

Dynamic Managerial Compensation: a Mechanism Design Approach*

Daniel Garrett
Northwestern University
dfgarrett@gmail.com

Alessandro Pavan
Northwestern University
alepavan@northwestern.edu

This version: June 2009
("still" PRELIMINARY and INCOMPLETE)

Abstract

We characterize the optimal incentive scheme for a manager who faces costly effort decisions and whose ability to generate profits for the firm varies stochastically over time. The optimal contract is obtained as the solution to a dynamic mechanism design problem with hidden actions and persistent shocks to the agent's productivity. When the agent is risk-neutral, the optimal contract can often be implemented with a simple pay package that is linear in the firm's profits. Furthermore, the power of the incentive scheme typically increases over time, thus providing a possible justification for the frequent practice of putting more stocks and options in the package of managers with a longer tenure in the firm. In contrast to other explanations proposed in the literature (e.g., declining disutility of effort or career concerns), the optimality of seniority-based reward schemes is not driven by variations in the agent's preferences or in his outside option. It results from an optimal allocation of the manager's informational rents over time. Building on the insights from the risk-neutral case, we then explore the properties of optimal incentive schemes for risk-averse managers. We find that, other things equal, risk-aversion reduces the benefit of inducing higher effort over time. Whether (risk-averse) managers with a longer tenure receive more or less high-powered incentives than younger ones then depends on the interaction between the degree of risk aversion and the dynamics of the impulse responses for the shocks to the manager's type.

JEL classification: D82

Keywords: dynamic mechanism design, adverse selection, moral hazard, incentives, optimal pay scheme, risk-aversion, stochastic process.

*We thank Mike Fishman, Paul Grieco, Igal Hendel, Bill Rogerson, and seminar participants at Northwestern, Toronto (CETC), and LBS. A special thank is to Yuliy Sannikov, for useful comments and suggestions.

1 Introduction

This paper contributes to the literature on managerial compensation by adopting a mechanism design approach to characterize the dynamics of the optimal incentive scheme.

We consider an environment in which the firm's shareholders (the principal) hire a manager (the agent) whose ability to generate profits for the firm varies stochastically over time. This could reflect, for example, the possibility that the value of the manager's expertise/competence changes in response to variations in the business environment. It could also be the result of learning by doing. We assume that both the manager's ability to generate profits (his type) as well as his effort choices are the manager's private information. The firm's shareholders simply observe the dynamics of profits (equivalently, the value of their shares), which we assume to be verifiable, and pay the manager on the basis of this information.

Contrary to the literature on renegotiation (e.g. Laffont and Tirole, 1988, 1990), we assume that the firm's shareholders perfectly understand the value of commitment and hence adhere to the incentive scheme they offered when they hired the manager, even if, after certain contingencies, such a scheme need not be optimal anymore. However, contrary to this literature, we do not impose restrictions on the process governing the evolution of the agent's private information. In particular, we do not restrict the agent's type to be constant over time, nor do we restrict the agent's types to be independent. Allowing for general processes is important for it permits us to shed light on certain properties of the optimal scheme that are obscured, if not completely eliminated, by assuming perfectly correlated, or independent types (more below).

Our baseline model features an environment where both the firm's shareholders and the manager are risk-neutral. Because the firm contracts with the manager at the time the latter is already privately informed about his type, interesting dynamics emerge even without introducing risk aversion. In particular, we show that the power of incentives typically increases over time, which can explain the frequent practice of putting more stocks and options in the pay package of managers with a longer tenure in the firm. Contrary to other explanations proposed in the literature (e.g., declining disutility of effort or career concerns), in our model, the optimality of seniority-based

reward schemes is not driven by variations in the agent’s preferences, nor by variations in his outside option. It results from an optimal allocation of the manager’s informational rents over time. In other words, it originates in the firm’s desire to minimize the manager’s compensation while preserving his incentives for both effort and information revelation.

The driving assumption behind this result is that the effect of the manager’s initial type on the distribution of his future types (which we call the impulse response) declines over time. This assumption is satisfied, for instance, when the agent’s private information evolves according to an ARIMA process with impulse responses smaller than one. As discussed also in other recent works on dynamic mechanism design (e.g. Battaglini, 2005, Pavan, Segal, and Toikka, 2009(a)) this assumption implies that, to minimize the agent’s rents, it is more efficient to distort decisions downwards in the early stages of the relationship than in later ones. The reason is that an agent’s rent, which is provided to ensure he does not mimic another type, depends on the different expectations the two types have about their future types. When this difference declines with the time horizon, distorting decisions in the distant future becomes less effective at reducing informational rents. When applied to the situation studied in this paper, this principle of “vanishing distortions” leads to an effort policy that is closer to the first-best in the long run than in the short run. This follows from the fact that a type’s rent increases in the effort of lower types, as shown by Laffont and Tirole (1986) in a static setting.

A second prediction of the model is that the optimal contract under risk neutrality often takes the form of a simple (state-contingent) linear contract. In other words, in each period, the firm pays the manager a fixed salary plus a bonus that is linear in the firm’s profits (or, equivalently, in the firm’s stock price, provided the latter also depends on the manager’s effort). When the manager’s type follows an ARIMA process (more generally, any process where the impulse responses exhibit a certain separability with respect to the initial type), the slope of the linear scheme changes deterministically over time, i.e. it depends on the manager’s initial type and on the number of periods the manager has been working for the firm, but not on the shocks experienced over time.

More generally, the optimal contract requires that the manager be given the possibility of

proposing changes to his pay package in response to the shocks to his productivity (equivalently, to any privately observed shock to the environment that affects his ability to generate profits for the firm). The idea that a manager may be given the possibility to propose changes to his reward package seems appealing in light of the recent empirical literature on managerial compensation where it is found that this practice has become more frequent in the last decade (see, among others, Kuhnen and Zwiebel (2008), and Bebchuck and Fried (2004)).

While, under risk neutrality, the optimality of linear schemes holds across a variety of specifications of the process governing the evolution of the manager's productivity, there are instances where the optimal effort policy requires the use of stronger incentive schemes according to which the manager is paid a bonus only when the firm's profits exceed a certain threshold, where this threshold may depend on the history of the manager's reports about his type. While the power of these schemes is stronger, contrary to linear schemes, these "bonus" schemes would not be appropriate when profits are the result not only of the manager's type and effort, but also of unobservable noise shocks whose distribution is unaffected by the manager's effort.

Building on the insights from the risk-neutral case, in the second part of the paper we explore the properties of optimal incentive schemes for risk-averse managers. We find that risk-aversion tends to reduce (but not necessarily eliminate) the benefits of seniority-based incentive schemes whose power increases, on average, over time. The reason is that the uncertainty the agent faces about his future productivity given his current productivity increases with the time horizon. In other words, while the agent's current type is a fairly good predictor of his type in the next period, it is a fairly poor predictor of his type, say, five periods into the future. Furthermore, because incentives are forward-looking, the sensitivity of the agent's pay to his productivity in period t is increasing in all future effort levels and is independent of past effort choices. Reducing effort in the far future is thus more effective at reducing the agent's overall exposure to risk than reducing effort in the present or in the near future. Other things equal, risk aversion thus makes it more attractive for the principal to induce higher effort in the early stages of the relationship, when the agent faces little uncertainty about his ability to generate profits, than in later periods, where this

uncertainty (as perceived from the moment the contract is signed) is higher. Whether risk-averse managers with a longer tenure receive more or less high-powered incentive schemes than younger ones then depends on the interaction between the degree of risk-aversion and the impulse responses for the shocks to the manager's type.

Related literature.¹ The literature on managerial compensation is too large to be successfully summarized within the context of this paper. We refer to Prendergast (1999) for an excellent review and to Edmans and Gabaix (2009) for a survey of some recent developments. Of particular interest for our paper is the empirical literature on the use of seniority-based incentive schemes. This literature finds mixed evidence as to the effect of tenure on performance-related pay. While some papers suggest that managers with a longer tenure tend to have weaker incentives and explain this by the fact that the board of directors tends to be captured by CEOs over time (e.g. Hill and Phan, 1991), others point to the contrary (see, e.g., Lippert and Porter, 1997, but also Gibbons and Murphy, 1991). As one would expect, these differences often originate in the choices about which incentives are relevant (e.g. whether to consider stock options). At the theoretical level, our paper contributes to this literature by offering a new trade-off for the optimality of seniority-based incentives that, to the best of our knowledge, was not noticed before.

Obviously the paper is also related to the literature on “dynamic moral hazard” and to its application to dynamic managerial compensation. Seminal works in this literature include Lambert (1983), Rogerson (1985) and Spear and Srivastava (1987). These works provide some qualitative insights about the optimal policy, but do not provide a full characterization. This has been possible only in restricted settings: Phelan and Townsend (1991) characterize optimal policies numerically in a discrete-time model, while Sannikov (2008) uses a continuous-time setting with Brownian shocks to characterize the optimal policy as the solution to a differential equation. In contrast to these results, Holmstrom and Milgrom (1987) show that the optimal contract has a simple structure when (a) the agent does not value the timing of payments, (b) noise follows a Brownian motion and (c) the agent's utility is exponential; under these assumptions, the optimal contract is a simple

¹This part is even more preliminary than the rest. We apologize to those who believe their work should have been cited here and that we omitted to discuss.

linear aggregator of aggregate profit.

Contrary to this literature, we assume that, in each period, the agent observes the shock to his productivity before choosing effort. In this respect, the paper is most closely related to Laffont and Tirole (1986). This alternative approach permits one to use techniques from the mechanism design literature to solve for the optimal contract. In work independent from ours, Edmans and Gabaix (2008) show how this approach can be applied to a dynamic setting, allowing for risk aversion. However, they do not characterize the optimal effort policy, nor which policies are implementable.² Allowing for general processes and characterizing the optimal effort policies is essential to establishing results about the dynamics of the power of incentives and the optimality of linear, or quasi-linear, schemes. Characterizing the optimal effort policy also shows that details about the agent's preferences and the process for the shocks do matter for the structure of the optimal contract.

From a methodological standpoint, our paper uses recent results from the dynamic mechanism design literature to arrive to a characterization of the necessary and sufficient conditions for incentive compatibility. In particular, the approach here builds on the techniques developed in Pavan, Segal, and Toikka (2009,a,b). This paper provides a general treatment of dynamic mechanism design in which the principal has full commitment, and the agent's type may be correlated across time. It extends previous work, for example by Besanko (1985) and Battaglini (2005), to a setting with fairly general payoffs and stochastic processes. We refer the reader to Pavan, Segal, and Toikka (2009,a) for a more extensive review of the dynamic mechanism design literature.

An important dimension in which the paper makes some progress is the characterization of optimal mechanisms under risk aversion and correlated information. In this respect, the paper is also related to the literature on optimal dynamic taxation (also known as Mirrleesian taxation). Battaglini and Coate (2008) consider a discrete-time-two-type model with Markov transitions and show continuity in the optimal mechanism as preferences converge to risk neutrality. Zhang (2009) considers a model with finitely many types, but where contracting occurs in continuous-time and

²Other differences are that (a) they restrict attention to effort policies that depend at most on the current shocks, and (b) they assume contracting occurs at a time the agent does not possess any private information.

where the arrival rate of the transitions between types follows a Poisson process. For most of the analysis, he also restricts attention to two types and finds that many of the results derived for the i.i.d. case (studied, for instance, by Albanesi and Sleet, 2006) carry over to the environment with persistent types. In particular, the celebrated “immiserization result” according to which consumption converges to its lower bound, extends to a setting with correlated types. One qualitative difference with respect to the i.i.d. case is that the size of the “wedges”, i.e. the distortions due to the agent’s private information, is significantly larger when types are persistent. Consistent with Battaglini and Coate (2006), he also finds that, contrary to the risk-neutral case, distortions do not vanish as soon as the agent becomes a high type.

Our results appear broadly consistent with the aforementioned findings from the dynamic optimal taxation literature; however, by allowing for a continuum of types and by considering fairly general stochastic processes, we also uncover patterns of distortions that have not been noticed before (e.g. the possibility that, under risk aversion and sufficiently persistent shocks, effort actually declines over time, as it is the case when productivity follows a random walk). The techniques used to arrive to the characterization of the optimal contract are also different from those in the literature with finitely many types.

Lastly, the paper relates to the literature on the optimal use of financial instruments in dynamic principal-agent relationships. For instance, DeMarzo and Fishman (2007), DeMarzo and Sannikov (2006) and Sannikov (2007)³ study optimal financial contracts for a manager who privately observes the dynamics of cash-flows and can divert funds from investors for private consumption. In these papers it is typically optimal to induce the highest possible effort (which is equivalent to no stealing/no saving); the instrument which is then used to create incentives is the probability of terminating the project. One of the key findings is that the optimal contract can often be implemented using long-term debt, a credit line, and equity. The equity component represents a linear component to the incentive scheme which is used to make the agent indifferent as to whether or not to divert funds for private use. Since the agent’s cost of diverting funds is constant across time

³As in our work, and contrary to the other papers cited here, Sannikov (2007) allows the agent to possess some private information prior to signing the contract. Assuming the agent’s initial type can be either “bad” or “good”, he then characterizes the optimal separating menu where only good types are funded.

and output realizations, so is the equity share. In contrast, we provide an explanation for why and how this share typically changes over time. While these two papers suppose cash-flows are i.i.d., Tchisty (2006) explores the consequences of correlation and shows that the optimal contract can be implemented using a credit line with an interest that increases with the balance. DeMarzo and Sannikov (2008) consider an environment in which both investors and the agent learn about the firm’s true productivity (which evolves according to a Brownian motion). In this paper, as in ours, the agent’s private information is correlated over time.

The rest of the paper is organized as follows. Section 2 presents the baseline model. Section 2.2 characterizes the optimal mechanism. Section 3 extends the analysis to settings where the optimal effort policy is contingent on the entire history of shocks. Section 4 examines optimal schemes for risk-averse agents. All proofs omitted in the text are in the Appendix.

2 The Baseline Model

2.1 The environment

The firm’s shareholders (hereafter referred to as the principal) hire a manager (the agent) to work on a project over T periods, where T may be either finite or infinite. In each period t , the agent receives some private information $\theta_t \in \Theta_t$ about the environment or, equivalently, about his ability to generate profits for the firm, and then chooses effort level $e_t \in E \subset \mathbb{R}$. We will assume that $\Theta_t \subset \mathbb{R}$ is either equal to $[\underline{\theta}_t, \bar{\theta}_t]$ or, in case $\bar{\theta}_t = +\infty$, to $[\underline{\theta}_t, \bar{\theta}_t) \subset \mathbb{R}$ for some $-\infty < \underline{\theta}_t \leq \bar{\theta}_t \leq +\infty$.⁴

To simplify the exposition (and facilitate the characterization of the optimal effort policy) we will assume that $E = \mathbb{R}$.⁵

The principal’s profits π_t in period t , gross of any agent compensation, depend on the sequence of effort decisions $e^t \equiv (e_s)_{s=1}^t$ chosen by the agent in previous periods and on the agent’s current

⁴As it will become clear from the analysis in the subsequent sections, that Θ_t is bounded from below is to guarantee that expected payoffs, when expressed taking incentives into account, are well defined.

⁵That effort can take negative values should not raise concerns: here, e simply stands for the effect of the agent’s activity on the firm’s performance, so there is no reason to restrict it to be positive.

“type” θ_t .⁶ In particular, we assume that⁷

$$\pi_t = \theta_t + e_t + \sum_{\tau=1}^{t-1} \rho^\tau e_{t-\tau}$$

for some constant $\rho \geq 0$ that captures the persistence of the effect of the manager’s effort on the firm’s profits. The set of possible period- t profits will be denoted by

$$\Pi_t \equiv \left\{ \pi_t \in \mathbb{R} : \pi_t = \theta_t + e_t + \sum_{\tau=1}^{t-1} \rho^\tau e_{t-\tau}, \theta_t \in \Theta_t, e_s \in E, \forall s \leq t \right\}$$

Both θ^t and e^t are the agent’s private information. On the contrary, the stream of profits π^t are assumed to be verifiable, which implies that the agent can be rewarded as a function of the firm’s profits.

As is common in the literature, we equate the agent’s period- t consumption c_t with the payment from the principal (in other words, we assume away the possibility of hidden savings). Such a restriction is, however, without loss of generality under the assumption of risk-neutrality considered in this section.

In each period, the principal may condition the agent’s payment on the entire history of profits π^t . By choosing effort e_t in period t , the agent suffers a disutility $\psi(e_t)$. To ensure interior solutions and to validate a certain dynamic envelope theorem (more below), we will assume that ψ is a continuously differentiable function and that there exist scalars $\bar{e} \in \mathbb{R}_{++}$ and $K \in \mathbb{R}_{++}$ such that $\psi(e) = 0$ for all $e < 0$, ψ is thrice continuously differentiable over $(0, \bar{e})$ with $\psi''(e) > 0$ and $\psi'''(e) \geq 0$ for all $e \in (0, \bar{e})$ and $\psi(e) = Ke$ for all $e > \bar{e}$.⁸

The agent’s preferences over (lotteries over) streams of consumption levels c^T and streams of effort choices e^T are described by an expected utility function with (Bernoulli) utility given by

$$U^A(c^T, e^T) = \sum_{t=1}^T \delta^{t-1} [c_t - \psi(e_t)] \tag{1}$$

⁶From now on, we adopt the convention of denoting sequences of variables by their superscripts.

⁷Note that because θ_t is *not* restricted to be independent of the past shocks $\theta^{t-1} \equiv (\theta_1, \dots, \theta_{t-1})$, there is no loss of generality in assuming that π_t depends only on θ_t , as opposed to the entire history $\theta^t = (\theta_1, \dots, \theta_t)$. To see this, suppose that $\pi_t = f_t(\theta^t) + h_t(e^t)$ for some functions $f_t : \mathbb{R}^t \rightarrow \mathbb{R}$ and $h_t : \mathbb{R}^t \rightarrow \mathbb{R}$. It then suffices to change variables and simply let $\theta_t^{new} = f_t(\theta^t)$.

⁸These conditions are satisfied e.g. when $\bar{e} = K$ and $\psi(e) = (1/2)e^2$ for all $e \in [0, \bar{e}]$. More generally, note that the assumption that $\psi''' \geq 0$ guarantees that the principal’s relaxed program, as defined below, is concave.

where $\delta < 1$ is a discount factor. As standard, the aforementioned specification presumes time consistency. In what follows, we will thus assume that, after each history h_t , the agent maximizes the expectation of $U^A(c^T, e^T)$, where the expectation is taken with respect to whatever information is available to the agent after history h_t .

The principal's payoff is given by the discounted sum of the firm's profits, net of the agent's compensation:

$$U^P(\pi^T, c^T) = \sum_{t=1}^T \delta^{t-1} [\pi_t - c_t].$$

The function U^P also corresponds to the principal's Bermoulli function used to evaluate lotteries over (π^T, c^T) .

Throughout, we will also assume that $\delta\rho < 1$ and that $K > 1 + \sum_{s=1}^{T-1} (\delta\rho)^s$.

In each period t , θ_t is drawn from a cumulative distribution function $F_t(\cdot|\theta^{t-1})$ defined on Θ_t .⁹ Below, we will often find it convenient to describe the evolution of the agent's type through a collection of functions of independent shocks. More precisely, let $(\tilde{\varepsilon}_t)_{t=2}^T$ denote a collection of random variables, each distributed according to the c.d.f. G_t , strictly increasing on the interval $\mathcal{E}_t \subset \mathbb{R}$, where $\mathcal{E}_t = [\underline{\varepsilon}_t, \bar{\varepsilon}_t]$ if $\bar{\varepsilon}_t < +\infty$ and $[\underline{\varepsilon}_t, \bar{\varepsilon}_t)$ if $\bar{\varepsilon}_t = +\infty$ for some $-\infty < \underline{\varepsilon}_t \leq \bar{\varepsilon}_t \leq +\infty$, and such that $(\tilde{\theta}_1, \tilde{\varepsilon}^T)$ are jointly independent.¹⁰ Then, let $(z_t(\cdot))_{t=2}^T$ denote a collection of real-valued functions such that, for any $t \geq 2$, any θ_1 and any $\varepsilon^{t-1} \in \mathcal{E}^{t-1} \equiv \times_{s=2}^{t-1} \mathcal{E}_s$, the distribution of $z_t(\theta_1, (\varepsilon^{t-1}, \tilde{\varepsilon}_t))$ given $(\theta_1, \varepsilon^{t-1})$ is the same as that of θ_t given $\theta^{t-1} = z^{t-1}(\theta_1, \varepsilon^{t-1})$, where $z^{t-1}(\theta_1, \varepsilon^{t-1}) \equiv (\theta_1, z_2(\theta_1, \varepsilon_2), \dots, z_{t-1}(\theta_1, \varepsilon^{t-1}))$. As indicated in Pavan, Segal, and Toikka, 2009(b), any stochastic process (i.e. any collection of kernels $F = \langle F_t(\cdot|\cdot) \rangle_{t=1}^T$) admits at least one such representation.

We initially restrict attention to processes for which each z_t is separable in its first component.

Definition 1 *The process for $(\tilde{\theta}_t)_{t=1}^T$ given by the kernels*

$$F \equiv \langle F_t : \Theta^{t-1} \rightarrow \Delta(\Theta_t) \rangle_{t=1}^T$$

⁹Throughout, we assume that, for any t , any (θ^t) , any $s > t$, $\mathbb{E}[|\tilde{\theta}_s| | \theta^t] < +\infty$.

¹⁰The reason for restricting $\underline{\varepsilon}_t > -\infty$ is the same as for restricting Θ_t to be bounded from below; it guarantees that the agent's payoff in any incentive compatible mechanism can be conveniently expressed in integral form.

is separable in the first component (SFC) if it admits an independent-shock representation such that for each $t \geq 2$, the function $z_t : \Theta_1 \times \mathcal{E}^t \rightarrow \Theta_t$ takes the form

$$z_t(\theta_1, \varepsilon^t) = \gamma_t(\theta_1) + \phi_t(\varepsilon_2, \dots, \varepsilon_t)$$

for some functions $\gamma_t : \Theta_1 \rightarrow \mathbb{R}$ and $\phi_t : \mathcal{E}^t \rightarrow \mathbb{R}$.

