_t + U(t, \bar{Y}_t, a_t) = dI^a(t) = \zeta_t \sigma dZ_t^a ,$$

so that

$$dv_t = (\rho v_t - U(\bar{Y}_t, a_t)) dt + \zeta_t \sigma dZ_t^a ,$$

³¹Square integrability of Γ_s^2 can be established for any $\varepsilon \in [0, \varepsilon_0)$ following the same steps as in Lemma 7.3 of Cvitanic *et al.* (2009).

with $\gamma_t \triangleq \zeta_t e^{\rho t}$. Collecting the exponential terms in (44) leads to (11). ■

Proof. Proposition 2: The sufficient conditions are established comparing the equilibrium path $\{a_t^*\}_{t=0}^T$ with an arbitrary effort path $\{a_t\}_{t=0}^T$. We define $\delta_t \triangleq a_t - a_t^*$ and $\Delta_t \triangleq \int_0^t \delta_s ds = A_t - A_t^*$ as the differences in current and cumulative effort between the arbitrary and recommended paths. We also attach a star superscript to denote the value of the \mathbb{F}^Y -measurable stochastic processes along the equilibrium path. The Brownian motions generated by the two effort policies are related by

$$\begin{aligned} \sigma dZ_t^{a^*} &= \sigma dZ_t^a + [\hat{\eta}(Y_t - A_t, t) + a_t - \hat{\eta}(Y_t - A_t^*, t) - a_t^*] dt \\ &= \sigma dZ_t^a + \left[\delta_t - \frac{\sigma^{-2}}{h_t} \Delta_t \right] dt . \end{aligned}$$

By definition, the total reward from the optimal policy reads

$$\begin{aligned} I^{a^*}(T) &= \int_0^T U(t, \bar{Y}_t, a_t^*) dt + W(\bar{Y}_T) = V^{a^*}(0) + \int_0^T \zeta_t^* \sigma dZ_t^{a^*} \\ &= V^{a^*}(0) + \int_0^T \zeta_t^* \left[\delta_t - \frac{\sigma^{-2}}{h_t} \Delta_t \right] dt + \int_0^T \zeta_t^* \sigma dZ_t^a . \end{aligned}$$

Hence, the total reward from the arbitrary policy is given by

$$\begin{aligned} I^a(T) &= \int_0^T [U(t, \bar{Y}_t, a_t) - U(t, \bar{Y}_t, a_t^*)] dt + I^{a^*}(T) \\ &= \int_0^T [U(t, \bar{Y}_t, a_t) - U(t, \bar{Y}_t, a_t^*)] dt + V^{a^*}(0) \\ &\quad + \int_0^T \zeta_t^* \left[\delta_t - \frac{\sigma^{-2}}{h_t} \Delta_t \right] dt + \int_0^T \zeta_t^* \sigma dZ_t^a . \end{aligned}$$

Let us focus on the third term on the right hand side

$$\begin{aligned} - \int_0^T \zeta_t^* \frac{\sigma^{-2}}{h_t} \Delta_t dt &= - \int_0^T \zeta_t^* \frac{\sigma^{-2}}{h_t} \left(\int_0^t \delta_s ds \right) dt = \int_0^T \delta_t \left(- \int_t^T \zeta_s^* \frac{\sigma^{-2}}{h_s} ds \right) dt \\ &= \int_0^T \delta_t \left(e^{-\rho t} \frac{\sigma^{-2}}{h_t} p_t^* + \int_t^T \zeta_s^* \sigma dZ_s^{a^*} \right) dt , \end{aligned}$$

where the last equality follows from the definition of p and ξ .³² Changing the Brown-

³²Observe that this additional step is linked to the introduction of private information. Then the volatility ζ of the continuation value will differ on and off the equilibrium path. To the contrary, in problems without private information, the volatility remains constant because it only depends on observable output and not on past actions. This is why sufficiency holds without restriction in, e.g., Schättler and Sung (1993) or Sannikov (2008).

ian motion and taking expectation yields

$$\begin{aligned}
V^a(0) - V^{a^*}(0) &= E_0^a [I^a(T)] - V^{a^*}(0) \\
&= E_0^a \left[\int_0^T \left(U(t, \bar{Y}_t, a_t) - U(t, \bar{Y}_t, a_t^*) + \delta_t \left(\zeta_t^* + e^{-\rho t} \frac{\sigma^{-2}}{h_t} p_t^* \right) \right) dt \right] \\
&\quad + E_0^a \left[\int_0^T \left(\int_t^T \xi_s^* \left(\delta_s - \frac{\sigma^{-2}}{h_s} \Delta_s \right) ds \right) dt \right] \\
&= E_0^a \left[\int_0^T e^{-\rho t} \left(U(w_t, a_t) - U(w_t, a_t^*) + \delta_t \left(\gamma_t^* + \frac{\sigma^{-2}}{h_t} p_t^* \right) \right) dt \right] \\
&\quad + E_0^a \left[\int_0^T e^{\rho t} \xi_t^* \Delta_t \left(\delta_t - \frac{\sigma^{-2}}{h_t} \Delta_t \right) dt \right].
\end{aligned}$$