The set of SFC processes is quite large and it includes for example all moving average processes, and more generally any ARIMA process with arbitrary parameters.

2.2 The mechanism design problem

The principal's problem consists of choosing a mechanism detailing for each period t a recommendation for the agent's effort e_t and a level of consumption c_t that depend on the sequence of realized profits π^t and (possibly) on a sequence of messages about the environment sent by the agent over time.

By the revelation principle, we restrict attention to direct mechanisms for which a truthful and obedient strategy is optimal for the agent. Let $\Theta^t \equiv \times_{\tau=1}^t \Theta_\tau$ and $\Pi^t = \times_{\tau=1}^t \Pi_\tau$. A (deterministic) direct mechanism $\Omega = \langle \xi_t, s_t \rangle_{t=1}^T$ consists of a collection of functions $\xi_t : \Theta^t \times \Pi^{t-1} \rightarrow E$ and $s_t : \Theta^t \times \Pi^t \rightarrow \mathbb{R}$ such that $\xi_t(\theta^t, \pi^{t-1})$ is the recommended level of effort for period t given the agent's reports θ^t and the observed past profits π^{t-1} , while $s_t(\theta^t, \pi^{t-1}, \pi_t)$ is the principal's payment (i.e. the agent's consumption) at the end of period t given the reports θ^t and the observed profits $\pi^t = (\pi^{t-1}, \pi_t)$. Note that $s_t(\theta^t, \pi^{t-1}, \pi_t)$ depends also on the current performance π_t . Equivalently, the mechanism Ω specifies for each period t and each history (θ^t, π^{t-1}) a recommended effort level $\xi_t(\theta^t, \pi^{t-1})$ along with a contingent payment scheme $s_t(\theta^t, \pi^{t-1}, \cdot) : \Pi_t \rightarrow \mathbb{R}$. With a slight abuse of notation, henceforth we will denote by $e_t(\theta^t) \equiv \xi_t(\theta^t, \pi^{t-1}(\theta^{t-1}))$ and by $c_t(\theta^t) = s_t(\theta^t, \pi^t(\theta^t))$ respectively the equilibrium effort and the equilibrium consumption level for period t given θ^t , where $\pi^t(\theta^t) = (\pi_s(\theta^s))_{s=1}^t$ with $\pi_s(\theta^s)$ defined recursively by $\pi_s(\theta^s) = \theta_s + \sum_{\tau=0}^{s-1} \rho^\tau \xi_{s-\tau}(\theta^{s-\tau}, \pi^{s-\tau-1}(\theta^{s-\tau-1}))$.

The timing of play in each period t is the following:

- At the beginning of period t , the agent learns $\theta_t \in \Theta_t$;

- The agent then sends a report $\hat{\theta}_t \in \Theta_t$;
- Finally, the mechanism reacts by prescribing an effort choice $e_t = \xi_t(\theta^t, \pi^{t-1})$ and a reward scheme $s_t(\theta^t, \pi^{t-1}, \cdot) : \Pi_t \rightarrow \mathbb{R}$.

The mechanism Ω is offered to the agent at date 1, after he has observed the first realization θ_1 of the process governing the evolution of θ_t .¹¹ If the agent refuses to participate in the mechanism Ω , then both the agent and the principal receive their outside options, which we assume to be equal to zero. If, instead, the agent accepts Ω , then he is obliged to stay in the relationship in all subsequent periods.¹²

Because we will often find it convenient to describe the evolution of the agent's type through an independent-shock representation (described above), hereafter, we will also consider direct mechanisms in which the agent reports the shocks ε_t in each period $t \geq 2$ instead of his period- t type θ_t . We will then denote such mechanisms by $\hat{\Omega} = \langle \hat{\xi}_t, \hat{s}_t \rangle_{t=1}^T$ where $\hat{\xi}_t : \Theta_1 \times \mathcal{E}^t \times \Pi^{t-1} \rightarrow E$ and $\hat{s}_t : \Theta_1 \times \mathcal{E}^t \times \Pi^t \rightarrow \mathbb{R}$ have the same interpretation as the mappings ξ_t and s_t in the primitive representation (the one in terms of the θ^t). Likewise, we will denote by $\hat{c}_t(\theta_1, \varepsilon^t)$ and by $\hat{e}_t(\theta_1, \varepsilon^t)$ the consumption and effort choices that are implemented in equilibrium given $(\theta_1, \varepsilon^t)$.

The optimal mechanism

To ease the understanding of the properties of the optimal mechanism, we start by considering the optimal effort policy in the absence of any private information.¹³

Proposition 1 *Assume the agent does not possess any private information, i.e. both the evolution of the environment (as captured by the process for θ_t) and the agent's effort choices e^T are publicly observable and verifiable. The optimal contract for the principal then implements the following*

¹¹Allowing the agent to possess private information at the time of contracting is not only realistic, but essential to shed light on important aspects of the optimal contract such as the time-varying power of incentives. Furthermore, it permits one to derive interesting dynamics, even without assuming the agent is risk-averse.

¹²That participation must be guaranteed only in period one is clearly not restrictive when the principal can ask the agent to post bonds. Below, we will discuss also situations/implementations where, even in the absence of bonding, participation can be guaranteed after any history.

¹³Given the assumptions on ψ , $e_t^{FB} \in (0, \bar{e})$ for all t .

effort policy:

$$\psi'(e_t^{FB}) = 1 + \sum_{s=1}^{T-t} (\delta\rho)^s \quad \forall t, \forall(\theta_1, \varepsilon^t)$$

In particular, when $T = +\infty$, the optimal effort is stationary over time and is implicitly given by $\psi'(e^{FB}) = 1/[1 - \delta\rho]$. Likewise, the optimal effort is constant and given by $\psi'(e^{FB}) = 1$ when the manager's effort has only a transitory effect on the firm's performance, i.e. when $\rho = 0$. That the first-best effort policy is independent of any variation in the underlying environment is a consequence of the assumption of independence of the agent's disutility of effort from the underlying state θ_t .

Clearly, the same first-best effort policy is implemented in any environment in which the agent's initial type θ_1 is publicly observed and verifiable (equivalently, in any environment in which the agent contracts with the principal before learning θ_1), irrespective of the observability of effort choices and future shocks θ_t .¹⁴

Next, consider the case where the agent possesses relevant private information. Thus assume that both the evolution of the environment (as captured by the process for θ_t) and the agent's effort choices are the agent's private information. In addition, suppose that contracting between the agent and the principal occurs at a time at which the agent is already informed about his period-1 type θ_1 . The following proposition presents the main characterization result for this environment.

Proposition 2 *Assume the process governing the evolution of θ_t satisfies the SFC condition and that, for each t , $\gamma_t(\cdot)$ is differentiable and there exists $M \in \mathbb{R}_+$ such that $\sup_t \{|\gamma'_t(\theta_1)|\} \leq M$ for all θ_1 . For any θ_1 , let $D_{1,1}(\theta_1) \equiv 1$ and for any $t \geq 2$, let $D_{1,t}(\theta_1) \equiv \gamma'_t(\theta_1) - \rho\gamma'_{t-1}(\theta_1)$, with, $\gamma'_1(\theta_1) \equiv 1$, and suppose that for any t , any θ_1 , $D_{1,t}(\theta_1) \geq 0$. Finally, assume that F_1 is absolutely continuous with density $f_1(\theta_1) > 0$ for all $\theta_1 \in \Theta_1$ and denote by $\eta(\theta_1) \equiv [1 - F_1(\theta_1)]/f_1(\theta_1)$ its inverse hazard rate. Then consider the effort policy \hat{e}^* implicitly defined, for all t all θ_1 , by¹⁵*

$$\psi'(\hat{e}_t^*(\theta_1)) = 1 + \sum_{s=1}^{T-t} (\delta\rho)^s - \eta(\theta_1)D_{1,t}(\theta_1)\psi''(\hat{e}_t^*(\theta_1)) \quad \forall\theta_1, \forall t \geq 1, \quad (2)$$

¹⁴As we will show below, this property is however a consequence of the assumption of transferable utility, i.e. of the fact that both the agent's and the principal's preferences are linear in the transfers c_t .

¹⁵Throughout, ψ''_+ will denote the second right derivative of ψ .

unless $\psi''_+(0) \geq \left[1 + \sum_{s=1}^{T-t} (\delta\rho)^s\right] / [\eta(\theta_1)D_{1,t}(\theta_1)]$ in which case $\hat{e}_t^*(\theta_1) = 0$.

1. For any t and any θ_1 let

$$\alpha_t(\theta_1) \equiv \psi'(\hat{e}_t^*(\theta_1)) - \delta\rho\psi'(\hat{e}_{t+1}^*(\theta_1))$$

[if T is finite, then $\alpha_T(\theta_1) \equiv \psi'(\hat{e}_T^*(\theta_1))$]. Suppose the policy \hat{e}^* defined above satisfies the following single-crossing condition

$$\left[\sum_{t=1}^T \delta^{t-1} \gamma'_t(\theta_1) [\alpha_t(\theta_1) - \alpha_t(\hat{\theta}_1)] \right] [\theta_1 - \hat{\theta}_1] \geq 0 \quad \forall \theta_1, \hat{\theta}_1 \in \Theta_1. \quad (3)$$

Then the recommendation policy

$$\hat{\xi}_t^*(\theta_1, \varepsilon^t, \pi^{t-1}) = \hat{e}_t^*(\theta_1) \quad \forall (\theta_1, \varepsilon^t, \pi^{t-1}) \in \Theta_1 \times \mathcal{E}^t \times \Pi^{t-1}$$

together with the output-contingent reward scheme defined below are part of an optimal mechanism. The reward scheme is such that

$$\hat{s}_1^*(\theta_1, \pi_1) = S_1(\theta_1) + \alpha_1(\theta_1)\pi_1$$

while for any $t \geq 2$,

$$\hat{s}_t^*(\theta_1, \varepsilon^t, \pi^t) = \alpha_t(\theta_1)\pi_t$$

where

$$S_1(\theta_1) \equiv \sum_{t=1}^T \delta^{t-1} \left[\psi(\hat{e}_t^*(\theta_1)) + \int_{\underline{\theta}_1}^{\theta_1} D_{1,t}(s) \psi'(\hat{e}_t^*(s)) ds - \mathbb{E} [\alpha_t(\theta_1) \hat{\pi}_t^*(\theta_1, \varepsilon^t)] \right]$$

with $\hat{\pi}_t^*(\theta_1, \varepsilon^t) \equiv z_t(\theta_1, \varepsilon^t) + \hat{e}_t^*(\theta_1) + \sum_{\tau=1}^{t-1} \rho^\tau \hat{e}_{t-\tau}^*(\theta_1)$.

2. Suppose that for any t , either (a) $\rho = 0$ and the function $\eta(\cdot)D_{1,t}(\cdot)$ is non-increasing, or (b) $\psi(e) = ke^2/2$ for all $e \in [0, \bar{e}]$ and $\eta(\cdot)[D_{1,t}(\cdot) - \delta\rho D_{1,t+1}(\cdot)]$ is non-increasing [if T is finite, then for $t = T$, $\eta(\cdot)D_{1,T}(\cdot)$ is non-increasing]. Then the effort policy \hat{e}^* of part (1) satisfies the single-crossing condition (3).

Because this is one of the main results in the paper and because many of the subsequent results follow from arguments/techniques similar to those used to establish Proposition 2, the proof for this result is given below instead of being relegated to the Appendix. The reader interested only in the predictions of the model can however skip this proof and continue with the reading at page 21.

Proof. The structure of the proof is the following. Lemma 1 provides a necessary condition for incentive compatibility based on the application of a dynamic envelope theorem (as in Pavan, Segal, and Toikka, 2009(b)) to the agent’s optimization problem. Lemma 2 characterizes the effort policy \hat{e}^* that solves the principal’s *relaxed* problem, where the latter considers only the necessary condition established in Lemma 1 (along with a certain participation constraint) and ignores all remaining constraints. Lemma 3 shows that, when the solution to the relaxed program satisfies the single-crossing condition of (3), then (i) it can be implemented by the linear scheme described in the proposition, (ii) under this scheme all types find it optimal to participate, and (iii) the lowest type θ_1 receives a zero expected payoff in equilibrium. As discussed in more detail below, together these properties guarantee that the effort policy \hat{e}^* (equivalently, the recommendation policy $\hat{\xi}^*$) along with the linear reward scheme \hat{s}^* are part of an optimal mechanism. Finally, Lemma 4 completes the proof by establishing the result in Part 2.

Given the mechanism $\hat{\Omega} = \langle \hat{\xi}, \hat{s} \rangle$, let $V^{\hat{\Omega}}(\theta_1)$ denote the value function when the agent’s period one type is θ_1 . This is simply the supremum of the agent’s expected payoff over all possible reporting and effort strategies. The mechanism $\hat{\Omega}$ is incentive compatible if $V^{\hat{\Omega}}(\theta_1)$ coincides with the agent’s expected payoff under a truthful and obedient strategy for every $\theta_1 \in \Theta_1$. We then have the following result.

Lemma 1 *The mechanism $\hat{\Omega}$ is incentive compatible only if $V^{\hat{\Omega}}(\theta_1)$ is Lipschitz continuous and, for almost every $\theta_1 \in \Theta_1$,*

$$\frac{dV^{\hat{\Omega}}(\theta_1)}{d\theta_1} = \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} D_{1,t}(\theta_1) \psi'(\hat{e}_t(\theta_1, \hat{\varepsilon}^t)) \right].$$

Proof of the lemma. Consider the following fictitious environment. At any point in time, the agent can misreport his private information but is then “forced” to choose effort so as to perfectly

“hide” his lies. That is, at any period t , and for any given sequence of reports $(\hat{\theta}_1, \hat{\varepsilon}^t)$, the agent must exert effort e_t so that $\pi_t = \hat{\pi}_t(\hat{\theta}_1, \hat{\varepsilon}^t)$, where $\hat{\pi}_t(\hat{\theta}_1, \hat{\varepsilon}^t)$ is the equilibrium profit for period t given $(\hat{\theta}_1, \hat{\varepsilon}^t)$, as defined in the Proposition. Now let

$$\begin{aligned}
\hat{e}_t(\theta_1, \varepsilon^t; \hat{\theta}_1, \hat{\varepsilon}^t) &= \hat{\pi}_t(\hat{\theta}_1, \hat{\varepsilon}^t) - z_t(\theta_1, \varepsilon^t) - \rho \sum_{\tau=1}^{t-1} \rho^{\tau-1} \hat{e}_{t-\tau}(\theta_1, \varepsilon^{t-\tau}; \hat{\theta}_1, \hat{\varepsilon}^{t-\tau}) \\
&= \hat{\pi}_t(\hat{\theta}_1, \hat{\varepsilon}^t) - z_t(\theta_1, \varepsilon^t) \\
&\quad - \rho \left(\hat{e}_{t-1}(\theta_1, \varepsilon^{t-1}; \hat{\theta}_1, \hat{\varepsilon}^{t-1}) + \sum_{\tau=1}^{t-2} \rho^\tau \hat{e}_{t-1-\tau}(\theta_1, \varepsilon^{t-1-\tau}; \hat{\theta}_1, \hat{\varepsilon}^{t-1-\tau}) \right) \\
&= \hat{\pi}_t(\hat{\theta}_1, \hat{\varepsilon}^t) - z_t(\theta_1, \varepsilon^t) - \rho \left(\hat{\pi}_{t-1}(\hat{\theta}_1, \hat{\varepsilon}^{t-1}) - z_{t-1}(\theta_1, \varepsilon^{t-1}) \right)
\end{aligned} \tag{4}$$

denote the effort the agent must exert in period t to meet the target $\hat{\pi}_t(\hat{\theta}_1, \hat{\varepsilon}^t)$ when his true type is $(\theta_1, \varepsilon^t)$ given that he met the targets $(\hat{\pi}_s(\hat{\theta}_1, \hat{\varepsilon}^s))_{s=1}^{t-1}$ in all preceding periods, with $e_1(\theta_1; \hat{\theta}_1) = \hat{\pi}_1(\hat{\theta}_1) - \theta_1$.

Now fix $(\hat{\theta}_1, \hat{\varepsilon}^T)$ and let $(\hat{c}^T, \hat{\pi}^T)$ be the stream of equilibrium payments and profits that, given the mechanism $\hat{\Omega}$, correspond to the sequence of reports $(\hat{\theta}_1, \hat{\varepsilon}^T)$. For any $(\hat{\theta}_1, \hat{\varepsilon}^T)$ and given any sequence of true shocks $(\theta_1, \varepsilon^T)$, the agent’s payoff in this fictitious environment is given by

$$\begin{aligned}
\hat{U}^A(\theta_1, \varepsilon^T; \hat{\theta}_1, \hat{\varepsilon}^T) &= \sum_{t=1}^T \delta^{t-1} [\hat{c}_t - \psi(\hat{e}_t(\theta_1, \varepsilon^t; \hat{\theta}_1, \hat{\varepsilon}^t))] \\
&= \hat{c}_1 - \psi(\hat{\pi}_1 - \theta_1) \\
&\quad + \sum_{t=2}^T \delta^{t-1} [\hat{c}_t - \psi(\hat{\pi}_t - z_t(\theta_1, \varepsilon^t) - \rho(\hat{\pi}_{t-1} - z_{t-1}(\theta_1, \varepsilon^{t-1})))] \\
&= \hat{c}_1 - \psi(\hat{\pi}_1 - \theta_1) \\
&\quad + \sum_{t=2}^T \delta^{t-1} [\hat{c}_t - \psi(\hat{\pi}_t - \gamma_t(\theta_1) - \phi_t(\varepsilon^t) - \rho(\hat{\pi}_{t-1} - \gamma_{t-1}(\theta_1) - \phi_{t-1}(\varepsilon^{t-1})))]
\end{aligned}$$

The assumptions that ψ and γ_t are differentiable and equi-Lipschitz continuous imply that \hat{U}^A is equi-Lipschitz continuous and differentiable in θ_1 . Now suppose the mechanism $\hat{\Omega}$ is incentive compatible in the unrestricted setting where the agent is free to choose any effort he wants at any point in time. It is then necessarily incentive compatible also in this fictitious setting where effort is pinned down by $(\theta_1, \varepsilon^T; \hat{\theta}_1, \hat{\varepsilon}^T)$ according to (4). The result in the Lemma then follows from Proposition 1 in Pavan, Segal, and Toikka (2009,b): Letting $\hat{U}^A(\theta_1, \varepsilon^T)$ denote the agent’s payoff when he follows a truthtelling and obedient strategy, we have that $\hat{\Omega}$ is incentive compatible only

if $V^{\hat{\Omega}}$ is Lipschitz continuous and, for almost every $\theta_1 \in \Theta_1$,

$$\begin{aligned} \frac{dV^{\hat{\Omega}}(\theta_1)}{d\theta_1} &= \mathbb{E} \left[\frac{\partial \hat{U}^A(\theta_1, \tilde{\varepsilon}^T)}{\partial \theta_1} \right] \\ &= \mathbb{E} \left[\psi'(\hat{e}_1(\theta_1)) + \sum_{t=2}^T \delta^{t-1} [\gamma'_t(\theta_1) - \rho \gamma'_{t-1}(\theta_1)] \psi'(\hat{e}_t(\theta_1, \tilde{\varepsilon}^t)) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} D_{1,t}(\theta_1) \psi'(\hat{e}_t(\theta_1, \tilde{\varepsilon}_t)) \right], \end{aligned}$$

which establishes the result. ■

Now, one can think of the principal's problem as involving the choice of a pair of contingent policies $\langle \hat{\pi}, \hat{c} \rangle$ so as to maximize her expected payoff

$$\mathbb{E}[\hat{U}^P] = \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left[\hat{\pi}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) - \hat{c}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) \right] \right]$$

subject to all IC and IR constraints. Because both the principal's and the agent's preferences are quasilinear, $\mathbb{E}[\hat{U}^P]$ can be rewritten as expected total surplus, net of the agent's expected (intertemporal) rent:

$$\mathbb{E}[\hat{U}^P] = \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left[\hat{\pi}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) - \psi(\hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t)) \right] \right] - \mathbb{E}[V^{\hat{\Omega}}(\tilde{\theta}_1)]. \quad (5)$$

Using the result in the previous Lemma, and integrating by parts, the agent's expected (intertemporal) rent can in turn be written as

$$\begin{aligned} \mathbb{E}[V^{\hat{\Omega}}(\tilde{\theta}_1)] &= V^{\hat{\Omega}}(\underline{\theta}_1) + \mathbb{E} \left[\frac{1 - F(\tilde{\theta}_1)}{f(\tilde{\theta}_1)} \frac{dV^{\hat{\Omega}}(\tilde{\theta}_1)}{d\theta_1} \right] \\ &= V^{\hat{\Omega}}(\underline{\theta}_1) + \mathbb{E} \left[\eta(\tilde{\theta}_1) \sum_{t=1}^T \delta^{t-1} D_{1,t}(\tilde{\theta}_1) \psi'(\hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t)) \right]. \end{aligned} \quad (6)$$

Finally, substituting (6) into (5), we have that

$$\begin{aligned} \mathbb{E}[\hat{U}^P] &= \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left(\hat{\pi}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) - \psi(\hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t)) - \eta(\tilde{\theta}_1) D_{1,t}(\tilde{\theta}_1) \psi'(\hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t)) \right) \right] - V^{\hat{\Omega}}(\underline{\theta}) \quad (7) \\ &= \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left(z_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) + \hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) + \sum_{\tau=1}^{t-1} \rho^\tau \hat{e}_{t-\tau}(\tilde{\theta}_1, \tilde{\varepsilon}^{t-\tau}) - \psi(\hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t)) \right) \right] \\ &\quad - V^{\hat{\Omega}}(\underline{\theta}) \end{aligned}$$

Next, consider a relaxed program for the principal that consists of choosing an effort policy \hat{e} and a constant $V^{\hat{\Omega}}(\underline{\theta}) \geq 0$ so as to maximize $\mathbb{E}[\hat{U}^P]$. The solution to this relaxed program is given in the following lemma.

Lemma 2 *Suppose that $D_{1,t}(\theta_1) \geq 0$ for any $\theta_1 \in \Theta_1$ and any t . The (almost-unique) solution to the principal's relaxed program is then given by $V^{\hat{\Omega}}(\underline{\theta}) = 0$ along with the effort policy \hat{e}^* in the Proposition.*

Proof of the Lemma. The result follows directly from pointwise maximization of (7). The assumptions that ψ is a continuously differentiable function with $\psi(e) = 0$ for all $e < 0$, $\psi''(e) > 0$ and $\psi'''(e) \geq 0$ for all $e \in [0, \bar{e}]$, $\psi'(e) = K$ for all $e > \bar{e}$, together with $D_{1,t}(\theta_1) \geq 0$ for all θ_1 , imply that, for all t all $(\theta_1, \varepsilon^t)$, the principal's payoff \hat{U}^P is strictly increasing in e_t for all $e_t < \hat{e}_t^*(\theta_1)$, and strictly decreasing in e_t for all $e_t > \hat{e}_t^*(\theta_1)$, where $\hat{e}_t^*(\theta_1)$ is implicitly given by (2) when $\psi''_+(0) < \left[1 + \sum_{s=1}^{T-t} (\delta\rho)^s\right] / [\eta(\theta_1)D_{1,t}(\theta_1)]$ and by $\hat{e}_t^*(\theta_1) = 0$ otherwise. ■

To prove the result in part 1, it then suffices to show that, when the effort policy in (2) satisfies the single-crossing condition (3), it can be implemented by the linear scheme proposed in the Proposition. That is, it suffices to show that, under this scheme, (i) the agent finds it optimal to participate in period one, (ii) the agent finds it optimal to report all his private information truthfully and to obey the principal's recommendations, and (iii) the lowest period-1 type's expected payoff is equal to his outside option, i.e. $V^{\hat{\Omega}}(\underline{\theta}) = 0$. This is shown in the following lemma.

Lemma 3 *Assume the effort policy \hat{e}^* that solves the relaxed program (as implicitly given by (2)) satisfies the single-crossing condition (3). Then the mechanism $\hat{\Omega} = \langle \hat{\xi}_t, \hat{s}_t \rangle_{t=1}^T$, where $\hat{\xi}_t$ and \hat{s}_t are, respectively, the recommendation policy and the reward scheme described in the Proposition, implements the effort policy \hat{e}^* . It induces any type θ_1 to participate and gives the lowest period-1 type $\underline{\theta}_1$ an expected payoff of zero.*

Proof of the Lemma. Because neither $\hat{\xi}_t$ nor \hat{s}_t depend on ε^t , it is immediate that the agent finds it optimal to report the shocks truthfully. Furthermore, conditional upon reporting $\hat{\theta}_1$ in period 1, it is also immediate that, at any period $t \geq 1$, the agent finds it optimal to follow the

principal's recommendation and choose effort $\hat{e}_t^*(\hat{\theta}_1)$, irrespective of his true period-1 type θ_1 , the true shocks ε^t and the history of past performances π^{t-1} . To see this, note that at any period $t \geq 1$, and for any history $(\theta_1, \varepsilon^t, \hat{\theta}_1, \hat{\varepsilon}^t, \pi^{t-1}, e^{t-1})$, the problem that the agent faces in period t is to choose a (possibly contingent) plan $(e_t, e_{t+1}(\cdot), \dots, e_T(\cdot))$ to maximize

$$\mathbb{E} \left[\sum_{\tau=t}^T \delta^{\tau-t} \left(\alpha_\tau(\hat{\theta}_1) \left(\tilde{e}_\tau + \sum_{\kappa=1}^{\tau-1} \rho^\kappa \tilde{e}_{\tau-\kappa} + z_\tau(\theta_1, \tilde{\varepsilon}^\tau) \right) - \psi(\tilde{e}_\tau) \right) \mid \theta_1, \varepsilon^t \right]$$

The solution to this problem is given by the (non-contingent) effort policy implicitly defined by

$$\psi'(e_\tau) = \alpha_\tau(\hat{\theta}_1) + \sum_{s=1}^T (\delta\rho)^s \alpha_{\tau+s}(\hat{\theta}_1) \text{ all } \tau \geq t.$$

When the sequence $(\alpha_t(\hat{\theta}_1))_{t=1}^T$ is the one specified in the Proposition, the effort policy that solves these conditions is the policy \hat{e}^* that solves the relaxed program.

It remains to show that each type θ_1 finds it optimal to report truthfully and to participate, and that type $\underline{\theta}_1$ expects a zero payoff from the relationship. That each type θ_1 finds it optimal to participate is guaranteed by the fact that his expected payoff (under a truthful and obedient strategy) is given by

$$\sum_{t=1}^T \delta^{t-1} \int_{\underline{\theta}_1}^{\theta_1} D_{1,t}(s) \psi'(e_t^*(s)) ds,$$

which is non-negative because $D_{1,t}(\theta_1) \geq 0$ and $\psi'(e) \geq 0$. To see that each type θ_1 finds it optimal to report truthfully let

$$\begin{aligned} U(\theta_1; \hat{\theta}_1) &\equiv \sum_{t=1}^T \delta^{t-1} \int_{\underline{\theta}_1}^{\hat{\theta}_1} D_{1,t}(s) \psi'(e_t^*(s)) ds \\ &\quad + \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \alpha_t(\hat{\theta}_1) [z_t(\theta_1, \tilde{\varepsilon}^t) - z_t(\hat{\theta}_1, \tilde{\varepsilon}^t)] \right]. \end{aligned}$$

The function $U(\theta_1; \hat{\theta}_1)$ simply represents the payoff that type θ_1 obtains by mimicking type $\hat{\theta}_1$.

Next note that

$$\frac{\partial U(\theta_1; \hat{\theta}_1)}{\partial \theta_1} = \sum_{t=1}^T \delta^{t-1} \alpha_t(\hat{\theta}_1) \gamma'_t(\theta_1).$$

The single-crossing condition in the Proposition guarantees that

$$\left[\frac{dU(\theta_1; \theta_1)}{d\theta_1} - \frac{\partial U(\theta_1; \hat{\theta}_1)}{\partial \theta_1} \right] [\theta_1 - \hat{\theta}_1] \geq 0.$$

To see this note that

$$\begin{aligned}
\frac{dU(\theta_1; \theta_1)}{d\theta_1} - \frac{\partial U(\theta_1; \hat{\theta}_1)}{\partial \theta_1} &= \sum_{t=1}^T \delta^{t-1} D_{1,t}(\theta_1) \psi'(\hat{e}_t^*(\theta_1)) - \sum_{t=1}^T \delta^{t-1} \alpha_t(\hat{\theta}_1) \gamma'_t(\theta_1) \\
&= \sum_{t=1}^T \delta^{t-1} \left\{ \begin{array}{l} [\gamma'_t(\theta_1) - \rho \gamma'_{t-1}(\theta_1)] \psi'(\hat{e}_t^*(\theta_1)) \\ -\gamma'_t(\theta_1) [\psi'(\hat{e}_t^*(\hat{\theta}_1)) - \delta \rho \psi'(\hat{e}_{t+1}^*(\hat{\theta}_1))] \end{array} \right\} \\
&= \psi'(\hat{e}_1^*(\theta_1)) - \psi'(\hat{e}_1^*(\hat{\theta}_1)) + \delta \rho \psi'(\hat{e}_2^*(\hat{\theta}_1)) \\
&\quad + \delta \gamma'_2(\theta_1) \psi'(\hat{e}_2^*(\theta_1)) - \delta \rho \psi'(\hat{e}_2^*(\theta_1)) \\
&\quad + \dots \\
&= [\psi'(\hat{e}_1^*(\theta_1)) - \delta \rho \psi'(\hat{e}_2^*(\theta_1))] - [\psi'(\hat{e}_1^*(\hat{\theta}_1)) - \delta \rho \psi'(\hat{e}_1^*(\hat{\theta}_1))] \\
&\quad + \delta \gamma'_2(\theta_1) [\psi'(\hat{e}_2^*(\theta_1)) - \delta \rho \psi'(\hat{e}_3^*(\theta_1))] \\
&\quad - \delta \gamma'_2(\theta_1) [\psi'(\hat{e}_2^*(\hat{\theta}_1)) - \delta \rho \psi'(\hat{e}_3^*(\hat{\theta}_1))] \\
&\quad + \dots \\
&= \sum_{t=1}^T \delta^{t-1} \gamma'_t(\theta_1) [\alpha_t(\theta_1) - \alpha_t(\hat{\theta}_1)].
\end{aligned}$$

The result then follows from Lemma 12 in Pavan, Segal, and Toikka (2009,a).

Lemma 4 *Suppose that, for any t , either (a) $\rho = 0$ and the function $\eta(\cdot)D_{1,t}(\cdot)$ is non-increasing, or (b) $\psi(e) = ke^2/2$ for all $e \in [0, \bar{e}]$ and $\eta(\cdot)[D_{1,t}(\cdot) - \delta \rho D_{1,t+1}(\cdot)]$ is non-increasing [if T is finite, then $\eta(\cdot)D_{1,T}(\cdot)$ is non-increasing]. Then the effort policy \hat{e}^* implicitly given by (2) satisfies condition (3), i.e., for any $\theta_1, \hat{\theta}_1 \in \Theta_1$:*

$$\left[\sum_{t=1}^T \delta^{t-1} \gamma'_t(\theta_1) [\alpha_t(\theta_1) - \alpha_t(\hat{\theta}_1)] \right] [\theta_1 - \hat{\theta}_1] \geq 0$$

Proof of the lemma. We establish the result by showing that, under the assumptions in the lemma, $\alpha_t(\theta_1)$ is non-decreasing in θ_1 , for each $t \geq 1$. Consider first case (a). When $\rho = 0$, $\alpha_t(\theta_1) = \psi'(\hat{e}_t^*(\theta_1))$. It then suffices to show that the effort policy $\hat{e}_t^*(\theta_1)$ implicitly given by (2) is non-decreasing. To see that this is indeed the case, it is enough to recognize that the dynamic virtual surplus (as defined in 7) has increasing differences in e_t and $-\eta(\theta_1)D_{1,t}(\theta_1)$ and, by assumption, $\eta(\cdot)D_{1,t}(\cdot)$ is non-increasing.¹⁶

¹⁶The relevant terms of the dynamic virtual surplus are $e_t + \sum_{s=1}^{T-t} (\delta \rho)^s e_t - \psi(e_t) - \eta(\theta_1)D_t(\theta_1)\psi'(e_t)$. The result then follows from monotone comparative statics analysis (see Topkis (1998) and Milgrom and Shannon (1994)).

Next, consider case (b). For any $t < T$ and any $\theta'_1 > \theta''_1$,

$$\begin{aligned}
\alpha_t(\theta'_1) - \alpha_t(\theta''_1) &= [\psi'(\hat{e}_t^*(\theta'_1)) - \delta\rho\psi'(\hat{e}_{t+1}^*(\theta'_1))] - [\psi'(\hat{e}_t^*(\theta''_1)) - \delta\rho\psi'(\hat{e}_{t+1}^*(\theta''_1))] \\
&= \left[1 + \sum_{s=1}^{T-t} (\delta\rho)^s - \eta(\theta'_1)D_{1,t}(\theta'_1)k - \delta\rho \left(1 + \sum_{s=1}^{T-t-1} (\delta\rho)^s - \eta(\theta'_1)D_{1,t+1}(\theta'_1)k \right) \right] \\
&\quad - \left[1 + \sum_{s=1}^{T-t} (\delta\rho)^s - \eta(\theta''_1)D_{1,t}(\theta''_1)k - \delta\rho \left(1 + \sum_{s=1}^{T-t-1} (\delta\rho)^s - \eta(\theta''_1)D_{1,t+1}(\theta''_1)k \right) \right] \\
&= k [\eta(\theta''_1) (D_{1,t}(\theta'_1) - \delta\rho D_{1,t+1}(\theta''_1)) - \eta(\theta'_1) (D_{1,t}(\theta'_1) - \delta\rho D_{1,t+1}(\theta'_1))] \\
&\geq 0,
\end{aligned}$$

where the inequality follows from the assumption that $\eta(\cdot)[D_{1,t}(\cdot) - \delta\rho D_{1,t+1}(\cdot)]$ is non-increasing.

Likewise, when T is finite, then

$$\alpha_T(\theta'_1) - \alpha_T(\theta''_1) = \psi'(\hat{e}_T^*(\theta'_1)) - \psi'(\hat{e}_T^*(\theta''_1)) = k [\eta(\theta''_1)D_{1,T}(\theta''_1) - \eta(\theta'_1)D_{1,T}(\theta'_1)] \geq 0,$$

where the inequality follows from the assumption that $\eta(\cdot)D_{1,T}(\cdot)$ is non-increasing. ■

This completes the proof of the proposition. ■

Note that, because the agent is indifferent over the way the constant term $S(\theta_1)$ is distributed over time, an equivalent (linear) implementation consists in paying the agent in each period t a fixed wage

$$\psi(\hat{e}_t^*(\theta_1)) + \int_{\underline{\theta}_1}^{\theta_1} D_{1,t}(s)\psi'(\hat{e}_t^*(s))ds - \mathbb{E} [\alpha_t(\theta_1)\hat{\pi}_t^*(\theta_1, \tilde{\varepsilon}^t)]$$

plus a fraction $\alpha_t(\theta_1)$ of the current profits π_t , with $S_1(\theta_1)$ now defined by

$$S_1(\theta_1) = \psi(\hat{e}_1^*(\theta_1)) + \int_{\underline{\theta}_1}^{\theta_1} \psi'(\hat{e}_1^*(s))ds - \alpha_1(\theta_1)\hat{\pi}_1^*(\theta_1).$$

While the particular way the constant term $S_1(\theta_1)$ is distributed over time is clearly inconsequential for incentives, certain choices may have the advantage of guaranteeing that, if the agent has the option to leave the relationship at any point in time, he does not find it optimal to do so. To see this, suppose that $T = +\infty$ and that all shocks are strictly positive, i.e. $\theta_1, \varepsilon_s > 0$ for all s . Then front-loading the payment

$$- \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E} [\alpha_t(\theta_1)\hat{\pi}_t^*(\theta_1, \tilde{\varepsilon}^t)]$$

and then paying in each period

$$\psi(\hat{e}_t^*(\theta_1)) + \int_{\underline{\theta}_1}^{\theta_1} D_{1,t}(s)\psi'(\hat{e}_t^*(s))ds + \alpha_t(\theta_1)\pi_t$$

guarantees participation in each period, at any truthful history.

We now turn to the properties of the optimal effort policy.¹⁷ Because $D_{1,t} \geq 0$ and ψ' is convex, the optimal effort policy involves downward distortions. These distortions in turn depend on inverse hazard rate $\eta(\theta_1)$ of the first-period distribution F_1 and on the function $D_{1,t}$, which captures the effect of θ_1 on both θ_t and θ_{t-1} , taking into account the persistent effect of effort. When the process for θ_t satisfies condition SFC, these distortions are independent of the realizations of the shocks ε_t and of their distributions G_t . Whether $\hat{e}_t^*(\theta_1)$ increases or decreases with t then depends entirely on the dynamics of $D_{1,t}(\theta_1)$ as illustrated in the following examples, where the conditions of Part 2 of Proposition 2 are clearly satisfied.

Example 1 Suppose that $T = \infty$ and that θ_t evolves according to an AR(1) process

$$\theta_t = \beta\theta_{t-1} + \varepsilon_t$$

for some $\beta \in (0, 1)$ with $\beta > \rho \geq 0$. Then $D_{1,t}(\theta_1) = \beta^{t-2}(\beta - \rho)$ for all $\theta_1 \in \Theta_1$. It follows that $\hat{e}_t^*(\theta_1)$ increases over time and

$$\lim_{t \rightarrow \infty} \hat{e}_t^*(\theta_1) = 1/[1 - \delta\rho] = e^{FB} \quad \forall \theta_1.$$

Example 2 Assume that each θ_t is i.i.d., so that $D_{1,t}(\theta_1) = 0$ for all $t \geq 2$ and all θ_1 . Then effort is distorted only in the first period, i.e. $\hat{e}_1^*(\theta_1) < e_1^{FB}$ and $\hat{e}_t^* = e_t^{FB}$ for all $t \geq 2$.

Example 3 Suppose θ_t follows a random walk, i.e.

$$\theta_t = \theta_{t-1} + \varepsilon_t$$

and that effort has only a contemporaneous effect on the firm's profits (i.e. $\rho = 0$). Then $\hat{e}_t^*(\theta_1)$ is constant over time and coincides with the static optimal effort.

¹⁷Conditions similar to (2) have been derived in a two-period model by Baron and Besanko (1984) and Laffont and Tirole (1991). However these early work do not examine under what conditions (and under what contracts), the effort policies that solve the principal's relaxed program are implementable.

The result in Example 1 is actually quite general; many ARIMA(k,q,m) processes have the property that $\lim_{t \rightarrow \infty} D_{1,t} = 0$, where $D_{1,t}$ are nonnegative scalars decreasing in t that depend on the parameters (k,q,m) of the ARIMA process.

Example 2 is the case considered by Edmans and Gabaix (2008) in their baseline model, where it is also assumed that $\rho = 0$. However, contrary to the case considered here, they assume that contracting occurs before the agent learns his first-period type. As discussed above, together with risk neutrality this implies that the sequence of effort decisions is always efficient.

Finally, the random walk case of Example 3 is also a process that is sometimes considered in the literature. In this case, because effort is constant over time, the optimal mechanism can be implemented by offering in period one the same menu of linear contracts that the principal would offer in a static relationship, and then committing to using the contract selected in period one in each subsequent period. Each linear contract (indexed by θ_1) has a fixed payment of

$$S(\theta_1) \equiv \psi(\hat{e}^*(\theta_1)) + \int_{\theta_1}^{\theta_1} \psi'(\hat{e}^*(s)) ds - \alpha(\theta_1)[\theta_1 + \hat{e}^*(\theta_1)]$$

together with a piece-rate $\alpha(\theta_1)$. These contracts are reminiscent of those derived in Laffont and Tirole (1986) in a static regulatory setting. Contrary to the static case, the entire linear scheme $S(\theta_1) + \alpha(\theta_1)\tilde{\pi}_t$ — as opposed to the point $S(\theta_1) + \alpha(\theta_1)[\theta_1 + \hat{e}^*(\theta_1)]$ — is now used over time. This is a direct consequence of the fact that the firm's performance $\tilde{\pi}_t$ now changes stochastically over time in response to the shocks $\tilde{\varepsilon}^t$. Also note that while the optimal mechanism can be implemented by using in each period the static optimal contract for period one, this does not mean that the dynamic optimal mechanism coincides with a sequence of static optimal contracts, as in Baron and Besanko (1984). Rather the opposite. In fact, because the agent's type θ_t (and its distribution) changes over time, the sequence of static optimal contracts entails a different choice of effort for each period. What the result then implies is that, despite the lack of stationarity, it is optimal for the principal to commit to the same reward scheme (and to induce the same effort) as if the agent's type were constant over time.

Out of curiosity, also note that the optimal reward scheme (and the corresponding effort dynamics) when θ_t follows a random walk coincide with the one that the principal would offer in an

environment in which the shocks have only a transitory (as opposed to permanent) effect on the firm’s performance. More generally, assuming $\mathbb{E}[\tilde{\varepsilon}_t] = 0$ for all $t > 1$ and letting $(a_s^t)_{s,t}$ denote arbitrary scalars, the optimal contract is the same when $\theta_t = \gamma_t' \theta_1 + \sum_{s=2}^t a_s^t \varepsilon_s$ as when $\theta_t = \gamma_t' \theta_1 + \varepsilon_t$.

Seniority. While the examples above highlight interesting properties for the dynamics of effort, they also have important implications for the dynamics of the optimal reward scheme. What these examples have in common is the fact that the effect of the agent’s first-period type on his future types declines over time (strictly in the first example). We find this property of “declining correlation” to be reasonable for many stochastic processes describing the evolution of the agent’s productivity. As anticipated in the introduction, this property has implications for the dynamics of the optimal reward scheme. In particular, it helps understand why it may be optimal to reward managers with a longer tenure with a more high-powered incentive scheme, e.g. by giving them more equity in the firm. To illustrate, consider the case presented in Example 1 above, and note that in this case

$$\alpha_t(\theta_1) = 1 - \eta(\theta_1) (\beta - \rho) \beta^{t-2} [\psi''(\hat{e}_t^*(\theta_1)) - \delta \rho \beta \psi''(\hat{e}_{t+1}^*(\theta_1))]. \quad (8)$$

This term, which captures the power of the incentive scheme, is typically increasing in t (it is easy to see that this is the case, for example, when $\rho = 0$ —in which case $\alpha_t(\theta_1)$ reduces to $\psi'(\hat{e}_t^*(\theta_1))$ —or when ψ is quadratic).

Note that the reason why the power of the incentive scheme here increases over time is not driven by variations of the manager’s preferences. It is merely a consequence of the fact that, when he was hired, the manager possessed relevant private information about his ability to generate profits for the firm. In the case of an AR(1) process, the correlation between the manager’s initial type and his future types declines over time. This implies that, to minimize the informational rents that the firms’ shareholders must leave to the manager, it is optimal to (downward) distort the agent’s effort more when he is “young” than when he is “old”. Because the manager’s effort is increasing in the sensitivity α_t of his reward scheme to the firm’s performance π_t , this in turn implies that it is optimal to give the manager a more “high powered” incentive scheme when he is “senior” than when he is “young”.

Clearly, as mentioned in the introduction, other explanations for seniority have been suggested in the literature. Gibbons and Murphy (1991), for example, argue that career-concern incentives decline over time and, by implication, managers with a higher tenure must be provided with stronger “explicit contracts”, i.e. with more high-powered incentive schemes. In their model, explicit incentives are a substitute for career-concern incentives.¹⁸

Another explanation for the correlation between seniority and the power of the incentive scheme may come from the fact that the disutility of effort may decline over time, most notably as the result of learning by doing. While we find such explanations plausible in certain environments, what our results indicate is that, even in the absence of any assumption of time-variant preferences/technologies/career concerns, seniority may arise quite naturally as the result of an optimal intertemporal screening problem in settings in which the correlation between the manager initial type/talent and his future ones declines over time. We believe this is a plausible assumption for most environments of interest.