We know from the optimization property of a_t^* that the first expectation term is at most equal to zero. On the other hand, the sign of the second expectation term is ambiguous. In order to bound it, we introduce the predictable process³³ $\chi_t^* \triangleq \zeta_t^* - e^{\rho t} \xi_t^* A_t^*$ and define the Hamiltonian function

$$H(t, a, A; \chi^*, \xi^*, p^*) \triangleq U(w, a) + (\chi^* + e^{\rho t} \xi^* A) a - e^{\rho t} \xi^* \frac{\sigma^{-2}}{h_t} A^2 + \frac{\sigma^{-2}}{h_t} p^* a.$$

Taking a linear approximation of the Hamiltonian around A^* yields

$$\begin{aligned}
&H_t(a_t, A_t) - H_t(a_t^*, A_t^*) - \frac{\partial H_t(a_t^*, A_t^*)}{\partial A} \Delta_t \\
&= U(w_t, a_t) - U(w_t, a_t^*) + \delta_t \left(\underbrace{\chi_t^* + e^{\rho t} \xi_t^* A_t^*}_{=\zeta_t^*} + \frac{\sigma^{-2}}{h_t} p_t^* \right) + e^{\rho t} \xi_t^* \Delta_t \left(\delta_t - \frac{\sigma^{-2}}{h_t} \Delta_t \right),
\end{aligned}$$

so that

$$V^a(0) - V^{a^*}(0) = E_0^a \left[\int_0^T e^{-\rho t} \left(H_t(a_t, A_t) - H_t(a_t^*, A_t^*) - \frac{\partial H_t(a_t^*, A_t^*)}{\partial A} \Delta_t \right) dt \right]$$

is negative when the Hamiltonian function is jointly concave. Given that the agent seeks to maximize expected returns, imposing concavity ensures that a^* dominates any alternative effort path. Concavity is established considering the Hessian matrix of the Hamiltonian

$$\mathcal{H}(t, a, A) = \begin{pmatrix} U_{aa}(w_t, a_t) & e^{\rho t} \xi_t \\ e^{\rho t} \xi_t & -2e^{\rho t} \xi_t \frac{\sigma^{-2}}{h_t} \end{pmatrix},$$

which is negative semi-definite when $-2\frac{\sigma^{-2}}{h_t} U_{aa}(w_t, a_t) \geq e^{\rho t} \xi_t$, as stated in (16). ■

³³ χ^* is predictable since both ξ^* and A^* are \mathbb{F}^Y -predictable.

Proof. Corollary 1: Let b_t be defined as

$$b_t \triangleq E \left[- \int_0^T e^{-\rho s} \gamma_s \left(\frac{h_\varepsilon}{h_s} \right) ds \middle| \mathcal{F}_t^a \right] = b_0 + \int_0^t \xi_s \sigma dZ_s, \text{ for all } t \in [0, T],$$

where the second equality follows from (17). Then the definition of p_t in (15) implies that

$$p_t = e^{\rho t} \sigma^2 h_t \left[b_t + \int_0^t e^{-\rho s} \gamma_s \frac{\sigma^{-2}}{h_s} ds \right],$$

and so, as T goes to infinity, p_t solves the SDE³⁴

$$dp_t = \left[\rho p_t + \frac{d(\sigma^2 h_t)}{dt} \frac{\sigma^{-2}}{h_t} p_t + \gamma_t \right] dt + e^{\rho t} \sigma^2 h_t db_t = \left[p_t \left(\rho + \frac{\sigma^{-2}}{h_t} \right) + \gamma_t \right] dt + \vartheta_t \sigma dZ_t,$$

with $\vartheta_t \triangleq e^{\rho t} \sigma^2 h_t \xi_t$. ■

Proof. Proposition 3: Assume that

$$\begin{aligned} j(v, t) &= \left(\frac{e^{-\rho t}}{\rho} \right) [j_0(t) + j_1 \ln(-v)], \\ w(t, v) &= -\frac{\ln(k_t v)}{\theta} + \lambda \Rightarrow U(w, 1) = -k_t v. \end{aligned}$$

Observe first that our guess implies that

$$\gamma_t(v, w, a) = -U_a(w(t, v), 1) - \frac{\sigma^{-2}}{h_t} \lambda \theta v = -\lambda \theta U(w(t, v), 1) - \frac{\sigma^{-2}}{h_t} \lambda \theta v = \lambda \theta v \left(k_t - \frac{\sigma^{-2}}{h_t} \right).$$

Hence, differentiating the Incentive Constraint yields

$$\frac{\partial \gamma_t(v, w, a)}{\partial w} = -U_{aw}(w, a) = -\theta \gamma_t(v, w, a) - \frac{\sigma^{-2}}{h_t} \lambda \theta^2 v_t.$$

Therefore, the FOC for wages is equivalent to

$$\begin{aligned} & -e^{-\rho t} - \frac{\partial j}{\partial v} \theta v k_t - \sigma^2 \frac{\partial^2 j}{\partial v^2} \left[\left(\lambda \theta v \left(k_t - \frac{\sigma^{-2}}{h_t} \right) \right)^2 + (\lambda \theta v)^2 \left(k_t - \frac{\sigma^{-2}}{h_t} \right) \frac{\sigma^{-2}}{h_t} \right] \theta \\ &= \left(\frac{e^{-\rho t}}{\rho} \right) \left(-\rho - j_1 \theta k_t + \sigma^2 j_1 \left[\left(k_t - \frac{\sigma^{-2}}{h_t} \right)^2 + \left(k_t - \frac{\sigma^{-2}}{h_t} \right) \frac{\sigma^{-2}}{h_t} \right] \lambda^2 \theta^3 \right) \\ &= \left(\frac{e^{-\rho t}}{\rho} \right) \left(-\rho - j_1 \theta k_t + \sigma^2 j_1 \left[k_t \left(k_t - \frac{\sigma^{-2}}{h_t} \right) \right] \lambda^2 \theta^3 \right) = 0, \end{aligned}$$