3 Fully-contingent effort policies

Consider now an environment in which the process for θ_t does not satisfy the SFC condition. When this is the case, the optimal effort policy typically depends not only on θ_1 but also on the realization of the shocks ε^t . In many cases of interest, the optimal mechanism can still be implemented by a menu of linear contracts, but the agent must now be allowed to change the slope of these contracts over time in response to the shocks. To illustrate, assume that $\rho = 0$, so that effort has only a transitory effect on the firm’s performance, that $T < +\infty$,¹⁹ that the stochastic process governing the evolution of θ_t is *Markov* so that each kernel $F_t(\cdot|\theta^{t-1})$ depends on θ^{t-1} only through θ_{t-1} . Finally, assume that, for any t any θ_{t-1} , $F_t(\cdot|\theta_{t-1})$ is absolutely continuous and strictly increasing over Θ_t with density $f_t(\theta_t|\theta_{t-1}) > 0$ for all $\theta_t \in (\underline{\theta}_t, \bar{\theta}_t)$, and that, for each t , there exists an integrable function $B_t : \Theta_t \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ such that, for any $\theta^t \in \Theta^t$, $\partial F_t(\theta_t|\theta_{t-1})/\partial \theta_{t-1}$ exists and $|\partial F_t(\theta_t|\theta_{t-1})/\partial \theta_{t-1}| \leq B_t(\theta_t)$.²⁰

¹⁸For a detailed analysis of career concerns incentives, see Dewatripont, Jewitt and Tirole (1999).

¹⁹The results in this section actually extend to $T = +\infty$ under mild additional conditions.

²⁰Throughout, if $\theta_{t-1} = \underline{\theta}_{t-1}$, then $\partial F_t(\theta_t|\underline{\theta}_{t-1})/\partial \theta_{t-1}$ denotes the right derivative of F_t with respect to θ_{t-1} .

Following steps similar to those used in the proof of Proposition 2, it is easy to see that the solution to the principal's relaxed program is an effort policy \hat{e}^* that is implicitly defined by the following conditions²¹

$$\psi'(\hat{e}_t^*(\theta_1, \varepsilon^t)) = 1 - \eta(\theta_1) \frac{\partial z_t(\theta_1, \varepsilon^t)}{\partial \theta_1} \psi''(\hat{e}_t^*(\theta_1, \varepsilon^t)) \quad (9)$$

where (z, G) is *any* independent shock representation for the process that corresponds to the kernels $F = \langle F_t(\cdot|\cdot) \rangle_{t=1}^T$.

Equivalently, this condition can be expressed in terms of the primitive representation F as follows. Consider the mechanism Ω where in each period the agent reports θ_t (as opposed to ε_t). Following steps similar to those in the proof of Proposition 2 (see also Proposition 2 in Pavan, Segal, and Toikka (2009,a)), one can show that, in any IC mechanism, after almost every truthful history²² h_{t-1} , the value function $V^\Omega(\theta^{t-1}, \theta_t)$ is Lipschitz continuous in θ_t and, for almost every θ_t ,

$$\frac{\partial V^\Omega(\theta^t)}{\partial \theta_t} = \mathbb{E}_{\tilde{\theta}^T | \theta^t} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_t^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] \quad (10)$$

where for all t , $J_t^t(\theta^t) \equiv 1$, and for any $\tau > t$,

$$J_t^\tau(\theta^\tau) \equiv \sum_{\substack{K \in \mathbb{N}, l \in \mathbb{N}^{K+1}, k=1 \\ t=l_0 < \dots < l_K = \tau}} \prod_{k=1}^K I_{l_{k-1}}^{l_k}(\theta^{l_k}),$$

with

$$I_l^m(\theta^m) \equiv - \frac{\partial F_m(\theta_m | \theta^{m-1}) / \partial \theta_l}{f_m(\theta_m | \theta^{m-1})}.$$

The function $J_t^\tau(\theta^\tau)$ is an impulse-response function that captures the total effect of a variation of θ_t on the distribution of θ_τ taking into account all effects on intermediate types $(\theta_{t+1}, \dots, \theta_{\tau-1})$.

While condition (10) applies to *any* (differentiable) process, in the case of a Markov process, because each $I_l^m(\theta^m)$ is equal to zero for all $l < m-1$ and depends on θ^m only through (θ_m, θ_{m-1}) , the impulse response $J_t^\tau(\theta^\tau)$ reduces to a function of $(\theta_t, \dots, \theta_\tau)$ only and can be written as $J_t^\tau(\theta_t, \theta_{t+1}, \dots, \theta_\tau) =$

²¹Again, this presumes that the RHS of (9) evaluated at $e = 0$ is positive, which is the case when $\psi''_+(0)$ is not too high. When this is not the case, then $\hat{e}_t^*(\theta_1, \varepsilon^t) = 0$.

²²A truthful history h_{t-1} is one that is reached by reporting θ^{t-1} truthfully and following the principal's effort recommendations in each period $s = 1, \dots, t-1$. For simplicity, whenever there is no risk of confusion, we will denote a truthful history simply by θ^{t-1} .

$\Pi_{k=t+1}^\tau I_{k-1}^k(\theta_k, \theta_{k-1})$, with each I_{k-1}^k given by

$$I_{k-1}^k(\theta_k, \theta_{k-1}) = \frac{-\partial F_k(\theta_k | \theta_{k-1}) / \partial \theta_{k-1}}{f_k(\theta_k | \theta_{k-1})}.$$

Applying condition (10) to $t = 1$, we then have that

$$V^\Omega(\theta_1) = \mathbb{E}_{(\tilde{\theta}_2, \dots, \tilde{\theta}_T) | \theta_1} \left[\sum_{t=1}^T \delta^{t-1} \int_{\theta_1}^{\theta_1} J_1^t(s, \tilde{\theta}_2, \dots, \tilde{\theta}_t) \psi'(e_t(s, \tilde{\theta}_2, \dots, \tilde{\theta}_t)) ds \right] + V^\Omega(\underline{\theta}_1).$$

Integrating by parts, this implies that the expected ex-ante surplus for the agent is given by

$$\mathbb{E} \left[V^\Omega(\tilde{\theta}_1) \right] = \mathbb{E}_{\tilde{\theta}^T} \left[\eta(\tilde{\theta}_1) \sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(e_t(\tilde{\theta}^t)) \right] + V^\Omega(\underline{\theta}_1).$$

The principal's expected payoff is thus given by

$$\mathbb{E}[U^P] = \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left\{ \tilde{\theta}_t + \hat{e}_t(\tilde{\theta}^t) - \psi(\hat{e}_t(\tilde{\theta}^t)) - \eta(\tilde{\theta}_1) J_1^t(\tilde{\theta}^t) \psi'(e_t(\tilde{\theta}^t)) \right\} \right] - V^\Omega(\underline{\theta}_1).$$

Provided that $J_1^t(\theta^t) \geq 0$ for each t all θ^t , the optimal effort policy can then be obtained by pointwise maximization of $\mathbb{E}[U^P]$ and is given by

$$\psi'(e_t^*(\theta^t)) = 1 - \eta(\theta_1) J_1^t(\theta^t) \psi''(e_t^*(\theta^t))$$

if $1 - \eta(\theta_1) J_1^t(\theta^t) \psi''(0) > 0$ and by $e_t^*(\theta^t) = 0$ otherwise.

This condition is the analogue of (9) expressed in terms of the primitive representation (the one where the agent reports θ_t as opposed to ε^t). From the same arguments as in the previous section, it then follows that, if there exists a payment scheme s that implements the effort policy e^* and gives zero expected surplus to the lowest period-one type (i.e. such that $V^\Omega(\underline{\theta}_1) = 0$) then, together with the effort policy e^* , such a payment scheme is part of an optimal mechanism.

Now consider the following class of payment schemes. In each period t , the principal pays the agent a fixed amount $S_t(\theta^t)$ and a linear bonus $\alpha_t(\theta^t)\pi_t$, where both S_t and α_t are now allowed to depend on the entire history of reports θ^t (equivalently, S_t and α_t are chosen by the agent out of a menu, as a function of the observed shocks θ^t). In what follows, we show that when the desired effort policy e^* satisfies a certain single-crossing condition, which is the analogue of condition (3) in the previous section, then the policy e^* can be implemented by a reward scheme in this class. To see this, for any t , let

$$\alpha_t(\theta^t) = \psi'(e_t^*(\theta^t)).$$

The sequence of fixed payments $S_t(\theta^t)$ is then defined recursively as follows. For $t = T$, let

$$S_T(\theta^T) \equiv \psi(e_T^*(\theta^T)) + \int_{\underline{\theta}_T}^{\theta_T} \psi'(e_T^*(\theta^{T-1}, s)) ds - \alpha_T(\theta^T) \pi_T^*(\theta^T),$$

while for any $t < T$,

$$\begin{aligned} S_t(\theta^t) &\equiv \psi(e_t^*(\theta^t)) - \alpha_t(\theta^t) \pi_t^*(\theta^t) \\ &+ \int_{\underline{\theta}_t}^{\theta_t} \mathbb{E}_{(\tilde{\theta}_{t+1}, \dots, \tilde{\theta}_T) | s} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_t^\tau \left(s, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau \right) \psi'(e_t^*(\theta^{t-1}, s, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau)) \right] ds \\ &- \mathbb{E}_{\tilde{\theta}^T | \theta^t} \left[\sum_{\tau=t+1}^T \delta^{\tau-t} \left(S_\tau(\tilde{\theta}^\tau) + \alpha_\tau(\tilde{\theta}^\tau) \pi_\tau^*(\tilde{\theta}^\tau) - \psi(e_\tau^*(\tilde{\theta}^\tau)) \right) \right] \end{aligned} \quad (11)$$

where, for any $j = 1, \dots, T$, any $\theta^j \in \Theta^j$, $\pi_j^*(\theta^j) \equiv \theta_j + e_j^*(\theta^j)$.

Now suppose $t = T$ and that the history of past reports is $\hat{\theta}^{T-1}$. It is then immediate that, irrespective of the true shocks θ^T , if the agent reports $\hat{\theta}_T$ in period T , he then finds it optimal to choose effort $e_T^*(\hat{\theta}^{T-1}, \hat{\theta}_T)$. Because the environment is Markov, it is also immediate that, irrespective of whether the history of past reports $\hat{\theta}^{T-1}$ was truthful, an agent whose period- T type is θ_T always finds it optimal to report truthfully in period T . This follows from arguments similar to those used to establish Proposition 2. To see this, note that the *continuation* payoff that type θ_T obtains by reporting $\hat{\theta}_T$ is simply²³

$$u_T(\theta_T, \hat{\theta}_T; \hat{\theta}^{T-1}) \equiv \int_{\underline{\theta}_T}^{\hat{\theta}_T} \psi'(e_T^*(\hat{\theta}^{T-1}, s)) ds + \alpha_T(\hat{\theta}^{T-1}, \hat{\theta}_T) [\theta_T - \hat{\theta}_T]. \quad (12)$$

Now, let

$$u_T(\theta_T; \hat{\theta}^{T-1}) \equiv u_T(\theta_T, \theta_T; \hat{\theta}^{T-1}) = \int_{\underline{\theta}_T}^{\theta_T} \psi'(e_T^*(\hat{\theta}^{T-1}, s)) ds \quad (13)$$

denote the continuation payoff that type θ_T obtains by reporting truthfully. It is then immediate that

$$\begin{aligned} \left[\frac{du_T(\theta_T; \hat{\theta}^{T-1})}{d\theta_T} - \frac{\partial u_T(\theta_T, \hat{\theta}_T; \hat{\theta}^{T-1})}{\partial \theta_T} \right] &= \psi'(e_T^*(\hat{\theta}^{T-1}, \theta_T)) - \alpha_T(\hat{\theta}^{T-1}, \hat{\theta}_T) \\ &= \psi'(e_T^*(\hat{\theta}^{T-1}, \theta_T)) - \psi'(e_T^*(\hat{\theta}^{T-1}, \hat{\theta}_T)), \end{aligned}$$

and hence

$$\left[\frac{du_T(\theta_T; \hat{\theta}^{T-1})}{d\theta_T} - \frac{\partial u_T(\theta_T, \hat{\theta}_T; \hat{\theta}^{T-1})}{\partial \theta_T} \right] [\theta_T - \hat{\theta}_T] \geq 0 \quad (14)$$

²³In what follows, by continuation payoff, we mean the discounted sum of the future flow payoffs.

if and only if $e_T^*(\hat{\theta}^{T-1}, \cdot)$ is increasing. As it is well known, condition (14) guarantees that truthtelling is optimal (see, e.g. Garcia, 2005).

Now, by induction, suppose that, irrespective of whether he has reported truthfully in the past, at any period $\tau > t$, the agent finds it optimal to report θ_τ truthfully. Then, consider the agent's incentives in period t . Take any history of reports $\hat{\theta}^{t-1}$. Again, because the environment is Markov, it is irrelevant whether this history corresponds to the truth or not. Then suppose the agent's true type in period t is θ_t and he announces $\hat{\theta}_t$. His continuation payoff is then given by

$$\begin{aligned} u_t(\theta_t, \hat{\theta}_t; \hat{\theta}^{t-1}) &= u_t(\hat{\theta}_t; \hat{\theta}^{t-1}) + \alpha_t(\hat{\theta}^{t-1}, \hat{\theta}_t)[\theta_t - \hat{\theta}_t] \\ &\quad + \mathbb{E}_{\tilde{\theta}_{t+1}|\theta_t} \left[u_{t+1}(\tilde{\theta}_{t+1}; \hat{\theta}^{t-1}, \hat{\theta}_t) \right] - \mathbb{E}_{\tilde{\theta}_{t+1}|\hat{\theta}_t} \left[u_{t+1}(\tilde{\theta}_{t+1}; \hat{\theta}^{t-1}, \hat{\theta}_t) \right] \end{aligned} \quad (15)$$

where, for any period $l \geq 1$ and any $(\theta_l, \hat{\theta}^{l-1})$,

$$u_l(\theta_l; \hat{\theta}^{l-1}) = \int_{\underline{\theta}_l}^{\theta_l} \mathbb{E}_{(\tilde{\theta}_{l+1}, \dots, \tilde{\theta}_T)|s} \left[\sum_{\tau=l}^T \delta^{\tau-l} J_\tau^\tau \left(s, \tilde{\theta}_{l+1}, \dots, \tilde{\theta}_\tau \right) \psi'(e_l^*(\hat{\theta}^{l-1}, s, \tilde{\theta}_{l+1}, \dots, \tilde{\theta}_\tau)) \right] ds \quad (16)$$

is the equilibrium continuation payoff under a truthful and obedient strategy starting from period l onwards, given the current type θ_l and the history of past reports $\hat{\theta}^{l-1}$. It follows that

$$\frac{du_t(\theta_t; \hat{\theta}^{t-1})}{d\theta_t} = \mathbb{E}_{(\tilde{\theta}_{t+1}, \dots, \tilde{\theta}_T)|\theta_t} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_\tau^\tau \left(\theta_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau \right) \psi'(e_t^*(\hat{\theta}^{t-1}, \theta_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau)) \right]$$

and that²⁴

$$\begin{aligned} \frac{\partial u_t(\theta_t, \hat{\theta}_t; \hat{\theta}^{t-1})}{\partial \theta_t} &= \alpha_t(\hat{\theta}^{t-1}, \hat{\theta}_t) + \frac{\partial \mathbb{E}_{\tilde{\theta}_{t+1}|\theta_t} \left[u_{t+1}(\tilde{\theta}_{t+1}; \hat{\theta}^{t-1}, \hat{\theta}_t) \right]}{\partial \theta_t} \\ &= \mathbb{E}_{(\tilde{\theta}_{t+1}, \dots, \tilde{\theta}_T)|\theta_t} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_\tau^\tau \left(\theta_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau \right) \psi'(e_t^*(\hat{\theta}^{t-1}, \theta_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau)) \right]. \end{aligned} \quad (17)$$

Once again, a sufficient condition for $u_t(\theta_t; \hat{\theta}^{t-1}) \geq u_t(\theta_t, \hat{\theta}_t; \hat{\theta}^{t-1})$ for any $\hat{\theta}_t$ is that

$$\left[\frac{du_t(\theta_t; \hat{\theta}^{t-1})}{d\theta_t} - \frac{\partial u_t(\theta_t, \hat{\theta}_t; \hat{\theta}^{t-1})}{\partial \theta_t} \right] [\theta_t - \hat{\theta}_t] \geq 0, \quad (18)$$

or equivalently that

$$\mathbb{E}_{(\tilde{\theta}_{t+1}, \dots, \tilde{\theta}_T)|\theta_t} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_\tau^\tau \left(\theta_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau \right) [\alpha_\tau(\hat{\theta}^{t-1}, \theta_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau) - \alpha_\tau(\hat{\theta}^{t-1}, \hat{\theta}_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau)] \right] [\theta_t - \hat{\theta}_t] \geq 0. \quad (19)$$

²⁴The expression in (17) is obtained by integration by parts, using (16).

This condition is the equivalent of condition (3) in the previous section. Note that, this condition is satisfied, for example, when the effort policy is *strongly monotone*, i.e. when at any period t , $e_t^*(\theta^t)$ is nondecreasing in θ^t . We then have the following result.

Proposition 3 *Assume the evolution of θ_t is governed by a Markov process satisfying the assumptions described above and that, for each period t , $\pi_t = \theta_t + e_t$.*

1. *Any effort policy satisfying the single-crossing condition (19) for any t , any $(\hat{\theta}^{t-1}, \hat{\theta}_t, \theta_t)$, can be implemented by the following linear pay package: In every period t , given any history of reports θ^t and any history of observed performances π^t , the principal pays the agent*

$$s_t(\theta^t, \pi^t) = S_t(\theta^t) + \alpha_t(\theta^t)\pi_t,$$

where $\alpha_t(\theta^t) \equiv \psi'(e_t(\theta^t))$ and where the fixed payment $S_t(\cdot)$ is as in (11).

2. *Let e^* be the effort policy implicitly defined, for all t and all $\theta^t \in \Theta^t$, by*

$$\psi'(e_t^*(\theta^t)) = 1 - \eta(\theta_1)J_1^t(\theta^t) \psi''(e_t^*(\theta^t)) \quad (20)$$

unless $\psi''(0) \geq 1/[\eta(\theta_1)J_1^t(\theta^t)]$, in which case $e_t^(\theta^t) = 0$. Assume e^* satisfies the single-crossing condition of (19). Then e^* , together with the linear pay package s^* described in part (1), are part of an optimal mechanism.*

A few remarks are in order. First, note that the result in Proposition 3 complements that in Proposition 2: while Proposition 3 does not restrict the process for θ_t to satisfy the SFC condition, it restricts θ_t to follow a Markov process, a property that is not required by Proposition 2.

Second, note that the linear scheme in Proposition 3 has the appealing property of guaranteeing that, even if the agent has the option of leaving the relationship at any point in time, he never finds it optimal to do so, i.e. it guarantees participation at any period, after any history.

Third note that a key distinction between the linear scheme of Proposition 3 and that of Proposition 2 is that the agent must now be allowed to propose changes to his pay package over time. These changes are in response to the shocks θ^t . This finding is consistent with some of the

recent literature on managerial compensation which documents that CEO compensation is often proposed by CEOs themselves (see e.g. Bebchuck and Fried, 2004). In our setting, the firm's shareholders (the principal) set in advance broad restrictions on the CEO's pay package but then delegate to the latter the choice of the specific terms of the reward scheme so as to permit him to respond to (unverifiable) variations in the environment. In particular, the optimal mechanism involves offering the CEO a menu of linear contracts with *memory*, in the sense that the set of possible packages available for period t depends on the reward packages selected in past periods (as indexed by θ^{t-1}).

Fourth, note that a form of seniority is likely to hold also in this environment, albeit only in expectation. For example, suppose ψ is quadratic. Then, by inspecting (20), one can see that the power of the incentive scheme, as captured by e_t increases, on average, with the manager's tenure, provided that $\mathbb{E}_{\tilde{\theta}_t|\theta^{t-1}} \left[J_1^t \left(\theta^{t-1}, \tilde{\theta}_t \right) \right] \leq J_1^{t-1} (\theta^{t-1})$. As discussed in the introduction, this property is satisfied by many stochastic processes that feature a correlation between θ_1 and θ_t declining in t .

Lastly note that, while the possibility of implementing the policy e^* that solves the relaxed program (as given by (20)) with a menu of linear schemes is certainly appealing, such a possibility cannot be taken for granted. In fact, in many cases of interest, e^* does not satisfy the single-crossing condition of (19). To see this, assume that, for any $t > 1$ and any θ_{t-1} , $I_{t-1}^t(\cdot, \theta_{t-1})$ is continuous and $\lim_{\theta_t \rightarrow \underline{\theta}_t} I_{t-1}^t(\theta_t, \theta_{t-1}) = \lim_{\theta_t \rightarrow \bar{\theta}_t} I_{t-1}^t(\theta_t, \theta_{t-1}) = 0$.²⁵ Then for any $1 < s \leq \tau$, any θ_{-s}^τ , $\lim_{\theta_s \rightarrow \underline{\theta}_s} J_1^\tau(\theta_1, \dots, \theta_s, \dots, \theta_\tau) = \lim_{\theta_s \rightarrow \bar{\theta}_s} J_1^\tau(\theta_1, \dots, \theta_s, \dots, \theta_\tau) = 0$. This in turn implies that $\lim_{\theta_s \rightarrow \bar{\theta}_s} e_\tau^*(\theta_{-s}^\tau, \theta_s) = \lim_{\theta_s \rightarrow \underline{\theta}_s} e_s^*(\theta_{-s}^\tau, \theta_s) = e_s^{FB}$. The policy $e_\tau^*(\theta_{-s}^\tau, \theta_s)$ is then typically non-monotone in θ_s , for any $\tau \geq s$ any θ_{-s}^τ , which makes it difficult (if not impossible) to satisfy (19).

Motivated by the aforementioned considerations about the possible difficulties of implementing the optimal effort policy with linear schemes, we now consider an alternative implementation based on the “trick” used to establish Lemma 1 in the proof of Proposition 2. The idea is to charge the agent a sufficiently large penalty L whenever, given the announcements θ^t , the observed profits

²⁵Note that, under our assumption of full support (i.e. F_t strictly increasing) over Θ_t , these conditions hold, for example, when $\bar{\theta}_t < +\infty$ and when F_t is an atomless distribution with density strictly positive over $[\underline{\theta}_t, \bar{\theta}_t]$.

are different from the equilibrium ones $\pi_t^*(\theta^t)$. To see how this permits one to relax condition (19), suppose that in all periods $t < T$ the principal uses the same reward scheme as in Part 1 in Proposition 3, whereas at $t = T$, she uses the following scheme

$$s_T(\theta^T, \pi^T) = \begin{cases} \psi(e_T^*(\theta^T)) + \int_{\underline{\theta}_T}^{\theta_T} \psi'(e_T^*(\theta^{T-1}, s))ds & \text{if } \pi_T = \pi_T^*(\theta^T) \\ -L & \text{otherwise} \end{cases}. \quad (21)$$

Note that, conditional on “meeting the target”, under the new scheme, for any sequence of reports θ^T , the agent receives exactly the same compensation he would have obtained under the original linear scheme by choosing effort in period t so as to attain profits $\pi^T(\theta^T)$. Provided that L is large enough, it is then immediate that deviations from the equilibrium strategy are less profitable under the new scheme than under the original linear one. In particular, the agent’s continuation payoff in period T , after he has reported $(\hat{\theta}^T)$ and experienced a shock θ_T in period t , is now given by

$$\begin{aligned} \hat{u}_T(\theta_T, \hat{\theta}_T; \hat{\theta}^{T-1}) &\equiv \int_{\underline{\theta}_T}^{\hat{\theta}_T} \psi'(e_T^*(\hat{\theta}^{T-1}, s))ds + \psi\left(\pi_T(\hat{\theta}^{T-1}, \hat{\theta}_T) - \hat{\theta}_T\right) - \psi\left(\pi_T(\hat{\theta}^{T-1}, \hat{\theta}_T) - \theta_T\right) \\ &= \int_{\underline{\theta}_T}^{\hat{\theta}_T} \psi'(e_T^*(\hat{\theta}^{T-1}, s))ds + \psi(e_T^*(\hat{\theta}^{T-1}, \hat{\theta}_T)) - \psi(e_T^*(\hat{\theta}^{T-1}, \hat{\theta}_T) + \hat{\theta}_T - \theta_T) \end{aligned}$$

rather than $u_T(\theta_T, \hat{\theta}_T; \hat{\theta}^{T-1})$ as in Equation (12). Irrespective of whether $\hat{\theta}^{T-1}$ was truthful or not, incentive compatibility is then ensured in period T (i.e., the agent finds it optimal to report θ_T truthfully and then choose the equilibrium level of effort $e_T^*(\hat{\theta}^{T-1}, \theta_T)$) if the effort policy e^* satisfies the analogue of condition (14) with $u_T(\theta_T, \hat{\theta}_T; \theta^{T-1})$ now replaced by the function $\hat{u}_T(\theta_T, \hat{\theta}_T; \theta^{T-1})$, that is, if²⁶

$$\left[\psi'(e_T^*(\hat{\theta}^{T-1}, \theta_T)) - \psi'(e_T^*(\hat{\theta}^{T-1}, \hat{\theta}_T) + \hat{\theta}_T - \theta_T)\right] \left[\theta_T - \hat{\theta}_T\right] \geq 0. \quad (22)$$

Note that condition (22) is clearly weaker than condition (19) which requires $[\psi'(e_T^*(\hat{\theta}^{T-1}, \theta_T)) - \psi'(e_T^*(\hat{\theta}^{T-1}, \hat{\theta}_T))][\theta_T - \hat{\theta}_T] \geq 0$. Moving from the linear scheme to this alternative scheme thus permits one to implement effort policies that are not necessarily monotone in the shock θ_T . It is easy to see that condition (22) is equivalent to requiring that the profit function $\pi_T(\hat{\theta}^{T-1}, \cdot)$ (as opposed to the effort policy $e_T^*(\hat{\theta}^{T-1}, \cdot)$) being non-decreasing. Absent the dependence on history, this is the same result found by Laffont and Tirole (1993, A1.4) for the static case.

²⁶ As mentioned above, note that the payoff under truthtelling under the new scheme is exactly the same as under the original scheme. That is $u_T(\theta_T; \hat{\theta}^{T-1})$ continues to be as in (13).

Now suppose the principal replaces the entire linear scheme s^* with the incentive scheme s recursively defined, for each t , as follows

$$s_t(\theta^t, \pi^t) = \begin{cases} \psi(e_t^*(\theta^t)) + \int_{\hat{\theta}_t}^{\theta_t} \mathbb{E}_{(\tilde{\theta}_{t+1}, \dots, \tilde{\theta}_T) | s} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_t^\tau(s, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau) \psi'(e_t^*(\theta^{t-1}, s, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau)) \right] ds \\ - \mathbb{E}_{\tilde{\theta}^T | \theta^t} \left[\sum_{\tau=t+1}^T \delta^{\tau-t} \left(s_\tau(\tilde{\theta}^\tau, \pi^{\tau*}(\tilde{\theta}^\tau)) - \psi(e_\tau^*(\tilde{\theta}^\tau)) \right) \right] & \text{if } \pi_t \geq \pi_t^*(\theta^t) \\ -L & \text{otherwise.} \end{cases} \quad (23)$$

where $\pi^{\tau*}(\tilde{\theta}^\tau) \equiv (\pi_s^*(\theta^s))_{s=1}^\tau$. Note that, for $t = T$, this scheme is the same as the one in (21).

Now suppose, by induction, that under the scheme s defined above, truthful reporting is optimal for the agent in each period $\tau > t$, irrespective of the period- τ history (recall that, because the environment is Markov, if truthful reporting is optimal on the equilibrium path, i.e. at a truthful period- τ history, then it is optimal at *all* period- τ histories). Provided L is large enough, the agent's period- t continuation payoff under this scheme when his period- t type is θ_t , he reports $\hat{\theta}_t$, and the sequence of past reports is $\hat{\theta}^{t-1}$, is then given by

$$\begin{aligned} \hat{u}_t(\theta_t, \hat{\theta}_t; \hat{\theta}^{t-1}) &= u_t(\hat{\theta}_t; \hat{\theta}^{t-1}) + \psi(e_t^*(\hat{\theta}^{t-1}, \hat{\theta}_t)) - \psi(e_t^*(\hat{\theta}^{t-1}, \theta_t) + \hat{\theta}_t - \theta_t) \\ &\quad + \mathbb{E}_{\tilde{\theta}_{t+1} | \theta_t} \left[u_{t+1}(\tilde{\theta}_{t+1}; \hat{\theta}^{t-1}, \hat{\theta}_t) \right] - \mathbb{E}_{\tilde{\theta}_{t+1} | \hat{\theta}_t} \left[u_{t+1}(\tilde{\theta}_{t+1}; \hat{\theta}^{t-1}, \hat{\theta}_t) \right], \end{aligned}$$

where, for any period $l \geq 1$ and any $(\theta_l, \hat{\theta}^{l-1})$, $u_l(\theta_l; \hat{\theta}^{l-1})$ continues to denote the equilibrium continuation payoff, as defined in (16). Incentive compatibility is then guaranteed in period t if condition (18) holds, that is, if

$$\mathbb{E}_{(\tilde{\theta}_{t+1}, \dots, \tilde{\theta}_T) | \theta_t} \left[\begin{aligned} &\psi'(e_t^*(\hat{\theta}^{t-1}, \theta_t)) - \psi'(e_t^*(\hat{\theta}^{t-1}, \hat{\theta}_t) + \hat{\theta}_t - \theta_t) \\ &+ \sum_{\tau=t+1}^T \delta^{\tau-t} J_t^\tau(\theta_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau) \\ &\times [\psi'(e_\tau^*(\hat{\theta}^{t-1}, \theta_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau)) - \psi'(e_\tau^*(\hat{\theta}^{t-1}, \hat{\theta}_t, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau))] \end{aligned} \right] [\theta_t - \hat{\theta}_t] \geq 0. \quad (24)$$

Note that this condition is the same as that in (19) with the initial term $\psi'(e_t^*(\hat{\theta}^{t-1}, \theta_t)) - \psi'(e_t^*(\hat{\theta}^{t-1}, \hat{\theta}_t))$ replaced by $\psi'(e_t^*(\hat{\theta}^{t-1}, \theta_t)) - \psi'(e_t^*(\hat{\theta}^{t-1}, \hat{\theta}_t) + \hat{\theta}_t - \theta_t)$. We then have the following result.

Proposition 4 *Any effort policy satisfying the single-crossing condition (24) for any t any $(\hat{\theta}^{t-1}, \hat{\theta}_t, \theta_t)$, can be implemented by the non-linear pay scheme given in (23).*

As an illustration of how the scheme s given in (23) may help implementing effort policies e^* that solve the principal's relaxed program but that cannot be implemented with the linear scheme s^* of Proposition 3, consider the following example.

Example 4 Suppose that, for any $e \in [0, \bar{e}]$, $\psi(e) = e^2/2$. Let θ_1 be a non-negative random variable with distribution F strictly increasing and absolutely continuous on the interval $[\underline{\theta}_1, \bar{\theta}_1] \subset \mathbb{R}_{++}$ with hazard rate $\eta(\theta_1)$ nonincreasing and such that $\eta(\theta_1) \leq \theta_1$ for each θ_1 .²⁷ Now suppose that, for any $t \geq 2$, $\theta_t = \theta_1 \times \prod_{\tau=2}^t \varepsilon_\tau$, where $\tilde{\varepsilon}^T \equiv (\tilde{\varepsilon}_\tau)_{\tau=2}^T$ is a collection of jointly independent random variables, each independent of θ_1 , each distributed according to the function G strictly increasing and absolutely continuous with density g strictly positive over \mathbb{R}_+ . Let e^* be the effort policy that solves the relaxed program as given in (20). Then the policy e^* cannot be implemented by the linear scheme of Proposition 3 but it can be implemented by the non-linear scheme of Proposition 4.

4 Risk aversion

We now show how the optimal mechanism must be adjusted to accommodate the possibility that the agent is risk averse. We restrict attention here to the case where T is finite. To simplify the notation, we omit discounting, i.e. set $\delta = 1$. We start by assuming that the agent's preferences are represented by a Bernoulli function

$$U^A(c^T, e^T) = \mathcal{V} \left(\sum_{t=1}^T c_t \right) - \sum_{t=1}^T \psi(e_t)$$

where \mathcal{V} is a strictly increasing and (weakly) concave function. This representation is quite common in the literature (e.g. Holmstrom and Milgrom's (1987) seminal paper on linearity and aggregation in dynamic contracting). As is well known, this representation permits one to introduce risk aversion while at the same time avoiding any complication stemming from the desire of consumption smoothing: it is thus appropriate for a setting where the agent cares only about his total compensation and not the way this is distributed over time. We will come back to an alternative representation that accommodates preferences for consumption smoothing at the end of the section.

For the stochastic process for θ_t , we adopt a general independent-shock representation and assume each $z_t(\theta_1, \varepsilon^t)$ is differentiable and equi-Lipschitz continuous.

Since the agent's productivity affects neither the marginal contribution of effort to profits, nor the agent's disutility from it, and since the agent's preferences over money are also separable from

²⁷This condition is satisfied, for instance, when θ_1 is distributed uniformly over the interval $[1, 3/2]$.

the agent's productivity, the first-best solution to the principal's problem is constant both over productivities and time. It is given in the following proposition.

Proposition 5 *Assume, as in Proposition 1, that the agent does not possess private information. The optimal contract for the principal implements the effort policy given by $\hat{e}_t = \hat{e}^{FB}$ for all t , where \hat{e}^{FB} solves*

$$\psi'(\hat{e}^{FB}) = \mathcal{V}'(\mathcal{V}^{-1}(T\psi(\hat{e}^{FB}))). \quad (25)$$

Proposition 5 shows that the first-best effort now depends on the agent's marginal utility of additional payments, evaluated at the payment $\mathcal{V}^{-1}(T\psi(\hat{e}^{FB}))$ that just compensates him for his total disutility of effort. This foreshadows one of the new roles for information rents that we explore below – further payments to cover the agent's information rents when the agent has private information will lower the agent's marginal utility of money, making incentives for effort more costly.

To examine the effects of information rents, one can follow steps similar to those used to establish Proposition 2. The characterization of incentive compatibility is unaffected by the introduction of risk aversion and that the agent's value function in period one remains equal to

$$V^{\hat{\Omega}}(\theta_1) = V^{\hat{\Omega}}(\underline{\theta}_1) + \mathbb{E} \left[\sum_{t=1}^T \int_{\underline{\theta}_1}^{\theta_1} D_{1,t}(s, \tilde{\varepsilon}^t) \psi'(\hat{e}_t(s, \tilde{\varepsilon}^t)) ds \right],$$

where $D_{1,1}(\theta_1) \equiv 1$ and, for any $t > 1$,

$$D_{1,t}(\theta_1, \varepsilon^t) \equiv \frac{\partial z_t(\theta_1, \varepsilon^t)}{\partial \theta_1} - \rho \frac{\partial z_{t-1}(\theta_1, \varepsilon^{t-1})}{\partial \theta_1}$$

with $z_1(\theta_1) \equiv \theta_1$. Note that these $D_{1,t}(\theta_1, \varepsilon^t)$ functions reduce to the corresponding $D_{1,t}(\theta_1)$ functions of Section 2.2 when the stochastic process for θ_t satisfies the SFC condition.

A similar characterization applies to each period $t > 1$. For example, incentive compatibility at any *truthful* history²⁸ $h_{T-1} = (\theta_1, \varepsilon^{T-1})$ implies that $V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \varepsilon_T)$ is Lipschitz continuous in

²⁸Note that incentive compatibility at a truthful history h_t means that the agent's value function in the mechanism $\hat{\Omega}$ after reaching history h_t is equal to the agent's expected payoff when, starting from history h_t the agent follows a truthful and obedient strategy in each period $\tau \geq t$. Also recall that incentive-compatibility in period one, given θ_1 , implies incentive compatibility at almost all (i.e. with probability one) truthful period t -history, $t = 1, \dots, T$.

ε_T and for a.e. ε_T ,

$$\frac{\partial V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \varepsilon_T)}{\partial \varepsilon_T} = \frac{\partial z_T(\varepsilon^{T-1}, \varepsilon_T)}{\partial \varepsilon_T} \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-1}, \varepsilon_T)),$$

which in turn implies that

$$V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \varepsilon_T) = V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \underline{\varepsilon}_T) + \int_{\underline{\varepsilon}_T}^{\varepsilon_T} \frac{\partial z_T(\varepsilon^{T-1}, s)}{\partial \varepsilon_T} \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-1}, s)) ds.$$

Furthermore, using the fact that incentive compatibility implies that $V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \varepsilon_T)$ must coincide with the equilibrium payoff with probability one, we have that, for almost every (h_{T-1}, ε_T) ,

$$\mathcal{V} \left(\sum_{t=1}^T \hat{c}_t(\theta_1, \varepsilon^t) \right) - \sum_{t=1}^T \psi(\hat{e}_t(\theta_1, \varepsilon^t)) = V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \underline{\varepsilon}_T) + \int_{\underline{\varepsilon}_T}^{\varepsilon_T} \frac{\partial z_T(\varepsilon^{T-1}, s)}{\partial \varepsilon_T} \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-1}, s)) ds.$$

This implies that in almost every state $(\theta_1, \varepsilon^T)$ the utility $\mathcal{V} \left(\sum_{t=1}^T \hat{c}_t(\theta_1, \varepsilon^t) \right)$ that the agent assigns to the total payment $\sum_{t=1}^T \hat{c}_t(\theta_1, \varepsilon^t)$ is uniquely determined by the effort policy \hat{e} up to a constant $V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \underline{\varepsilon}_T)$ which may depend on $(\theta_1, \varepsilon^{T-1})$ but is independent of ε_T . Iterating backwards, and noting that for each period t and any history h_t

$$V^{\hat{\Omega}}(h_t) = \mathbb{E}[V^{\hat{\Omega}}(h_t, \tilde{\varepsilon}_{t+1})],$$

the dependence of the constant $V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \underline{\varepsilon}_T)$ on the history $(\theta_1, \varepsilon^{T-1})$ also turns out to be uniquely determined by the effort policy \hat{e} up to a scalar K that does not depend on anything.²⁹

Letting $D_{t,t}(\theta_1, \varepsilon^t) \equiv \partial z_t(\theta_1, \varepsilon^t) / \partial \varepsilon_t$ for each $1 < t \leq T$, as well as

$$D_{t,s}(\theta_1, \varepsilon^s) \equiv \frac{\partial z_s(\theta_1, \varepsilon^s)}{\partial \varepsilon_t} - \rho \frac{\partial z_{s-1}(\theta_1, \varepsilon^{s-1})}{\partial \varepsilon_t}$$

for any $s > t$, these arguments lead to Proposition 6.

Proposition 6 *In any incentive-compatible mechanism $\hat{\Omega}$, the total payment to the agent in each state $(\theta_1, \varepsilon^T)$ is given by:*

$$\sum_{t=1}^T \hat{c}_t(\theta_1, \varepsilon^t) = \mathcal{V}^{-1} \left(\sum_{t=1}^T \psi(\hat{e}_t(\theta_1, \varepsilon^t)) + V^{\hat{\Omega}}(\underline{\theta}_1) + \mathbb{E}_{\underline{\varepsilon}^T} \left[\int_{\underline{\theta}_1}^{\theta_1} \sum_{t=1}^T D_{1,t}(s, \tilde{\varepsilon}^t) \psi'(\hat{e}_t(s, \tilde{\varepsilon}^t)) ds \right] + \sum_{t=2}^T \hat{H}_t(\theta_1, \varepsilon^t) \right)$$

²⁹See Proposition 9 in PST for a similar result for quasi-linear settings.

with

$$\begin{aligned} \hat{H}_t(\theta_1, \varepsilon^t) &\equiv \mathbb{E}_{(\tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_t}^{\varepsilon_t} \sum_{\tau=t}^T D_{t,\tau}(\theta_1, \varepsilon^{t-1}, s, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-1}, s, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau)) ds \right] \\ &\quad - \mathbb{E}_{(\tilde{\varepsilon}_t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_t}^{\tilde{\varepsilon}_t} \sum_{\tau=t}^T D_{t,\tau}(\theta_1, \varepsilon^{t-1}, s, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-1}, s, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau)) ds \right] \end{aligned}$$

Using the characterization in Proposition 6, we then have that, in any incentive-compatible mechanism, the principal's expected payoff can be expressed as:

$$\begin{aligned} \mathbb{E}[\hat{U}^P] &= \mathbb{E} \left[\sum_{t=1}^T \hat{\pi}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) \right] - \mathbb{E} \left[\sum_{t=1}^T \hat{c}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) \right] \tag{26} \\ &= \mathbb{E} \left[\sum_{t=1}^T \left(z_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) + \hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) \right) \right] \\ &\quad - \mathbb{E} \left[\mathcal{V}^{-1} \left(\begin{aligned} &\sum_{t=1}^T \psi(\hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t)) + V^{\hat{\Omega}}(\underline{\theta}_1) \\ &+ \mathbb{E}_{\tilde{\varepsilon}^T} \left[\int_{\underline{\theta}_1}^{\tilde{\theta}_1} \sum_{t=1}^T D_{1,t}(s, \tilde{\varepsilon}^t) \psi'(\hat{e}_t(s, \tilde{\varepsilon}^t)) ds \right] + \sum_{t=2}^T \hat{H}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) \end{aligned} \right) \right] \end{aligned}$$

The expression in (26) is the analogue of *dynamic virtual surplus* for the case of a risk-averse agent (it is easy to see that, when \mathcal{V} is the identity function and the process for θ_t satisfies the SFC condition, (26) reduces to the same expression as in (7) by standard integration by parts).

We now turn to the possibility of using “quasi-linear” schemes (i.e. pay packages that are convex in a linear aggregator of the firm's profits) to implement a desired effort policy. We start with the following result.

Proposition 7 *Let \hat{e} be any policy that depends only on time t and on the agent's first-period type θ_1 . Suppose that $\mathbb{E}D_{1,t}(\theta_1, \tilde{\varepsilon}^t) \geq 0$ for any t any θ_1 , and that the policy \hat{e} satisfies the following single-crossing condition*

$$\left[\sum_{t=1}^T \mathbb{E}D_{1,t}(\theta_1, \tilde{\varepsilon}^t) [\alpha_t(\theta_1) - \alpha_t(\hat{\theta}_1)] \right] [\theta_1 - \hat{\theta}_1] \geq 0 \tag{27}$$

for any $\theta_1, \hat{\theta}_1 \in \Theta_1$, where for any $t < T$ and any θ_1 ,

$$\alpha_t(\theta_1) \equiv \psi'(\hat{e}_t^*(\theta_1)) - \rho \psi'(\hat{e}_{t+1}^*(\theta_1)),$$

while for $t = T$, $\alpha_T(\theta_1) \equiv \psi'(\hat{e}_T^*(\theta_1))$. Then the effort policy \hat{e} can be implemented by a “quasi-linear” payment scheme \hat{s}^* according to which the total payment the agent receives when he reports

$(\theta_1, \varepsilon^T)$ and the sequence of observed profits is π^T is given by

$$\sum_{t=1}^T \hat{s}_t^*(\theta_1, \varepsilon^t, \pi^t) = \mathcal{V}^{-1} \left(S(\theta_1) + \sum_{t=1}^T \alpha_t(\theta_1) \pi_t \right), \quad (28)$$

where

$$S_1(\theta_1) \equiv \sum_{t=1}^T \left[\psi(\hat{e}_t^*(\theta_1)) + \int_{\underline{\theta}_1}^{\theta_1} \mathbb{E} D_{1,t}(s, \tilde{\varepsilon}^t) \psi'(\hat{e}_t^*(s)) ds - \alpha_t(\theta_1) \mathbb{E} [\hat{\pi}_t(\theta_1, \tilde{\varepsilon}^t)] \right],$$

with $\hat{\pi}_t(\theta_1, \varepsilon^t) \equiv z_t(\theta_1, \varepsilon^t) + \hat{e}_t(\theta_1) + \sum_{\tau=1}^{t-1} \rho^\tau \hat{e}_{t-\tau}(\theta_1)$.