³⁴The change with respect to time of σ^{-2}/h_t is given by

$$\frac{d(\sigma^{-2}/h_t)}{dt} = \frac{d(\sigma^{-2}(h_0 + t\sigma^{-2})^{-1})}{dt} = -\sigma^{-4}(h_0 + t\sigma^{-2})^{-2} = -\left(\frac{\sigma^{-2}}{h_t} \right)^2 < 0.$$

implying the following quadratic equation for k_t

$$-\rho - k_t \left(j_1 \theta + \sigma^2 j_1 \frac{\sigma^{-2}}{h_t} \lambda^2 \theta^3 \right) + k_t^2 (\sigma^2 j_1 \lambda^2 \theta^3) = 0 . \quad (45)$$

The remaining step consists in checking that the HJB equation is indeed satisfied

$$\begin{aligned} & e^{-\rho t} (1 - w) + \frac{\partial j}{\partial t} + \frac{\partial j}{\partial v} (\rho v - U(w, 1)) + \left(\frac{\sigma^2}{2} \right) \frac{\partial^2 j}{\partial v^2} \gamma^2 \\ = & e^{-\rho t} \left[\left(1 + \frac{\ln(-v)}{\theta} + \frac{\ln(-k_t)}{\theta} - \lambda \right) - [j_0(t) + j_1 \ln(-v)] + \left(\frac{1}{\rho} \right) j_0'(t) \right. \\ & \left. + \left(\frac{1}{\rho} \right) j_1 (\rho + k_t) - \left(\frac{\sigma^2}{2} \right) \left(\frac{1}{\rho} \right) j_1 \left(\lambda \theta \left(k_t - \frac{\sigma^{-2}}{h_t} \right) \right)^2 \right] = 0 , \end{aligned}$$

when $j_1 = \theta^{-1}$ and

$$j_0'(t) - \rho j_0(t) = -\rho \left(1 - \lambda + \frac{\ln(-k_t)}{\theta} \right) - \frac{\rho + k_t}{\theta} + \frac{\theta (\sigma \lambda)^2}{2} \left(k_t - \frac{\sigma^{-2}}{h_t} \right)^2 . \quad (46)$$

The quadratic equation (24) is obtained reinserting j_1 in (45)

$$-\rho - k_t \left(1 + \frac{(\lambda \theta)^2}{h_t} \right) + k_t^2 (\sigma \lambda \theta)^2 = 0 .$$

The relevant solution is unique and given by the negative root because wages are not well defined when $k_t > 0$. The ODE described in the Proposition is obtained noticing that the quadratic equation above implies that

$$\frac{(\sigma \theta \lambda)^2}{2} \left(k_t - \frac{\sigma^{-2}}{h_t} \right)^2 = \rho + k_t + \frac{(\sigma \theta \lambda)^2}{2} \left(\left(\frac{\sigma^{-2}}{h_t} \right)^2 - k_t^2 \right) ,$$

and reinserting this expression into (46).

As usual, the unique solution to the ODE is pinned down by its boundary condition. The value function as $t \rightarrow \infty$ must converge to the solution of the problem without parameter uncertainty, i.e., when h_t is infinite. It can be derived solving the following HJB

$$0 = \max_{\{a, w\}} \left\{ e^{-\rho t} (a - w) + \frac{\partial l}{\partial t} + \frac{\partial l}{\partial v} (\rho v - U(w, a)) + \left(\frac{\sigma^2}{2} \right) \frac{\partial^2 l}{\partial v^2} \gamma(v, w, a)^2 \right\} ,$$

with

$$\gamma(v, w, a) \geq -U_a(a, w) , \text{ for all } a > 0 .$$

The solution is of the form $\rho l(t, v) = e^{-\rho t} \left[l_0 + \frac{\ln(-v)}{\theta} \right]$ with

$$\rho l_0 = \rho \left(1 - \lambda + \frac{\ln(-k_\infty)}{\theta} \right) + \frac{\theta (\sigma \lambda)^2}{2} k_\infty^2 ,$$

where $k_\infty = \lim_{t \rightarrow \infty} k(t) = \left(\frac{1}{2(\sigma\lambda\theta)^2}\right) (1 - \sqrt{1 + 4\rho})$. One can easily verify that the desired convergence of $j_0(t)$ to l_0 as $t \rightarrow \infty$ holds true when the boundary condition $\lim_{t \rightarrow \infty} j'_0(t) = 0$ is satisfied. ■

Proof. Corollary 2: Given that $\vartheta_t^* = \lambda\theta\gamma_t^*(v) = (\lambda\theta)^2 v \left(k_t - \frac{\sigma^{-2}}{h_t}\right)$ and $U_{aa}(w_t, a_t^*) = (\lambda\theta)^2 (-k_t v)$, the sufficient condition of Proposition 2 are satisfied when

$$2k_t v - v \left(k_t - \frac{\sigma^{-2}}{h_t}\right) = v \left(k_t + \frac{\sigma^{-2}}{h_t}\right) > 0 \Leftrightarrow -k_t > \frac{\sigma^{-2}}{h_t}. \quad (47)$$

The explicit solution of the quadratic equation for k_t reads

$$2k_t = \frac{1}{(\sigma\lambda\theta)^2} + \frac{\sigma^{-2}}{h_t} - \sqrt{\left(\frac{1}{(\sigma\lambda\theta)^2} + \frac{\sigma^{-2}}{h_t}\right)^2 + \frac{4\rho}{(\sigma\lambda\theta)^2}}, \quad (48)$$

and so

$$\frac{dk(t)}{dt} = \left(\frac{1}{2}\right) \left[1 - \frac{\frac{1}{(\sigma\lambda\theta)^2} + \frac{\sigma^{-2}}{h_t}}{\sqrt{\left(\frac{1}{(\sigma\lambda\theta)^2} + \frac{\sigma^{-2}}{h_t}\right)^2 + \frac{4\rho}{(\sigma\lambda\theta)^2}}} \right] \underbrace{\frac{d(\sigma^{-2}h_t^{-1})}{dt}}_{<0} < 0. \quad (49)$$

Since σ^{-2}/h_t is decreasing in t , condition (47) is satisfied for all t provided that $-k_0 > \sigma^{-2}/h_0$, i.e.

$$-\frac{1}{(\sigma\lambda\theta)^2} - 3 \left(\frac{\sigma^{-2}}{h_0}\right) + \sqrt{\left(\frac{1}{(\sigma\lambda\theta)^2} + \left(\frac{\sigma^{-2}}{h_0}\right)\right)^2 + \frac{4\rho}{(\sigma\lambda\theta)^2}} > 0,$$

which, after some straightforward simplifications, leads to the requirement (27). ■

Proof. Corollary 3: The statement immediately follows from

$$\frac{1}{2} > \frac{dk(\sigma^{-2}/h_t)}{d(\sigma^{-2}/h_t)} = \left(\frac{1}{2}\right) \left[1 - \frac{\frac{1}{(\sigma\lambda\theta)^2} + \frac{\sigma^{-2}}{h_t}}{\sqrt{\left(\frac{1}{(\sigma\lambda\theta)^2} + \frac{\sigma^{-2}}{h_t}\right)^2 + \frac{4\rho}{(\sigma\lambda\theta)^2}}} \right] > 0,$$

and the solution for wages $w_t^*(v) = -\ln(k_t v)/\theta + \lambda$. ■

Proof. Corollary 4: Reinserting the law of motion (28) for v into (29) and applying Ito's lemma yields³⁵

$$\begin{aligned} dw_t^* &= -\left(\frac{1}{\theta}\right) \left[\left[\left(\frac{1}{k_t}\right) \frac{dk_t}{dt} + \rho + k_t - \frac{(\sigma\theta\lambda)^2}{2} \left(k_t - \frac{\sigma^{-2}}{h_t}\right)^2 \right] dt + \lambda\theta \left(k_t - \frac{\sigma^{-2}}{h_t}\right) \sigma dZ_t \right] \\ &= -\left(\frac{1}{\theta}\right) \left[\left[\left(\frac{1}{k_t}\right) \frac{dk_t}{dt} - \frac{(\sigma\theta\lambda)^2}{2} \left(\left(\frac{\sigma^{-2}}{h_t}\right)^2 - k_t^2\right) \right] dt + \lambda\theta \left(k_t - \frac{\sigma^{-2}}{h_t}\right) \sigma dZ_t \right]. \end{aligned}$$

³⁵See the proof of Proposition 4 for the intermediate step linking the two equalities.

The statement for the volatility component is established reinserting $dk(\sigma^{-2}/h_t)/d(\sigma^{-2}/h_t)$ into

$$-\frac{d\left(k(t) - \frac{\sigma^{-2}}{h(t)}\right)}{dt} = -\left(\frac{dk(t)}{dt} - \frac{d(\sigma^{-2}/h(t))}{dt}\right) = -\left(\underbrace{\frac{dk(\sigma^{-2}/h(t))}{d(\sigma^{-2}/h(t))}}_{\in(0,1/2)} - 1\right) \underbrace{\frac{d(\sigma^{-2}/h(t))}{dt}}_{<0} < 0.$$

The sign of the deterministic trend is established remembering that the sufficient condition (16) holds if and only if $-k_t > \sigma^{-2}/h_t$. Hence, $(\sigma^{-2}/h_t)^2 - k_t^2 < 0$, and so the trend is negative. ■

Proof. Corollary 5: Let $\Psi(t)$ denote the following function

$$\Psi(t) \triangleq \rho \left((1 - \lambda) + \frac{\ln(-k_t)}{\theta} \right) - \frac{(\sigma\lambda)^2 \theta}{2} (s^2 - k_t^2),$$

so that

$$j'_0(t) - \rho j_0(t) + \Psi(t) = 0.$$

Differentiating $\Psi(t)$ with respect to time yields³⁶

$$\Psi'(t) = \rho \left(\frac{1}{k_t} \right) \frac{dk(t)}{dt} - \frac{(\sigma\lambda)^2 \theta}{2} \left(-\frac{\sigma^{-2}}{h_t} - k_t \frac{dk(t)}{dt} \right) > 0.$$