The proof follows from steps similar to those that establish Proposition 2, adjusted for the fact that the stochastic process for θ_t is here not restricted to satisfy the SFC condition and for the fact that the agent's payoff is now allowed to be concave in his total reward.

The value of the proposition is twofold. Firstly, it guarantees a form of continuity in the optimal mechanism and in the players' payoff with respect to the agent's preferences. In particular, it implies that when \mathcal{V} is sufficiently close to the identity function, the principal can guarantee herself a payoff arbitrarily close to the one she obtains under risk neutrality by choosing to implement the same effort policy as in Proposition 2 and by adjusting the reward scheme as indicated in (28). More generally, the proposition shows how one can adjust the linear reward scheme identified in the baseline model to implement any effort policy that depends only on time t and the agent's first period report θ_1 , provided that such a policy satisfies the single-crossing condition of (27).

We now turn to the characterization of the optimal effort policy. We start by considering policies that depend only on θ_1 , and then turn to general policies. To facilitate the characterization of the necessary conditions, we consider an example in which θ_t follows an ARIMA process (in which case the $D_{t,s}$ functions are scalars) and where the inverse of the agent's utility function over consumption is quadratic.

Example 5 *Suppose that $T < \infty$ and that $\rho = 0$ so that $\pi_t = \theta_t + e_t$, $t = 1, \dots, T$. In addition, suppose that the process governing the evolution of θ_t is ARIMA. Suppose, further, that for any $t \geq 2$,*

$$\theta_t = z_t(\theta_1, \varepsilon^t) = D_{1,t} \theta_1 + \sum_{s=2}^t D_{s,t} \varepsilon_s$$

with $D_{\tau,t} \geq 0$, for any τ, t . Suppose also that θ_1 is distributed according to the c.d.f. F_1 , strictly increasing on an interval $\Theta_1 = [\underline{\theta}_1, \bar{\theta}_1] \subset \mathbb{R}$, and each ε_t distributed according to the c.d.f. G_t , strictly increasing on a compact interval $\mathcal{E}_t \subset \mathbb{R}$. Let the agent's utility function over consumption be given by $\mathcal{V}(c) = \frac{1}{\alpha} \sqrt{2\alpha c + \beta^2} - \frac{\beta}{\alpha}$ with $\alpha, \beta > 0$ and note that this function is chosen so that $\mathcal{V}^{-1}(u) = \frac{\alpha}{2} u^2 + \beta u$. Let $K > 1/\beta$ and assume that, in addition to the assumptions stated above, the function ψ is such that $\log \psi'$ is strictly concave on $(0, \bar{e})$, and that for any $e \geq \bar{e}$, $\psi(e) = \psi(\bar{e}) + (e - \bar{e})K$. There exists an essentially unique³⁰ policy \hat{e}^* that maximizes the principal's expected payoff (as given in 26) among those that depend on θ_1 only. This policy satisfies, for any $t = 1, \dots, T$, almost any θ_1 ,

$$\begin{aligned} & \psi'(\hat{e}_t^*(\theta_1)) \left[\alpha \left(\sum_{s=1}^T \psi(\hat{e}_s^*(\theta_1)) + \int_{\underline{\theta}_1}^{\theta_1} \sum_{s=1}^T D_{1,s} \psi'(\hat{e}_s^*(q)) dq \right) + \beta \right] \\ \geq & 1 - \frac{\psi''(\hat{e}_t^*(\theta_1)) D_{1,t}}{f(\theta_1)} \int_{\underline{\theta}_1}^{\bar{\theta}_1} \left[\alpha \left(\sum_{s=1}^T \psi(\hat{e}_s^*(q)) + \int_{\underline{\theta}_1}^q \sum_{s=1}^T D_{1,s} \psi'(\hat{e}_s^*(r)) dr \right) + \beta \right] f(q) dq \\ & - \alpha \sum_{s=2}^t \left[\left(\sum_{\tau=s}^T D_{s,\tau} \psi'(\hat{e}_\tau^*(\theta_1)) \right) D_{s,t} \psi''(\hat{e}_t^*(\theta_1)) \right] \text{Var}(\varepsilon_s), \end{aligned} \quad (29)$$

with the inequality holding as equality if $\hat{e}_t^*(\theta_1) > 0$. When this policy satisfies the single-crossing condition (27), it can be implemented by the “quasi-linear” payment scheme of Proposition 7.

To shed light on what lies behind Condition (29), recall that the choice of the optimal effort policy trades off two concerns: (1) limiting the agent's intertemporal informational rent (as perceived from a period-1 perspective) and (2) insuring the agent against the risk associated with variations in his reward that are necessary to guarantee incentive-compatibility. To see this more clearly, consider the case where $T = 2$. When applied to $t = 1$ (and assuming an interior solution), Condition (29) becomes

$$\begin{aligned} & \psi'(\hat{e}_1^*(\theta_1)) \left[\alpha \left(\sum_{s=1}^2 \psi(\hat{e}_s^*(\theta_1)) + \int_{\underline{\theta}_1}^{\theta_1} \sum_{s=1}^2 D_{1,s} \psi'(\hat{e}_s^*(q)) dq \right) + \beta \right] \\ = & 1 - \frac{\psi''(\hat{e}_1^*(\theta_1))}{f(\theta_1)} \int_{\underline{\theta}_1}^{\bar{\theta}_1} \left[\alpha \left(\sum_{s=1}^2 \psi(\hat{e}_s^*(q)) + \int_{\underline{\theta}_1}^q \sum_{s=1}^2 D_{1,s} \psi'(\hat{e}_s^*(r)) dr \right) + \beta \right] f(q) dq. \end{aligned} \quad (30)$$

Because θ_1 is known at the time of contracting, the agent does not face any risk concerning his period-one performance and hence the optimal choice of effort for period one is determined uniquely

³⁰The qualifier “essentially” is due to the fact that the optimal policy is determined only almost everywhere.

by the desire to limit the agent's informational rent. It is easy to see that Condition (30) reduces to Condition (2) for the risk-neutral case when $\alpha = 0$ and $\beta = 1$. Note that, as in the risk-neutral case, the policy \hat{e}_1^* is increasing in θ_1 . This simply follows from the fact that the (measure of the) set of types $(\theta_1, \bar{\theta}_1)$ to whom the principal must give a higher rent when he increases $e_1(\theta_1)$ is decreasing in θ_1 . However, contrary to the risk-neutral case, distortions do not vanish "at the top". In fact, when applied to $\bar{\theta}_1$, Condition (30) becomes

$$\psi'(\hat{e}_1^*(\bar{\theta}_1)) \left[\alpha \left(\sum_{s=1}^2 \psi(\hat{e}_s^*(\bar{\theta}_1)) + \int_{\theta_1}^{\bar{\theta}_1} \sum_{s=1}^2 D_{1,s} \psi'(\hat{e}_s^*(q)) dq \right) + \beta \right] = 1, \quad (31)$$

while efficiency (under risk aversion) requires that

$$\psi'(\hat{e}_1^*(\bar{\theta}_1)) \left[\alpha \left(\sum_{s=1}^2 \psi(\hat{e}_s^*(\bar{\theta}_1)) \right) + \beta \right] = 1.$$

The reason is that, with risk aversion, the rent the principal must pay to type $\bar{\theta}_1$ to discourage him from mimicking a lower type reduces $\bar{\theta}_1$'s marginal utility of money; this in turn makes higher effort more costly to sustain which explains why the principal finds it optimal to distort (downwards) $\bar{\theta}_1$'s effort (see also Battaglini and Coate, 2008, for a similar result in a two-type model).

Next, consider the optimal effort policy for $t = 2$. When applied to $t = 2$, Condition (29) becomes (assuming again an interior solution):

$$\begin{aligned} & \psi'(\hat{e}_2^*(\theta_1)) \left[\alpha \left(\sum_{s=1}^2 \psi(\hat{e}_s^*(\theta_1)) + \int_{\theta_1}^{\theta_1} \sum_{s=1}^2 D_{1,s} \psi'(\hat{e}_s^*(q)) dq \right) + \beta \right] \\ &= 1 - \frac{\psi''(\hat{e}_2^*(\theta_1)) D_{1,2}}{f(\theta_1)} \int_{\theta_1}^{\bar{\theta}_1} \left[\alpha \left(\sum_{s=1}^2 \psi(\hat{e}_s^*(q)) + \int_{\theta_1}^q \sum_{s=1}^2 D_{1,s} \psi'(\hat{e}_s^*(r)) dr \right) + \beta \right] f(q) dq \\ & \quad - \alpha \psi'(\hat{e}_2^*(\theta_1)) \psi''(\hat{e}_2^*(\theta_1)) \text{Var}(\varepsilon_2). \end{aligned} \quad (32)$$

The key difference between (32) and (30) is the last term on the right-hand side of the equality. This term captures the principal's concern about exposing the agent to the risk associated with the uncertainty the latter faces about his second period's productivity. Other things equal, this term contributes to reducing effort, as anticipated in the introduction.

To further appreciate the distinctions/similarities between the risk-neutral and the risk-averse case, take the specification of the example, and suppose that $\theta_t = \gamma \theta_{t-1} + \varepsilon_t$, $\gamma \in (0, 1]$, with




Figure 1: Optimal (shock-independent) effort policies: AR1

θ_1 uniformly distributed over $[0, 1]$. In addition, assume that $\beta = 1$, and that $\psi(e) = \frac{e^2}{4}$ for all $e \in (0, \bar{e})$, with $\bar{e} > 4$ and $K = \frac{\bar{e}}{2}$. Then note that, under risk-neutrality ($\alpha = 0, \beta = 1$), the (fully)³¹ optimal policies are given by $\hat{e}_1^*(\theta_1) = 2 - (1 - \theta_1)$ and $\hat{e}_2^*(\theta_1) = 2 - (1 - \theta_1)\gamma$, with $e_1^{FB} = e_2^{FB} = 2$. As discussed in Section 2.2, to minimize the agent's informational rents, the principal finds it optimal to distort both e_1 and e_2 downward. Furthermore, because the effect of the agent's initial type on the distribution of his future types declines over time, it is optimal to distort more in the early stages of the relationship than in the later ones. This property leads to the seniority effect discussed in the previous sections. This effect can be seen easily within the context of this example: the smaller γ is (i.e. the smaller the effect of θ_1 on θ_2) the stronger the seniority effect, with $\hat{e}_2^*(\theta_1) = \hat{e}_1^*(\theta_1)$ [i.e. no seniority] when $\gamma = 1$ [random walk case] and $\hat{e}_2^*(\theta_1) = \hat{e}^{FB}$ when $\gamma = 0$ [θ_1 and θ_2 independent].

Now, to see how risk aversion affects the choice of effort, Figure 1 depicts the optimal (shock-independent) policies for the aforementioned specification with $\alpha = \beta = 1, \gamma = \frac{1}{2}$, and ε_2 uniformly distributed over $[0, 1]$.³²

While a form of seniority continues to hold (\hat{e}_2^* is on average higher than \hat{e}_1^*), risk aversion tends to depress \hat{e}_2^* , thus reducing the optimality of seniority-based reward schemes. Fur-

³¹Recall, from Proposition 2, that with risk-neutrality restricting the policy \hat{e} to depend only on θ_1 is without loss of optimality.

³²We approximated the solution using a sixth-order polynomial.

thermore, now there exist values of θ_1 for which $\hat{e}_2^*(\theta_1) < \hat{e}_1^*(\theta_1)$. To see the reason for this, note that, when evaluated at $\bar{\theta}_1$, equations (30) and (32) are symmetric except for the term $-\alpha\psi''(\hat{e}_2^*(\bar{\theta}_1))\psi'(\hat{e}_2^*(\bar{\theta}_1))\text{Var}(\varepsilon_2)$. As discussed above, this term captures the additional cost associated with a high second-period effort, stemming from the volatility of the agent's payment generated by the shock ε_2 to his second-period productivity. To better appreciate where this term comes from, recall, from Proposition 6, that incentive-compatibility requires that the total payment to the agent in each state $(\theta_1, \varepsilon_2)$ be given by

$$C(\theta_1, \varepsilon_2) = \mathcal{V}^{-1} \left(\sum_{t=1}^2 \psi(\hat{e}_t(\theta_1)) + \psi'(\hat{e}_2(\theta_1))[\varepsilon_2 - \mathbb{E}[\tilde{\varepsilon}_2]] + \int_{\theta_1}^{\theta_1} \left[\sum_{t=1}^2 D_{1,t}\psi'(\hat{e}_t(s)) \right] ds \right). \quad (33)$$

It is then immediate that reducing $\hat{e}_2(\theta_1)$ permits the principal to reduce the agent's exposure to the risk generated by ε_2 . For high values of θ_1 , this new effect dominates the rent-extraction effect documented in the previous section, thus resulting in $\hat{e}_2(\theta_1) < \hat{e}_1(\theta_1)$.

When the effect of θ_1 on θ_2 is small (i.e., for low values of γ), this new effect mitigates but does not overturn the optimality of seniority-based incentive schemes. When instead the effect of θ_1 on θ_2 is strong (i.e., for high values of γ) then this new effect can completely reverse the optimality of incentive schemes whose power increases with time. In the limit, when the shocks to the agent's productivity become fully persistent (θ_t follows a random walk, i.e. $\gamma = 1$), one can then easily see from (30) and (32) that $\hat{e}_2(\theta_1) < \hat{e}_1(\theta_1)$ for all θ_1 .

The aforementioned properties extend to $T > 2$. Figure 2 depicts the optimal policies for the same specification considered above but now letting $T = 3$. The left-hand side is for the case $\gamma = 1/2$, while the right-hand side is for the case $\gamma = 1$ (random walk).

When $\gamma = 1$, effort decreases over time, for all θ_1 . As anticipated in the introduction, this reflects the fact that reducing effort in period t is more effective in reducing the agent's exposure to risk than reducing effort in period $s < t$. When instead $\gamma = 1/2$ then, on average, effort is higher in the early periods than in later ones, but, the opposite is true for high values of θ_1 , as in the $T = 2$ case. Furthermore, effort in later periods can now be decreasing in θ_1 . This follows from the fact that, when t is high, reducing the agent's effort in period t has little effect on the agent's informational

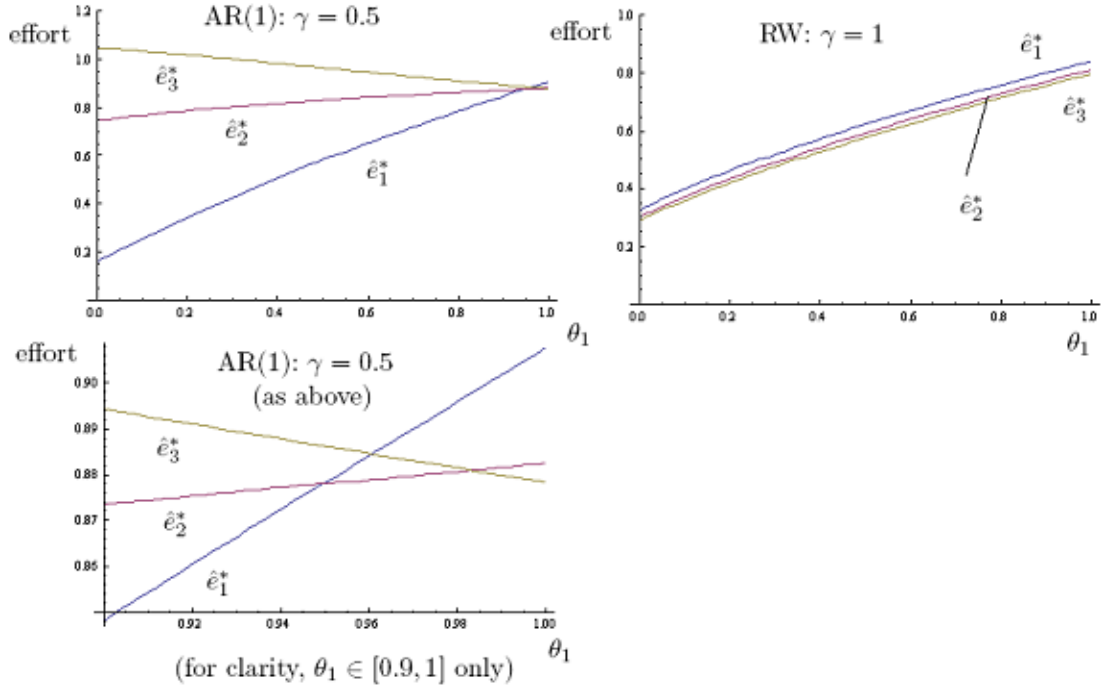


Figure 2: Optimal (shock-independent) effort policies: AR1 and random walk

rent and a strong effect on the agent's exposure to risk (while the opposite is true when t is small). Now, distorting effort in any period t to reduce the agent's rent is always relatively more effective for low types than for high ones (the reason is the same as in static settings): this explains why the optimal effort policy is increasing in θ_1 in the early stages of the relationship. Together with the fact that a reduction in effort in period s is an (imperfect) substitute for a reduction of effort in period $t > s$ on the agent's total exposure to risk, this implies that the optimal effort policy must eventually become decreasing for t sufficiently large.

Another way the principal could mitigate the effect of the volatility of the shocks to the agent's productivity is by conditioning the effort policy on the realization of these shocks. To gauge the effect of this additional flexibility, consider again the same specification assumed above. While a complete analytical characterization of the fully-optimal policy escapes us because of the complexity of the optimization problem, we could approximate the optimal policy with 6th-degree polynomials. The result for the $T = 2$ is depicted in Figure 3, where we considered the same parametrization as




Figure 3: Fully optimal first-period and average second-period effort policies

in Figure 1, but now allowed the second-period effort to depend on the shock ε_2 . Again, when the correlation between θ_1 and θ_2 is not too high (in the example, $\lambda = 1/2$), the optimality of seniority-based schemes is maintained: $\mathbb{E}\hat{e}_2(\theta_1, \tilde{\varepsilon}_2)$ is on average higher than $\hat{e}_1(\theta_1)$, with the inequality reversed for sufficiently high values of θ_1 .

Also note that the second-period effort is typically decreasing in the shock ε_2 , as illustrated in Figure 4. This negative correlation permits the principal to further reduce the agent’s exposure to risk, as one can see directly from (33).

The negative correlation between \hat{e}_2^* and ε_2 also suggests that, in certain environments such as the one considered in this example, it may be difficult to sustain the fully-optimal policy with linear or even “quasi-linear” schemes such as those of Proposition 7. When this is the case, one may need to resort to the type of schemes introduced in Proposition 4, adapted to the presence of risk aversion as indicated in Proposition 8 below. To facilitate the comparison with the results in the previous section, we revert here to the primitive representation where the agent reports θ_t , as opposed to the shocks ε_t . The following proposition then generalizes the results in Propositions 3 and 4 to the case of a (weakly) risk averse agent.

Proposition 8 *Suppose the agent’s type θ_t evolves according to a Markov process and that effort*




Figure 4: Fully-optimal second-period effort

has only a transitory effect on performance, so that $\pi_t = \theta_t + e_t$, all t .

1. Any policy e satisfying the single-crossing condition of (19) for any t , any $(\hat{\theta}^{t-1}, \hat{\theta}_t, \theta_t)$, can be implemented by the following “quasi-linear” scheme: given the reports θ^T and the observed performances π^T , the principal pays the agent a total reward

$$\sum_{t=1}^T s_t(\theta^t, \pi^t) = \mathcal{V}^{-1} \left(\sum_{t=1}^T [S_t(\theta^t) + \alpha_t(\theta^t)\pi_t] \right),$$

where the functions $S_t(\cdot)$ are as in (11) and where $\alpha_t(\theta^t) \equiv \psi'(e_t(\theta^t))$.

2. Any effort policy e satisfying the single-crossing condition (24) for any t , any $(\hat{\theta}^{t-1}, \hat{\theta}_t, \theta_t)$, can be implemented by the following “bonus” scheme: given the reports θ^T and the observed performances π^T , the principal pays the agent a total reward

$$\sum_{t=1}^T s_t(\theta^t, \pi^t) = \mathcal{V}^{-1} \left(\begin{aligned} & \sum_{t=1}^T \psi(e_t(\theta^t)) \\ & + \mathbb{E}_{(\tilde{\theta}_2, \dots, \tilde{\theta}_T) | \theta_1} \left[\int_{\theta_1}^{\theta_1} \sum_{t=1}^T J_1^t(s, \tilde{\theta}_2, \dots, \tilde{\theta}_t) \psi' \left(e_t \left(s, \tilde{\theta}_2, \dots, \tilde{\theta}_t \right) \right) ds \right] \\ & + \sum_{t=2}^T H_t(\theta^t) \end{aligned} \right)$$

if $\pi_t \geq \pi_t^*(\theta^t) \equiv \theta_t + e_t(\theta^t) \forall t$, and $\sum_{t=1}^T s_t(\theta^t, \pi^t) = -L$ otherwise, where, for any $t \geq 2$, any θ^t ,

$$H_t(\theta^t) \equiv \mathbb{E}_{(\tilde{\theta}_{t+1}, \dots, \tilde{\theta}_T) | \theta_t} \left[\int_{\theta_t}^{\theta_t} \sum_{\tau=t}^T J_\tau^t \left(\theta^{t-1}, s, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau \right) \psi' \left(e_\tau \left(\theta^{t-1}, s, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau \right) \right) ds \right] \\ - \mathbb{E}_{(\tilde{\theta}_t, \dots, \tilde{\theta}_T) | \theta_{t-1}} \left[\int_{\theta_t}^{\tilde{\theta}_t} \sum_{\tau=t}^T J_\tau^t \left(\theta^{t-1}, s, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau \right) \psi' \left(e_\tau \left(\theta^{t-1}, s, \tilde{\theta}_{t+1}, \dots, \tilde{\theta}_\tau \right) \right) ds \right].$$

Depending on whether the desired effort policy satisfies the stronger single-crossing condition of (24) or the weaker single-crossing condition of (19), it can be implemented either by the quasi-linear scheme of part (1) where the agent's compensation is a convex function of the linear aggregator $\sum_{t=1}^T [S_t(\theta^t) + \alpha_t(\theta^t)\pi_t]$, or by the bonus scheme of part (2) according to which the agent receives a positive bonus only upon meeting the firm's targets in each period.