Hence, if $\rho j_0(t) \leq \Psi(t)$, we have $j'_0(t) < 0$ and so $\rho j_0(\tau) < \Psi(\tau)$ for all $\tau \geq t$. But this contradicts the boundary condition $\lim_{t \rightarrow \infty} j_0(t) = \Psi(t)$. We can therefore conclude that $\rho j_0(t) > \Psi(t)$ which implies in turn that $j'_0(t) > 0$. Given that parameter precision is increasing in time, the claim stated in Corollary 5 follows. ■

Proof. Proposition 4: The proof proceeds one part at a time:

Part (i). This claim holds because the utility function has no wealth effect on the demand for leisure, and results formally because $\hat{\eta}_s$ ($s \leq t$) drops out of the FOC determining a_t . **Parts (ii)(C) and (iii)(A).** We construct a solution to the first-order condition in the way that implies the claims. If the claims (i) and (ii) are correct, since $\partial Y_s / \partial a_t = 1$ for $s \geq t$, the first-order condition for optimal effort at date t is

$$\frac{\partial U}{\partial a} (\hat{\eta}(Y_t - A_t, t), a_t) + \int_t^\infty e^{-\rho(s-t)} \frac{\partial}{\partial Y} E_t [U(\hat{\eta}(Y_s - A_s, s), a_s)] ds \quad \begin{cases} > 0 & \text{if } h_t < h_1 \\ = 0 & \text{if } h_t \in [h_1, h_2] \\ < 0 & \text{if } h_t > h_2 \end{cases} . \quad (50)$$

³⁶Remember that both $dk(t)/dt$ and $k(t)$ are negative.

Now let T be such that $h_T = h_2$. Then, $a_s = 0$ for $s \geq T$ and $A_s = A_T$, implying that at $t = T$ (50) becomes

$$-\frac{\partial U(\hat{\eta}(Y_T - A_T, T), 0)}{\partial a} = \int_T^\infty e^{-\rho(s-T)} \frac{\partial}{\partial Y} E_T [U(\hat{\eta}(Y_s - A_T, s), 0)] ds. \quad (51)$$

It is shown in the “final part” of the proof below that (51) is equivalent to

$$\lambda = \int_0^\infty e^{-\rho\tau} \underbrace{\left[\frac{\frac{1}{\sigma^2}}{h_T + \tau \frac{1}{\sigma^2}} \exp\left(\frac{\theta^2}{2} \left(\frac{\frac{1}{\sigma^2}}{h_T + \tau \frac{1}{\sigma^2}}\right)^2 (\tau\sigma^2 + h_T^{-1})\right)\right]}_{=g(\tau;T)} d\tau, \quad (52)$$

which does not depend on the posterior $\hat{\eta}_T$. This implies that the stopping time T does not vary with cumulative output Y_T . Since that $g(\tau; T)$ is strictly decreasing in T , the equality can be satisfied only for at most a single T . The RHS of (52) is continuous in T , and $\lim_{T \rightarrow \infty} g(\tau; T) = 0$. Therefore a solution for T exists if $\lambda < \int g(\tau, T) d\tau$, i.e., if (33) holds. This proves (iii)(A).and (ii)(C).

Part (ii)(B). Since a_t is continuous in t , there exists a $\delta > 0$ such that optimal effort is interior, i.e. $a_t \in (0, 1)$ for all $t \in (T - \delta, T)$. Similar steps as before (and also reported below in the “final part” of the proof) yield

$$\lambda f(t) = \int_0^\infty e^{-\rho\tau} f(t + \tau) \underbrace{\left[\frac{\frac{1}{\sigma^2}}{h_t + \tau \frac{1}{\sigma^2}} \exp\left(\frac{\theta^2}{2} \left(\frac{\frac{1}{\sigma^2}}{h_t + \tau \frac{1}{\sigma^2}}\right)^2 (\tau\sigma^2 + h_t^{-1})\right)\right]}_{=g(\tau;t)} d\tau, \quad (53)$$

where $f(t) = \exp(\lambda\theta a_t)$. Differentiating (53) yields

$$\lambda f'(t) = \int_0^\infty e^{-\rho\tau} \left[f'(t + \tau)g(\tau; t) + f(t + \tau) \frac{\partial g(\tau; t)}{\partial t} \right] d\tau.$$

Given that $\partial g(\tau; t)/\partial t < 0$ and that both $f(\cdot)$ and $g(\cdot)$ are nonnegative, if (a) (53) holds as an exact equality and if (b) $f(t) > 0$, then

$$f'(t + \tau) \leq 0 \text{ for } \tau > 0 \implies f'(t) < 0.$$

That is, a sufficient condition for the derivative at time t to be negative is that it is at most zero afterwards. This is easily established considering the limit as t goes to T . First, we know that $f'(T + \tau) = 0$ for all $\tau > 0$. Furthermore, since T is unique, $f(t) > 0$ for $t \in (T - \delta, T)$. Iterating this argument we conclude that a_t must be decreasing

$$a_t \in (0, 1) \implies f'(t) = \exp(\lambda\theta a_t) \lambda\theta \dot{a}_t < 0.$$

Part (iii)(C). If h_0 is small enough so that $h_0 < h_1$, we shall end up at the upper

bound. Since

$$\frac{\partial^2}{\partial Y_t \partial a_t} E_0 [U(\hat{\eta}(Y_t - A_t, t), a_t)] = \lambda^2 \theta \frac{\sigma^{-2}}{h} U < 0 ,$$

a sufficient condition for an initial horizontal segment in Figure 5 to exist is that (34) should hold.