4.1 Consumption smoothing

Finally, to see how the results in the previous sections may be affected by the agent's preferences for consumption smoothing, consider an alternative setting where the agent's payoff is given by

$$U^A(c^T, e^T) = \sum_{t=1}^T \delta^{t-1} [v(c_t) - \psi(e_t)].$$

For simplicity, assume here that effort has only a transitory effect on the firm's performance, i.e. $\pi_t = \theta_t + e_t$ for all t , and that θ_t evolves according to a Markov process as in the previous section. Following the same steps used to establish Proposition 6, one can show that, in each state θ^T – equivalently, $(\theta_1, \varepsilon^T)$ – the utility of the total payment to the agent is uniquely pinned down by the policy e up to a constant $V^\Omega(\underline{\theta}_1)$. The characterization of the optimal reward scheme in this environment then proceeds as follows. Given any effort policy e , for each state θ^T , let

$$\begin{aligned} \sum_{t=1}^T \delta^{t-1} v(c_t(\theta^t)) &= \sum_{t=1}^T \delta^{t-1} \psi(e_t(\theta^t)) + V^\Omega(\underline{\theta}_1) \\ &+ \mathbb{E}_{\tilde{\theta}^T | \theta_1} \left[\int_{\underline{\theta}_1}^{\theta_1} \sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(e_t(\tilde{\theta}^t)) ds \right] \\ &+ \sum_{t=2}^T H_t(\theta^t) \end{aligned} \tag{34}$$

denote the total utility of money that is necessary to sustain the policy e , where for all $t \geq 2$, all θ^t

$$H_t(\theta^t) \equiv \mathbb{E}_{\tilde{\theta}^T | \theta^t} \left[\int_{\underline{\theta}_t}^{\theta_t} \sum_{\tau=t}^T \delta^{t-1} J_t^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) ds \right] - \mathbb{E}_{\tilde{\theta}^T | \theta^{t-1}} \left[\int_{\underline{\theta}_t}^{\tilde{\theta}_t} \sum_{\tau=t}^T \delta^{t-1} J_t^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) ds \right].$$

Then let $c^{opt}(\cdot; e)$ denote the reward scheme that minimizes the expected payment to the principal, among all schemes that satisfy conditions (34), naturally adapted to the filtration generated by the history of reports θ^t . We then have the following result.

Proposition 9 *Suppose the agent's type θ_t evolves according to a Markov process and that effort has only a transitory effect on performance so that $\pi_t = \theta_t + e_t$, all t . Let e^* denote any effort policy that maximizes*

$$\mathbb{E}_{\tilde{\theta}^T} \left[\sum_{t=1}^T \delta^{t-1} \left(\tilde{\theta}_t + e(\tilde{\theta}^t) - c_t^{opt}(\tilde{\theta}^t; e) \right) \right]$$

1. *For any t any θ^t , let $\alpha_t(\theta^t) \equiv \psi'(e_t^*(\theta^t))$. Suppose the policy e^* satisfies the single-crossing condition of (19) for any t any $(\hat{\theta}^{t-1}, \hat{\theta}_t, \theta_t)$. Then e^* together with the “quasi-linear” reward scheme s^* defined below are part of an optimal mechanism. The scheme s^* is such that, in each period t , given the reports θ^t and the observed performances π^t , the principal pays the agent a reward $s_t^*(\theta^t, \pi^t) = v^{-1}(S_t^*(\theta^t) + \alpha_t(\theta^t)\pi_t)$, where the fixed payment $S_t^*(\theta^t)$ is now implicitly defined by*

$$v^{-1}(S_t^*(\theta^t) + \alpha_t(\theta^t)\pi_t^*(\theta^t)) = c_t^{opt}(\theta^t; e^*)$$

with $\pi_t^(\theta^t) \equiv \theta_t + e_t^*(\theta^t)$ for any t any θ^t .*

2. *Suppose instead that the policy e^* does not satisfy the single-crossing condition of (19) but satisfies the single-crossing condition of (24) for any t , any $(\hat{\theta}^{t-1}, \hat{\theta}_t, \theta_t)$. Then e^* can be implemented by the following “bonus” scheme s^* : in each period t , given the reports θ^t and the observed performances π^t , the principal pays the agent a reward $s_t^*(\theta^t, \pi^t) = c_t^{opt}(\theta^t; e^*)$ if $\pi_t = \pi_t^*(\theta^t)$ and charges the agent a penalty $L > 0$ otherwise.*

Both parts (1) and (2) follow directly from the preceding results along with the definition of $c^{opt}(\cdot; e)$. The only difference between this environment and the one examined at the beginning of this section is that, while in that setting the way the total payment is distributed over time is irrelevant for the agent (and hence for the principal), in the environment considered here it is essential to distribute the payments optimally over the entire relationship. The payment schemes in Proposition 9 guarantee that the agent has the right incentives to report his information truthfully and then exert the right level of effort, while at the same time inducing the level of intertemporal consumption smoothing that maximizes the agent's utility and hence minimizes the cost for the principal.

To get a sense of how the principal allocates optimally the agent's consumption over time (which is instrumental to the characterization of the optimal effort policy) one can use Rogerson's (1985) necessary conditions for optimality. Adapted to our environment, these conditions can be stated as follows.

Proposition 10 *Suppose that the effort policy e can be implemented by the reward scheme s and let c be the corresponding consumption policy. If s implements e at minimum cost for the principal, then the following inverse Euler equation must hold for any two adjacent periods $1 \leq t, t + 1 \leq T$, almost every θ^t ,*

$$\frac{1}{v'(c_t(\theta^t))} = \mathbb{E}_{\tilde{\theta}_{t+1}|\theta^t} \left[\frac{\delta}{v'(c_{t+1}(\theta^t, \tilde{\theta}_{t+1}))} \right]. \quad (35)$$

We now show how one can calculate the optimal effort policies. The payoff equivalence result of condition (34) permits one to determine the total utility of consumption (up to a constant $V^\Omega(\underline{\theta}_1)$) that must be given to the agent in each state θ^T , for any given effort policy e . However, because the agent's payoff now depends on the timing of the payments, the principal must now span the payments optimally over time, adding an additional dimension to the problem. One way to arrive to the optimal policy e^* is the one indicated in Proposition 9; using (34) and (35) one determines the optimal payment scheme for each possible effort policy e and then chooses the policy e^* that maximizes the principal's expected payoff. An alternative route, described below, involves using the utility of consumption as an additional control and then maximizing the principal's expected payoff with respect to effort and utility of consumption, subject to (34). This alternative approach often facilitates the computation, for it does not require computation of the cost-minimizing payment scheme for each possible effort policy e .

To illustrate, suppose for simplicity that $T = 2$ and $\delta = 1$. Denote the utility of consumption in each period by $u_1(\theta_1) = v(c_1(\theta_1))$ and $u_2(\theta_1, \theta_2) = v(c_2(\theta_1, \theta_2))$. Then, for any $\theta^T = (\theta_1, \theta_2)$, equation (34), evaluated at both (θ_1, θ_2) and $(\theta_1, \underline{\theta}_2)$, allows us to express each $u_1(\theta_1)$ and $u_2(\theta_1, \theta_2)$ as functions of the effort policy, the constant $V^\Omega(\underline{\theta}_1)$, and the function $u_2(\cdot, \underline{\theta}_2)$. Specifically, for all (θ_1, θ_2) ,

$$\begin{aligned}
u_1(\theta_1) &= -u_2(\theta_1, \underline{\theta}_2) + \psi(e_2(\theta_1, \underline{\theta}_2)) + \psi(e_1(\theta_1)) + V^\Omega(\underline{\theta}_1) \\
&\quad + \mathbb{E}_{\tilde{\theta}_2|\theta_1} \left[\int_{\underline{\theta}_1}^{\theta_1} \psi'(e_1(s)) + J_1^2(s, \tilde{\theta}_2) e_2(s, \tilde{\theta}_2) ds \right] \\
&\quad - \mathbb{E}_{\tilde{\theta}_2|\theta_1} \left[\int_{\underline{\theta}_2}^{\tilde{\theta}_2} \psi'(e_2(\theta_1, s)) ds \right],
\end{aligned} \tag{36}$$

and

$$\begin{aligned}
u_2(\theta_1, \theta_2) &= u_2(\theta_1, \underline{\theta}_2) + \psi(e_2(\theta_1, \theta_2)) - \psi(e_2(\theta_1, \underline{\theta}_2)) \\
&\quad + \int_{\underline{\theta}_2}^{\theta_2} \psi'(e_2(\theta_1, s)) ds.
\end{aligned} \tag{37}$$

The optimal mechanism can then be obtained by maximizing the principal's expected payoff

$$\mathbb{E}_{(\tilde{\theta}_1, \tilde{\theta}_2)} \left[\tilde{\theta}_1 + \tilde{\theta}_2 + e_1(\tilde{\theta}_1) + e_2(\tilde{\theta}_1, \tilde{\theta}_2) - v^{-1}(u_1(\tilde{\theta}_1)) - v^{-1}(u_2(\tilde{\theta}_1, \tilde{\theta}_2)) \right],$$

with respect to $e_1(\cdot)$, $e_2(\cdot, \cdot)$, $u_2(\cdot, \underline{\theta}_2)$, and the constant $V^\Omega(\underline{\theta}_1)$, where u_1 and u_2 are given by (36) and (37).

As in previous problems, at any optimum, $V^\Omega(\underline{\theta}_1) = 0$. As an illustration, consider the same specification as in Example 5, i.e. let $v(c) = \frac{1}{\alpha} \sqrt{2\alpha c + \beta^2} - \frac{\beta}{\alpha}$ for $\alpha, \beta > 0$, and assume that $J_1^2(\theta_1, \theta_2) = \frac{1}{2}$ (an equivalent choice was made for Figures 3 and 4 in the previous subsection). Using again sixth-order polynomials, we can then compute numerically the optimal effort policies. These policies have the same qualitative features as in Example 5. In particular, both $e_1(\theta_1)$ and $\mathbb{E}_{\tilde{\theta}_1|\theta_1} [e_2(\theta_1, \tilde{\theta}_2)]$ are increasing in θ_1 , and $e_2(\theta_1, \theta_2)$ is decreasing in θ_2 (because of the similarity, the figures for the present case are not displayed).

5 Appendix

Proof of Proposition 1. Because the agent's participation constraint clearly binds at the optimum, the principal's payoff coincides with the total surplus generated by the relationship, which is given by

$$W = \sum_{t=1}^T \delta^{t-1} [\pi_t(\theta_t, e^t) - \psi(e_t)] = \sum_{t=1}^T \delta^{t-1} \left[\theta_t + e_t + \sum_{\tau=1}^{t-1} \rho^\tau e_{t-\tau} - \psi(e_t) \right]$$

The result then follows from pointwise maximization of $\mathbb{E}[W]$ with respect to each $e_t(\theta_1, \varepsilon^t)$. ■

Proof of Example 4. First note that this environment satisfies all the conditions on the kernels F assumed at the beginning of the section and that, for any $t \geq 2$, any $\theta_t \in \Theta_t = \mathbb{R}_+$, any θ_{t-1} ,

$$F_t(\theta_t | \theta_{t-1}) = G\left(\frac{\theta_t}{\theta_{t-1}}\right).$$

Therefore, for each $k \geq 2$, each (θ_k, θ_{k-1}) ,

$$I_{k-1}^k(\theta_k, \theta_{k-1}) = \frac{-\partial G\left(\frac{\theta_k}{\theta_{k-1}}\right) / \partial \theta_{k-1}}{\frac{1}{\theta_{k-1}} g\left(\frac{\theta_k}{\theta_{k-1}}\right)} = \frac{\theta_k}{\theta_{k-1}}.$$

It follows that, for each $\tau > t$,

$$J_t^\tau(\theta_t, \dots, \theta_\tau) = \frac{\theta_\tau}{\theta_t}.$$

In each period t , the effort policy that solves the relaxed program is thus given by

$$e_t^*(\theta^t) = \max\left\{1 - \eta(\theta_1) \frac{\theta_t}{\theta_1}; 0\right\}.$$

It is immediate to see that, because $e_t^*(\theta^t)$ is decreasing in θ_t , it violates condition (19). Furthermore, by taking e.g. $t = T$, one can easily see that, for any $\hat{\theta}_1$ and any $\hat{\theta}_T < \theta_T < \hat{\theta}_1 / \eta(\hat{\theta}_1)$,

$$u_T(\theta_T, \hat{\theta}_T; \hat{\theta}^{T-1}) > u_T(\theta_T; \hat{\theta}^{T-1}) \equiv u_T(\theta_T, \theta_T; \hat{\theta}^{T-1}),$$

where

$$u_T(\theta_T; \hat{\theta}^{T-1}) = \int_0^{\theta_T} \psi'(e_T^*(\hat{\theta}^{T-1}, s)) ds = \int_0^{\theta_T} \left[1 - \eta(\hat{\theta}_1) \frac{s}{\hat{\theta}_1}\right] ds$$

and

$$\begin{aligned} u_T(\theta_T, \hat{\theta}_T; \hat{\theta}^{T-1}) &= \int_{\hat{\theta}_T}^{\hat{\theta}_T} \psi'(e_T^*(\hat{\theta}^{T-1}, s)) ds + \alpha_T(\hat{\theta}^{T-1}, \hat{\theta}_T)[\theta_T - \hat{\theta}_T] \\ &= \int_0^{\hat{\theta}_T} \left[1 - \eta(\hat{\theta}_1) \frac{s}{\hat{\theta}_1}\right] ds + \left[1 - \eta(\hat{\theta}_1) \frac{\hat{\theta}_T}{\hat{\theta}_1}\right] [\theta_T - \hat{\theta}_T] \end{aligned}$$

are, respectively, the continuation payoff that type θ_T obtains by reporting θ_T truthfully and the continuation payoff he obtains by reporting $\hat{\theta}_T < \theta_T$, under the linear scheme s^* of Proposition 3.

This proves that e^* cannot be implemented with the linear scheme.

Finally, to see that e^* can be implemented by the scheme s of Proposition 4, it suffices to show that the single-crossing condition (24) holds for any t , any $(\hat{\theta}^{t-1}, \hat{\theta}_t, \theta_t)$. To see this, note that, when $t = 1$, (24) is equivalent to

$$\mathbb{E}_{\tilde{\theta}_2, \dots, \tilde{\theta}_T | \theta_1} \left[\max \{1 - \eta(\theta_1); 0\} - \max \left\{ \max \{1 - \eta(\hat{\theta}_1); 0\} + \hat{\theta}_1 - \theta_1; 0 \right\} \right. \\ \left. + \sum_{t=2}^T \delta^{t-1} \frac{\tilde{\theta}_t}{\hat{\theta}_1} \left[\max \{1 - \eta(\theta_1) \frac{\tilde{\theta}_t}{\hat{\theta}_1}; 0\} - \max \{1 - \eta(\hat{\theta}_1) \frac{\tilde{\theta}_t}{\hat{\theta}_1}; 0\} \right] \right] [\theta_1 - \hat{\theta}_1] \geq 0.$$

This holds by the assumption that η is non-increasing. Now take any period $t > 1$. Condition (24) then requires that

$$\left[\max \left\{ 1 - \eta(\hat{\theta}_1) \frac{\theta_t}{\hat{\theta}_1}; 0 \right\} - \max \left\{ \max \left\{ 1 - \eta(\hat{\theta}_1) \frac{\hat{\theta}_t}{\hat{\theta}_1}; 0 \right\} + \hat{\theta}_1 - \theta_1; 0 \right\} \right] [\theta_t - \hat{\theta}_t] \geq 0$$

which holds because, by assumption, $\eta(\hat{\theta}_1) \leq \hat{\theta}_1$ for each $\hat{\theta}_1$. ■

Proof of Proposition 5. Let $\mathcal{W} = \mathcal{V}^{-1}$. The principal maximizes by choice of functions $(\hat{e}_t(\cdot))_{t=1}^T$

$$\mathbb{E} \left[\sum_{t=1}^T \left[\pi_t - \mathcal{W} \left(\sum_{t=1}^T \psi(\hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t)) \right) \right] \right] \\ = \mathbb{E} \left[\sum_{t=1}^T \left[z_t(\tilde{\theta}_1, \tilde{\varepsilon}_t) + \hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t) - \mathcal{W} \left(\sum_{t=1}^T \psi(\hat{e}_t(\tilde{\theta}_1, \tilde{\varepsilon}^t)) \right) \right] \right].$$

It is immediate by the convexity of \mathcal{W} and ψ that non-contingent effort policies \hat{e}_t are (almost- uniquely) optimal. For each t , the first-order condition for optimal effort is then

$$\mathcal{W}' \left(\sum_{\tau=1}^T \psi(\hat{e}_\tau) \right) \psi'(\hat{e}_t) = 1,$$

which is sufficient for an optimum by convexity of \mathcal{W} and ψ . Applying the inverse-function theorem, we have $\hat{e}_t = \hat{e}^{FB}$ for all t , where \hat{e}^{FB} is given by (25). ■

Proof of Proposition 6. We prove the result by backward induction, starting from $t = T$. Using the characterization of the necessary conditions for incentive compatibility in the main text, we have that (to ease the exposition, hereafter we drop the qualification “for almost every truthful history”):

$$V^{\hat{\Omega}}(\theta_1, \varepsilon^T) = V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \underline{\varepsilon}_T) + \int_{\underline{\varepsilon}_T}^{\varepsilon^T} D_{T,T}(\theta_1, \varepsilon^{T-1}, x) \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-1}, x)) dx$$

Also,

$$V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}) = V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-2}, \underline{\varepsilon}_{T-1}) + \mathbb{E}_{\tilde{\varepsilon}_T} \left[\int_{\underline{\varepsilon}_{T-1}}^{\varepsilon^{T-1}} \left(\begin{array}{c} D_{T-1, T-1}(\theta_1, \varepsilon^{T-2}, x) \psi'(\hat{e}_{T-1}(\theta_1, \varepsilon^{T-2}, x)) \\ + D_{T-1, T}(\theta_1, \varepsilon^{T-2}, x, \tilde{\varepsilon}_T) \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-2}, x, \tilde{\varepsilon}_T)) \end{array} \right) dx \right]$$

Using the fact that

$$V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}) = \mathbb{E}_{\tilde{\varepsilon}_T} \left[V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \tilde{\varepsilon}_T) \right]$$

we then have that

$$\begin{aligned} V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \underline{\varepsilon}_T) &= \mathbb{E}_{\tilde{\varepsilon}_T} \left[V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}, \tilde{\varepsilon}_T) - \int_{\underline{\varepsilon}_T}^{\tilde{\varepsilon}_T} D_{T, T}(\theta_1, \varepsilon^{T-1}, x) \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-1}, x)) dx \right] \\ &= V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-1}) - \mathbb{E}_{\tilde{\varepsilon}_T} \left[\int_{\underline{\varepsilon}_T}^{\tilde{\varepsilon}_T} D_{T, T}(\theta_1, \varepsilon^{T-1}, x) \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-1}, x)) dx \right] \\ &= V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-2}, \underline{\varepsilon}_{T-1}) \\ &\quad + \mathbb{E}_{\tilde{\varepsilon}_T} \left[\int_{\underline{\varepsilon}_{T-1}}^{\varepsilon^{T-1}} \left(\begin{array}{c} D_{T-1, T-1}(\theta_1, \varepsilon^{T-2}, x) \psi'(\hat{e}_{T-1}(\theta_1, \varepsilon^{T-2}, x)) \\ + D_{T-1, T}(\theta_1, \varepsilon^{T-2}, x, \tilde{\varepsilon}_T) \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-2}, x, \tilde{\varepsilon}_T)) \end{array} \right) dx \right] \\ &\quad - \mathbb{E}_{\tilde{\varepsilon}_T} \left[\int_{\underline{\varepsilon}_T}^{\tilde{\varepsilon}_T} D_{T, T}(\theta_1, \varepsilon^{T-1}, x) \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-1}, x)) dx \right] \end{aligned}$$

Therefore,

$$\begin{aligned} V^{\hat{\Omega}}(\theta_1, \varepsilon^T) &= V^{\hat{\Omega}}(\theta_1, \varepsilon^{T-2}, \underline{\varepsilon}_{T-1}) \\ &\quad + \mathbb{E}_{\tilde{\varepsilon}_T} \left[\int_{\underline{\varepsilon}_{T-1}}^{\varepsilon^{T-1}} \left(\begin{array}{c} D_{T-1, T-1}(\theta_1, \varepsilon^{T-2}, x) \psi'(\hat{e}_{T-1}(\theta_1, \varepsilon^{T-2}, x)) \\ + D_{T-1, T}(\theta_1, \varepsilon^{T-2}, x, \tilde{\varepsilon}_T) \psi'(\hat{e}_T(\theta_1, \varepsilon^{T-2}, x, \tilde{\varepsilon}_T)) \end{array} \right) dx \right] \\ &\quad + \hat{H}_T(\theta_1, \varepsilon^T) \end{aligned}$$

This establishes the first step of the induction. Now suppose that there exists a $t \leq T-1$ such that the following representation holds for all periods $s, t \leq s < T$:

$$\begin{aligned} V^{\hat{\Omega}}(\theta_1, \varepsilon^T) &= V^{\hat{\Omega}}(\theta_1, \varepsilon^{s-1}, \underline{\varepsilon}_s) \tag{38} \\ &\quad + \mathbb{E}_{(\tilde{\varepsilon}_{s+1}, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_s}^{\varepsilon_s} \sum_{\tau=s}^T D_{s, \tau}(\theta_1, \varepsilon^{s-1}, x, \tilde{\varepsilon}_{s+1}, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{s-1}, x, \tilde{\varepsilon}_{s+1}, \dots, \tilde{\varepsilon}_\tau)) dx \right] \\ &\quad + \sum_{\tau=s+1}^T \hat{H}_\tau(\theta_1, \varepsilon^\tau) \end{aligned}$$