Final part of the proof. It remains for us to show that (51) implies (52) and (53).

Derivation of (52).— First we show that (51) implies

$$\lambda \exp(-\theta \hat{\eta}(Y_T - A_T, T)) = \int_T^\infty e^{-\rho(s-t)} \frac{\sigma^{-2}}{h_s} E_T [U(\hat{\eta}(Y_s - A_T, s), 0)] ds . \quad (54)$$

Observe that

$$\frac{\partial}{\partial Y} E_t [U(\hat{\eta}(Y_s - A_s, s), a_s)] = E_t [U(\hat{\eta}(Y_s - A_s + \Delta, s), a_s)] \left(-\theta \frac{\sigma^{-2}}{h_s} \right) .$$

Hence, since

$$\begin{aligned} \hat{\eta}(Y_s - A_s, s) &= \frac{h_0 m_0 + \frac{1}{\sigma^2} (Y_s - A_s) h_t}{h_t} \frac{h_t}{h_s} + \frac{Y_{s-t} - A_{s-t}}{\sigma^2 h_s} \\ &= \hat{\eta}(Y_t - A_t, t) + \frac{\sigma^{-2}}{h_s} (Y_{s-t} - A_{s-t} - \hat{\eta}(Y_t - A_t, t) (s - t)) , \end{aligned}$$

we have for any $s \geq t$

$$\begin{aligned} &E_t [U(\hat{\eta}(Y_s - A_s, s), a_s)] \\ &= \exp(-\theta \hat{\eta}(Y_t - A_t, t)) E_t \left[-\exp \left(-\theta \left(\frac{\sigma^{-2}}{h_s} [Y_{s-t} - A_{s-t} - \hat{\eta}(Y_t - A_t, t) (s - t)] - \lambda a_s \right) \right) \right] z . \end{aligned}$$

Reinserting this expression into the RHS of (54) and rearranging yields

$$\lambda = \int_T^\infty e^{-\rho(s-t)} \frac{\sigma^{-2}}{h_s} E_T \left[\exp \left(-\theta \left(\frac{\sigma^{-2}}{h_s} [Y_{s-T} - \hat{\eta}(Y_T - A_T, T) (s - T)] \right) \right) \right] ds .$$

The expectation is derived noticing that the distribution of $Y_{s-T} = \int_T^s dY_\tau$ can be expressed as

$$\varphi_Y(Y_{s-T} | \hat{\eta}_T) = \int \varphi_Y(Y_{s-T} | \eta_T) \varphi_\eta(\eta_T) d\eta = N((s - T) \hat{\eta}_T, (s - T) \sigma^2 + h_T^{-1})$$

because $\varphi_Y(Y_{s-T} | \eta_T) = N((s - T) \eta_T, (s - T) \sigma^2)$ and $\varphi_\eta(\eta_T) = N(\hat{\eta}_T, h_T^{-1})$. Hence the expectation is taken over a lognormally distributed variable so that

$$E_T \left[\exp \left(-\theta \left(\frac{\sigma^{-2}}{h_s} [Y_{s-T} - \hat{\eta}_T (s - T)] \right) \right) \right] = \exp \left(\frac{\theta^2}{2} \left(\frac{\sigma^{-2}}{h_s} \right)^2 [(s - T) \sigma^2 + h_T^{-1}] \right) .$$

The optimality condition is therefore given by (52).

Derivation of (53).— We have

$$\begin{aligned} -\frac{\partial U(\hat{\eta}(Y_t - A_t, t), a_t)}{\partial a} &= \int_t^\infty e^{-\rho(s-t)} \frac{\partial}{\partial Y} E_t [U(\hat{\eta}(Y_s - A_s, s), a_s)] ds, & \text{i.e.,} \\ \lambda \exp(\lambda \theta a_t) &= \int_t^\infty e^{-\rho(s-t)} \frac{\sigma^{-2}}{h_s} \exp(\lambda \theta a_s) \exp\left(\frac{\theta^2}{2} \left(\frac{\sigma^{-2}}{h_s}\right)^2 [(s-t)\sigma^2 + h_t^{-1}]\right) ds, \end{aligned}$$

which, upon a change of variable to $\tau = s - t$ can be rewritten as (53). ■

Appendix B: Additional results

Derivation of (15): We first change variable and define $\tilde{p}_t \triangleq (\sigma^{-2}/h_t) p_t$. Then $\tilde{p}_t = -E \int_t^T e^{-\rho(s-t)} \gamma_s \frac{\sigma^{-2}}{h_s} ds$, so that differentiating with respect to time leads to

$$\frac{d\tilde{p}_t}{dt} = \rho \tilde{p}_t + \frac{\sigma^{-2}}{h_t} \gamma_t = \rho \tilde{p}_t - \frac{\sigma^{-2}}{h_t} (U_a(w_t, a_t) + \tilde{p}_t),$$

where the second equality follows after substitution of $\gamma_t = -U_a(w_t, a) - \tilde{p}_t$. Integrating this expression, we obtain

$$\tilde{p}_t = E_a \left[\int_t^T e^{-\rho(s-t) + \int_t^s \frac{\sigma^{-2}}{h_\tau} d\tau} \frac{\sigma^{-2}}{h_s} U_a(w_s, a_s) ds \right].$$

To simplify the integral in the exponent, we observe that

$$\frac{\sigma^{-2}}{h_\tau} = \frac{\sigma^{-2}}{h_0 + \tau \sigma^{-2}} = \frac{d \ln h_t}{d\tau} \implies \exp\left(\int_t^s \frac{\sigma^{-2}}{h_\tau} d\tau\right) = \exp(\ln h_s - \ln h_t) = \frac{h_s}{h_t}.$$

Therefore

$$\tilde{p}_t = E_a \left[\int_t^T e^{-\rho(s-t)} \left(\frac{h_s}{h_t}\right) \left(\frac{\sigma^{-2}}{h_s}\right) U_a(w_s, a_s) ds \right] = \frac{\sigma^{-2}}{h_t} E_a \left[\int_t^T e^{-\rho(s-t)} U_a(w_s, a_s) ds \right],$$

which, given the definition of \tilde{p}_t , is equivalent to (15). Observe, however, that when $a_t = 0$ for some t then (13) is not representable as (15).

Extending the HJB eq.(19) to include $\hat{\eta}$: The HJB equations defined in (19) and (21) can be extended to include $\hat{\eta}$ and would still be satisfied. To see this, define $X_t \triangleq Y_t - A_t$ and $g(X_t, t) \triangleq e^{-\rho t} \hat{\eta}(X_t, t)/\rho$. This function satisfies the HJB equations below because

$$\begin{aligned} e^{-\rho t} \hat{\eta}(X_t, t) + \frac{\partial g}{\partial t} + \hat{\eta}_{X_t}(X_t, t) \frac{\partial g}{\partial X_t} + \frac{\sigma_t^2}{2} \frac{\partial^2 g}{\partial X_t^2} &= e^{-\rho t} \left[\begin{array}{c} \hat{\eta}(X_t, t) - \hat{\eta}(X_t, t) \\ + \frac{1}{\rho} \hat{\eta}_t(X_t, t) + \frac{1}{\rho} \hat{\eta}_{X_t}(X_t, t) \hat{\eta}(X_t, t) \end{array} \right] \\ &= \left(\frac{e^{-\rho t}}{\rho}\right) [\hat{\eta}_t(X_t, t) + \hat{\eta}_{X_t}(X_t, t) \hat{\eta}(X_t, t)] = 0, \end{aligned}$$

where the last equality follows from (7).

Appendix C: Details of the piece-rate simulation in Figure 8

As explained in the text, we simulate the piece rate model using a discrete-time solution. We consider agents with a finite lifetime horizon T and establish the properties of interest when T goes to infinity.

Last Period.— The Zero Profit Condition (ZPC hereafter) on the RHS of (36) holds when

$$b_{0,T} = (1 - b_{1,T}) E [y_T | y^T] = (1 - b_{1,T}) (\hat{\eta}_T + a_T^*) ,$$

where $y^T \triangleq \{y_0, y_1, \dots, y_{T-1}\}$ denote the output history at the beginning of period T . For the utility function in (20), the agent's utility is maximized when he provides full effort $a_T = 1$, which is incentive compatible iff $b_{1,T} \geq \lambda$. Minimizing the income variance yields

$$\begin{aligned} b_{0,T} &= (1 - \lambda) (\hat{\eta}_T + 1) , \\ b_{1,T} &= \lambda . \end{aligned}$$

Previous Periods.

Claim 1 *The sequence $\{b_{1,t}\}_{t=1}^T$ is deterministic. Hence, output history and cross-agent differences in beliefs $\hat{\eta}_t$ affect only the mean, not the variance of wages*

Proof. The proof is established recursively. From the discussion above, we know that $b_{1,T} = \lambda$, independently of the output history. We now establish that if $\{b_{1,s}\}_{s=t+1}^T$ is deterministic, so is $b_{1,t}$. By the definition of preferences and by the ZPC

$$U(w_s, a_s) = \exp(-\theta [(1 - b_{1,s}) (\hat{\eta}_s + a_s^*) + b_{1,s} y_s - \lambda a_s]) ,$$

where recommended a_s^* and actual a_s efforts are allowed to differ. Given that y_s is independent of a_t for all $s > t$, we have³⁷

$$\frac{\partial U(w_s, a_s)}{\partial a_t} = -\theta \left(\frac{\partial \hat{\eta}_s}{\partial a_t} \right) (1 - b_{1,s}) U(w_s, a_s) + \theta \left(\frac{\partial b_{1,s}}{\partial a_t} \right) \varepsilon_s U(w_s, a_s) .$$

The second term on the RHS is equal to zero under the premise that $\{b_{1,s}\}_{s=t+1}^T$ is deterministic. Then after dividing by θ , the agent's FOC reads

$$(b_1 - \lambda) E_{t-1} [U_t] + \sum_{s=t+1}^T \beta^{s-t} \frac{h_\varepsilon}{h_s} (1 - b_{1,s}) E_{t-1} [U_s] = 0 . \quad (55)$$

³⁷Remember that we define output y_t in period as

$$y_t = a_t + \eta + \varepsilon_t .$$

Observe that the premise is again required in order to take $(1 - b_{1,s})$ out of the expectation term. Because the optimal contract minimizes the variance of income, the agent's FOC also defines the optimal indexation to performance $b_{1,t}$. Rearranging yields the simplified optimality condition

$$b_{1,t} = \max(0, \lambda - R_t) , \quad (56)$$

which, after a certain date, will always admit an interior solution where

$$R_t = \sum_{s=t+1}^T \beta^{s-t} \frac{h_\varepsilon}{h_s} (1 - b_{1,s}) \frac{E_{t-1}[U_s]}{E_{t-1}[U_t]} \quad (57)$$

is the reputational concern. The ZPC implies that in every period

$$b_{0,t} = (1 - b_{1,t}) (\hat{\eta}_t + a_t^*) ,$$

and so utilities along the equilibrium path are equal to

$$U(w_s, a_s) = \exp(-\theta [\hat{\eta}_t + b_{1,t} \varepsilon_t + a_t^* (1 - \lambda)]) .$$

According to our parametric assumption, conditional on beginning-of-date- t information,

$$-\theta [\hat{\eta}_s + b_{1,s} \varepsilon_s] \sim N(-\theta \hat{\eta}_t, \theta^2 (h_t^{-1} - h_s^{-1} + b_{1,s}^2 \sigma_\varepsilon^2)) .$$

Furthermore, we know that full effort is sustainable at time T . Since incentives are more easily provided in previous periods due to reputational concerns, full effort is implementable at all $t \leq T$, and will be recommended because the higher the action, the better off the agent is. Hence, we can set $a_t^* = 1$ for all $t \leq T$, implying that

$$E_{t-1}[U_s] = -\exp\left(-\theta \hat{\eta}_t + \frac{\theta^2}{2} (h_t^{-1} - h_s^{-1} + b_{1,s}^2 \sigma_\varepsilon^2)\right) \exp(1 - \lambda) ,$$

and, since

$$E_{t-1}[U_t] = -\exp\left(-\theta \hat{\eta}_t + \frac{\theta^2}{2} b_{1,t}^2 \sigma_\varepsilon^2\right) \exp(1 - \lambda) ,$$

we have

$$\frac{E_{t-1}[U_s]}{E_{t-1}[U_t]} = \exp\left(\frac{\theta^2}{2} (h_t^{-1} - h_s^{-1} + [b_{1,s}^2 - b_{1,t}^2] \sigma_\varepsilon^2)\right) .$$

Substituting into (57) we finally obtain

$$R_t = \sum_{s=t+1}^T \beta^{s-t} \frac{h_\varepsilon}{h_s} (1 - b_{1,s}) \exp\left(\frac{\theta^2}{2} (h_t^{-1} - h_s^{-1} + [b_{1,s}^2 - b_{1,t}^2] \sigma_\varepsilon^2)\right) , \quad (58)$$

which is independent of output history. Given that $b_{1,t} = 1 - R_t$, the claim is proved. \blacksquare

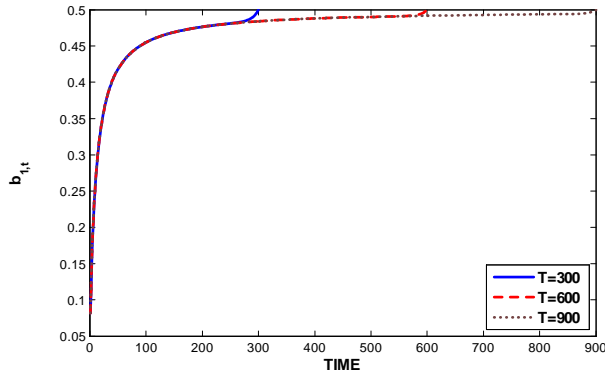


Figure 10: A PLOT OF $b_{1,t}$ FOR THREE SEPARATE HORIZONS

Mean wages and the standard-deviation band.— Since equilibrium effort $a_t = 1$, mean wages are unity when $\hat{\eta}_t = 0$. As Chernoff (1968) shows, the variance of $\hat{\eta}_t$ is $h_0^{-1} - h_t^{-1}$. The piece-rate variance is $\sigma b_{1,t}$. Therefore the one-Standard Deviation band reads

$$\sqrt{h_0^{-1} - h_t^{-1} + \sigma^2 b_{1,t}^2} \quad (59)$$

For Holmström's model, the one-Standard Deviation band is narrower and equal to

$$\sqrt{h_0^{-1} - h_t^{-1}} \quad (60)$$

Comparison to full commitment.—Figure 8 compares the above to a continuous-time formulation with (ρ, σ) given. Taking period length Δ , the discrete-time piece-rate model chooses the discount factor

$$\sigma_\varepsilon^2 = \Delta \sigma^2 \quad \text{and} \quad \beta = \frac{1}{1 + \rho \Delta}$$

and solve the discrete case for $\Delta = 1$. A simulation is done in Figure 10, with three horizons separately, $T = 300, 600$ and 900 , respectively. As we had asserted when discussing the results, the simulations show that $b_{1,t}$ converges rather slowly to its limit of $\lambda = 0.5$. Furthermore, the simulated values in the early periods hardly depend on the horizon length thereby justifying our approximation of the infinite horizon problem.