We then want to show that it holds also for $s = t - 1$. Note that, by incentive compatibility,

$$\begin{aligned}
V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-1}, \underline{\varepsilon}_t) &= \mathbb{E}_{\tilde{\varepsilon}_t} \left[V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-1}, \tilde{\varepsilon}_t) \right] \\
&\quad - \mathbb{E}_{(\tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_t}^{\tilde{\varepsilon}_t} \sum_{\tau=t}^T D_{t,\tau}(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau)) dx \right] \\
&= V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-1}) \\
&\quad - \mathbb{E}_{(\tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_t}^{\tilde{\varepsilon}_t} \sum_{\tau=t}^T D_{t,\tau}(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau)) dx \right]
\end{aligned}$$

and that, again by incentive compatibility,

$$\begin{aligned}
V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-1}) &= V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-2}, \underline{\varepsilon}_{t-1}) \\
&\quad + \mathbb{E}_{(\tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_{t-1}}^{\varepsilon_{t-1}} \sum_{\tau=t-1}^T D_{t-1,\tau}(\theta_1, \varepsilon^{t-2}, x, \tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-2}, x, \tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_\tau)) dx \right].
\end{aligned}$$

Therefore,

$$\begin{aligned}
V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-1}, \underline{\varepsilon}_t) &= V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-2}, \underline{\varepsilon}_{t-1}) & (39) \\
&\quad + \mathbb{E}_{(\tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_{t-1}}^{\varepsilon_{t-1}} \sum_{\tau=t-1}^T D_{t-1,\tau}(\theta_1, \varepsilon^{t-2}, x, \tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-2}, x, \tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_\tau)) dx \right] \\
&\quad - \mathbb{E}_{(\tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_t}^{\tilde{\varepsilon}_t} \sum_{\tau=t}^T D_{t,\tau}(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau)) dx \right]
\end{aligned}$$

Using (38) for $s = t$ and combining it with (39), we then have that

$$\begin{aligned}
V^{\hat{\Omega}}(\theta_1, \varepsilon^T) &= V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-2}, \underline{\varepsilon}_{t-1}) \\
&+ \mathbb{E}_{(\tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_{t-1}}^{\varepsilon_{t-1}} \sum_{\tau=t-1}^T D_{t-1, \tau}(\theta_1, \varepsilon^{t-2}, x, \tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-2}, x, \tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_\tau)) ds \right] \\
&- \mathbb{E}_{(\tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_t}^{\tilde{\varepsilon}_t} \sum_{\tau=t}^T D_{t, \tau}(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau)) dx \right] \\
&+ \mathbb{E}_{(\tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_t}^{\varepsilon_t} \sum_{\tau=t}^T D_{t, \tau}(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-1}, x, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_\tau)) dx \right] \\
&+ \sum_{\tau=t+1}^T \hat{H}_\tau(\theta_1, \varepsilon^\tau) \\
&= V^{\hat{\Omega}}(\theta_1, \varepsilon^{t-2}, \underline{\varepsilon}_{t-1}) \\
&+ \mathbb{E}_{(\tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_T)} \left[\int_{\underline{\varepsilon}_{t-1}}^{\varepsilon_{t-1}} \sum_{\tau=t-1}^T D_{t-1, \tau}(\theta_1, \varepsilon^{t-2}, x, \tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_\tau) \psi'(\hat{e}_\tau(\theta_1, \varepsilon^{t-2}, x, \tilde{\varepsilon}_t, \dots, \tilde{\varepsilon}_\tau)) ds \right] \\
&+ \sum_{\tau=t}^T \hat{H}_\tau(\theta_1, \varepsilon^\tau)
\end{aligned}$$

which proves that the representation in (38) holds also for $s = t - 1$. The result then follows directly from the fact that the agent's payoff under truth-telling must coincide with the value function almost surely. ■

Proof of Example 5. Using (26), the principal's expected payoff in any incentive-compatible mechanism $\hat{\Omega}$ implementing a policy \hat{e} that is contingent on θ_1 only (i.e. such that there exists a sequence of functions $\hat{e}_t : \Theta_1 \rightarrow \mathbb{R}$, $t = 2, \dots, T$, such that $\hat{e}_t(\theta_1, \varepsilon^t) = \hat{e}_t(\theta_1)$ for all ε^t) is given by $\mathbb{E}[W(\tilde{\theta}_1)]$, where, for any $\theta_1 \in \Theta_1$,

$$\begin{aligned}
W(\theta_1) &\equiv \mathbb{E} \left[\sum_{s=1}^T \left(D_{1,s} \theta_1 + \sum_{\tau=2}^s D_{\tau,s} \tilde{\varepsilon}_\tau \right) \right] + \sum_{s=1}^T \hat{e}_s(\theta_1) \\
&- \frac{\alpha}{2} \mathbb{E} \left[\left(\begin{aligned} &\sum_{s=1}^T \psi(\hat{e}_s(\theta_1)) + V^{\hat{\Omega}}(\theta_1) \\ &+ \int_{\underline{\theta}_1}^{\theta_1} \sum_{s=1}^T D_{1,s} \psi'(\hat{e}_s(z)) dz + \sum_{s=2}^T \hat{H}_s(\theta_1, \tilde{\varepsilon}^s) \end{aligned} \right)^2 \right] \\
&- \beta \mathbb{E} \left[\begin{aligned} &\sum_{s=1}^T \psi(\hat{e}_s(\theta_1)) + V^{\hat{\Omega}}(\theta_1) \\ &+ \int_{\underline{\theta}_1}^{\theta_1} \sum_{s=1}^T D_{1,s} \psi'(\hat{e}_s(\tilde{\theta}_1)) d\tilde{\theta}_1 + \sum_{s=2}^T \hat{H}_s(\theta_1, \tilde{\varepsilon}^s) \end{aligned} \right],
\end{aligned}$$

where, for $s = 2, \dots, T$,

$$\hat{H}_s(\theta_1, \varepsilon^s) = [\varepsilon_s - \mathbb{E}[\tilde{\varepsilon}_s]] \sum_{\tau=s}^T D_{s,\tau} \psi'(\hat{e}_\tau(\theta_1)).$$

For each $s = 2, \dots, T$, let $\sigma_s^2 \equiv \text{Var}(\tilde{\varepsilon}_s)$. Then

$$\begin{aligned} W(\theta_1) &= \sum_{s=1}^T \left(D_{1,s} \theta_1 + \sum_{\tau=2}^s D_{\tau,s} \mathbb{E}[\tilde{\varepsilon}_\tau] \right) + \sum_{s=1}^T \hat{e}_s(\theta_1) \\ &\quad - \frac{\alpha}{2} \left(\sum_{s=1}^T \psi(\hat{e}_s(\theta_1)) + V^{\hat{\Omega}}(\underline{\theta}_1) + \int_{\underline{\theta}_1}^{\theta_1} \sum_{s=1}^T D_{1,s} \psi'(\hat{e}_s(z)) dz \right)^2 \\ &\quad - \frac{\alpha}{2} \sum_{s=2}^T \sigma_s^2 \left[\sum_{\tau=s}^T D_{s,\tau} \psi'(\hat{e}_\tau(\theta_1)) \right]^2 \\ &\quad - \beta \left(\sum_{s=1}^T \psi(\hat{e}_s(\theta_1)) + V^{\hat{\Omega}}(\underline{\theta}_1) + \int_{\underline{\theta}_1}^{\theta_1} \sum_{s=1}^T D_{1,s} \psi'(\hat{e}_s(z)) dz \right). \end{aligned}$$

The principal's relaxed program then consists of choosing a vector of effort functions $\hat{e}_t : \Theta_1 \rightarrow \mathbb{R}$, $t = 1, \dots, T$, along with a scalar $V^{\hat{\Omega}}(\underline{\theta}_1) \geq 0$, so as to maximize $\mathbb{E}[W(\tilde{\theta}_1)]$. It is immediate that, at the optimum, $V^{\hat{\Omega}}(\underline{\theta}_1) = 0$. Furthermore, given that $\psi'(\bar{e}) > 1/\beta$, it is also immediate that any policy $\hat{e} = (\hat{e}_t(\cdot))_{t=1}^T$ that maximizes $\mathbb{E}[W(\tilde{\theta}_1)]$ must have the property that, for any t , $\hat{e}_t(\theta_1) \in [0, \bar{e}]$ for almost every $\theta_1 \in \Theta_1$.

Now let $g : [0, K] \rightarrow \mathbb{R}$ be the function defined by $g(0) = 0$, $g(y) = \psi'^{-1}(y)$, all $y \in (0, K)$, and $g(K) = \bar{e}$. For any $t = 1, \dots, T$, any $\theta_1 \in \Theta_1$, then let $u_t(\theta_1) \equiv \psi'(\hat{e}_t(\theta_1))$ and $x_t(\theta_1) \equiv \int_{\underline{\theta}_1}^{\theta_1} u_t(z) dz$. Omitting the first term, which does not depend on the effort policy, the principal's relaxed problem thus consists in choosing functions $u : \Theta_1 \rightarrow [0, K]^T$ and $x : \Theta_1 \rightarrow \mathbb{R}_+^T$ that maximize

$$\int_{\underline{\theta}_1}^{\bar{\theta}_1} L(\theta_1, u(\theta_1), x(\theta_1)) d\theta_1$$

where, for any $(\theta_1, u, x) \in \Theta_1 \times [0, K]^T \times \mathbb{R}_+^T$,

$$L(\theta_1, u, x) \equiv f(\theta_1) \left[\begin{array}{c} \sum_{s=1}^T g(u_s) - \frac{\alpha}{2} \left(\sum_{s=1}^T \psi(g(u_s)) + \sum_{s=1}^T D_{1,s} x_s \right)^2 \\ - \frac{\alpha}{2} \sum_{s=2}^T \sigma_s^2 \left(\sum_{\tau=s}^T D_{s,\tau} u_\tau \right)^2 - \beta \left(\sum_{s=1}^T \psi(g(u_s)) + \sum_{s=1}^T D_{1,s} x_s \right) \end{array} \right]$$

under the constraint that

$$x_t(\theta_1) = \int_{\underline{\theta}_1}^{\theta_1} u_t(z) dz, \quad \forall t = 1, \dots, T, \quad \forall \theta_1 \in \Theta_1. \quad (40)$$

We solve this problem with optimal control treating u as the vector of control variables and x as the vector of state variables. First we verify that a solution to this optimal control problem exists by applying the Tonelli existence theorem.³³ To this aim, we first show that, for any $(\theta_1, x) \in \Theta_1 \times \mathbb{R}_+^T$, the function $L(\theta_1, \cdot, x)$ is strictly concave. Note that, for any $y \in (0, K)$,

$$g''(y) = \frac{d}{dy} \left[\frac{1}{\psi''(g(y))} \right] = \frac{-\psi'''(g(y))}{[\psi''(g(y))]^3} < 0.$$

This implies that $\sum_{s=1}^T g(u_s)$ is strictly concave in u . Next, note that, for any $y \in (0, K)$,

$$\frac{d^2}{du^2} \psi(g(y)) = \frac{d}{du} \left[\frac{\psi'(g(y))}{\psi''(g(y))} \right] = \frac{h'(g(y))}{\psi''(g(y))}$$

where $h(z) \equiv \frac{\psi'(z)}{\psi''(z)}$, and so $h'(z) \equiv \frac{d}{dz} \left(\frac{\psi'(z)}{\psi''(z)} \right) = \frac{d}{dz} \left(\left\{ \frac{d \log \psi'(z)}{dz} \right\}^{-1} \right) > 0$, since $\log \psi'$ is concave.

Hence $\psi(g(y))$ is convex. Therefore, $\sum_{s=1}^T \psi(g(u_s))$ is convex in u . Moreover, since $(\cdot)^2$ is convex and is increasing whenever its argument is non-negative, $-\left(\sum_{s=1}^T \psi(g(u_s)) + \sum_{s=1}^T D_{1,s} x_s \right)^2$ is concave in u as well. The same argument implies that, for any s , the function $-\left(\sum_{\tau=s}^T D_{s,\tau} u_\tau \right)^2$ is concave in u so that $-\frac{\alpha}{2} \sum_{s=2}^T \sigma_s^2 \left(\sum_{\tau=s}^T D_{s,\tau} u_\tau \right)^2$ is weakly concave. Together these observations imply that L is strictly concave in u , as required. Moreover, L is continuous.³⁴ Finally, that u is bounded renders the ‘‘coercivity’’ condition of Tonelli’s theorem unnecessary.³⁵ Thus a solution exists. Finally, that L is weakly jointly concave in (u, x) and strictly concave in u implies that the solution is essentially unique.

By Proposition 2.1 of Clarke (1989), that u is bounded and that L is strictly concave in u for each x , then guarantees that each $x_t(\cdot)$ is continuously differentiable and that the Pontryagin principle applies.

The Hamiltonian function is given by

$$\mathcal{H} = L(\theta_1, u(\theta_1), x(\theta_1)) + \mu(\theta_1)^\top u(\theta_1).$$

³³See, for example, Theorem 3.7 of Buttazzo, Giaquinta and Hildebrandt (1998).

³⁴All that is required for a Tonelli-type existence theorem is that L be measurable in θ_1 for all admissible x and $u \in [0, K]^T$, and continuous in (x, u) for almost every θ_1 . See Theorem 3.6 of Buttazzo, Giaquinta and Hildebrandt (1998).

³⁵See, for instance, Theorem 3.7 of Buttazzo, Giaquinta and Hildebrandt (1998). The role of the coercivity condition in Tonelli’s result is exactly to guarantee that the controls u are essentially bounded.

where μ is the vector of co-state variables associated with the law of motions given by

$$\dot{x}(\theta_1) = u(\theta_1) \quad \text{a.e. } \theta_1. \quad (41)$$

Pointwise maximization of the Hamiltonian then requires that for almost every $\theta_1 \in \Theta_1$, all $t = 1, \dots, T$,³⁶

$$f(\theta_1) \left[\begin{array}{c} g'(u_t(\theta_1)) - \psi'(g(u_t(\theta_1))) g'(u_t(\theta_1)) \left[\alpha \left(\sum_{s=1}^T \psi(g(u_s(\theta_1))) + \sum_{s=1}^T D_{1,s} x_s(\theta_1) \right) + \beta \right] \\ -\alpha \sum_{s=2}^t \left[\left(\sum_{\tau=s}^T D_{s,\tau} u_\tau(\theta_1) \right) D_{s,t} \sigma_s^2 \right] \end{array} \right] + \mu_t(\theta_1) \leq 0, \quad (42)$$

with inequality satisfied as equality if $u_t(\theta_1) > 0$. Furthermore, for almost every $\theta_1 \in \Theta_1$, any $t = 1, \dots, T$, the adjoint equations

$$\dot{\mu}_t(\theta_1) = f(\theta_1) D_{1,t} \left[\alpha \left(\sum_{s=1}^T \psi(g(u_s(\theta_1))) + \sum_{s=1}^T D_{1,s} x_s(\theta_1) \right) + \beta \right] \quad (43)$$

must hold. Finally, the following boundary and transversality conditions must be satisfied:

$$\mu_t(\bar{\theta}_1) = x_t(\theta_1) = 0, \quad t = 1, \dots, T. \quad (44)$$

Combining together (42)-(44), and using absolute continuity of the co-state variables, gives (30).
■

Proof of Proposition 10. First note that if s implements e , then e can also be implemented by the following “bonus” scheme:

$$\hat{s}_t(\theta^t, \pi^t) = \begin{cases} c_t(\theta^t) & \text{if } \pi_t = \theta_t + e_t(\theta^t) \\ -L & \text{otherwise} \end{cases} \quad (45)$$

with $L > 0$ arbitrarily large. So, without loss, assume s itself satisfies condition (45).

Now, suppose there exists a period t and a (positive measure) set $Q \subset \Theta^t$ such that, for any $\theta^t \in \Theta^t$,

$$\frac{1}{v'(c_t(\theta^t))} > \mathbb{E}_{\tilde{\theta}_{t+1}|\theta^t} \left[\frac{\delta}{v'(c_{t+1}(\theta^t, \tilde{\theta}_{t+1}))} \right].$$

³⁶We abstract from the constraints that $u_t \leq K$. It is in fact immediate from the fact that $K = \psi'(\bar{e}) > 1/\beta$ that these constraints never bind.

The argument for the case where the inequality is reversed is symmetric. Then consider the following alternative scheme $s^\#$. For any $\tau \neq t, t+1$, any (θ^τ, π^τ) , $s_\tau^\#(\theta^\tau, \pi^\tau) = s_\tau(\theta^\tau, \pi^\tau)$; in period t ,

$$s_t^\#(\theta^t, \pi^t) = \begin{cases} s_t(\theta^t, \pi^t) & \text{if } \theta^t \notin Q \\ v^{-1}(v(c_t(\theta^t)) + k) & \text{if } \theta^t \in Q \text{ and } \pi_t = \theta_t + e_t(\theta^t) \text{ ;} \\ -L & \text{if } \theta^t \in Q \text{ and } \pi_t \neq \theta_t + e_t(\theta^t) \end{cases}$$

and in period $t+1$,

$$s_{t+1}^\#(\theta^{t+1}, \pi^{t+1}) = \begin{cases} s_{t+1}(\theta^{t+1}, \pi^{t+1}) & \text{if } \theta^t \notin Q \\ v^{-1}(v(c_{t+1}(\theta^{t+1})) - k/\delta) & \text{if } \theta^t \in Q \text{ and } \pi_{t+1} = \theta_{t+1} + e_{t+1}(\theta^{t+1}) \text{ .} \\ -L & \text{if } \theta^t \in Q \text{ and } \pi_{t+1} \neq \theta_{t+1} + e_{t+1}(\theta^{t+1}) \end{cases}$$

Clearly, this scheme preserves incentives for both truthful revelation and obedience and, in equilibrium, gives the agent the same payoff as the original scheme s .³⁷ The difference between the ex-ante expected cost to the principal under this scheme and under the original scheme s is given by

$$\Delta(k) \equiv F(\tilde{\theta}^t \in Q) \mathbb{E}_{\{\tilde{\theta}^t: \tilde{\theta}^t \in Q\}} \left[\begin{array}{c} \delta^{t-1} \left(v^{-1} \left(v \left(c_t \left(\tilde{\theta}^t \right) \right) + k \right) - c_t \left(\tilde{\theta}^t \right) \right) \\ + \delta^t \mathbb{E}_{\tilde{\theta}_{t+1} | \tilde{\theta}^t} \left[v^{-1} \left(v \left(c_{t+1} \left(\tilde{\theta}^{t+1} \right) \right) - \frac{k}{\delta} \right) - c_{t+1} \left(\tilde{\theta}^{t+1} \right) \right] \end{array} \right]$$

where $F(\tilde{\theta}^t \in Q)$ denotes the ex-ante probability that $\tilde{\theta}^t \in Q$ and $\mathbb{E}_{\{\tilde{\theta}^t: \tilde{\theta}^t \in Q\}}[\cdot]$ denotes the conditional expectation of $[\cdot]$ over Θ^t given the sigma-algebra generated by the event that $\tilde{\theta}^t \in Q$.

Clearly, $\Delta(0) = 0$ and

$$\frac{\partial \Delta(0)}{\partial k} = F(\tilde{\theta}^t \in Q) \mathbb{E}_{\{\tilde{\theta}^t: \tilde{\theta}^t \in Q\}} \delta^{t-1} \left[\frac{1}{v'(c_t(\tilde{\theta}^t))} + \mathbb{E}_{\tilde{\theta}_{t+1} | \tilde{\theta}^t} \left(\frac{\delta}{v'(c_{t+1}(\tilde{\theta}^t, \tilde{\theta}_{t+1}))} \right) \right] > 0$$

The principal can then reduce her expected payment to the agent by switching to a scheme $s^\#$ with $k < 0$ arbitrarily small, contradicting the assumption that s is cost-minimizing. ■

References

- [1] Albanesi, Stefania and Christopher Sleet (2006), ‘Dynamic Optimal Taxation with Private Information’, *Review of Economic Studies*, 73, pp 1-30.
- [2] Baron, D. P., and D. Besanko (1984): ‘Regulation and Information in a Continuing Relationship,’ *Information Economics and Policy*, 1(3), 267–302.

³⁷Since the choice of L in the scheme s was arbitrary, it may be chosen large enough that incentives are still preserved in $s^\#$.

- [3] Battaglini, Marco (2005). ‘Long-Term Contracting with Markovian Consumers’, *American Economic Review*, 95(3), 637-658.
- [4] Battaglini, Marco and Stephen Coate (2008), ‘Pareto efficient income taxation with stochastic abilities’, *Journal of Public Economics* 92, 844-868.
- [5] Bebchuk, Lucian A. and Jesse Fried (2004). *Pay without Performance: The Unfulfilled Promise of Executive Compensation*, Harvard University Press.
- [6] Besanko, David (1985), ‘Multi-Period Contracts between Principal and Agent with Adverse Selection’, *Economics Letters*, 17, 33-37.
- [7] Clarke, Frank (1989), *Methods of Dynamic and Nonsmooth Optimization*, Society for Industrial and Applied Mathematics, Philadelphia, PA.
- [8] DeMarzo, Peter and Michael Fishman (2007), ‘Optimal Long-Term Financial Contracting’, *Review of Financial Studies*, 20, 2079-2127.
- [9] De Marzo, Peter and Yuliy Sannikov (2006), ‘Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model’ *Journal of Finance* 61, 2681-2724.
- [10] DeMarzo, Peter and Yuliy Sannikov (2008), ‘Learning in Dynamic Incentive Contracts’, mimeo, Stanford University.
- [11] Dewatripont, Mathias, Ian Jewitt and Tirole, Jean (1999). "The Economics of Career Concerns, Part I: Comparing Information Structures," *Review of Economic Studies* 66(1), 183-98, January.
- [12] Edmans, Alex and Xavier Gabaix (2008), ‘Tractability and Detail-Neutrality in Incentive Contracting’, mimeo Stern, New York University and Wharton, University of Pennsylvania.
- [13] Edmans, Alex and Xavier Gabaix (2009), ‘Is CEO Pay Really Inefficient? A Survey of New Optimal Contracting Theories’ *European Financial Management* (forthcoming).

- [14] Garcia, D. (2005), ‘Monotonicity in Direct Revelation Mechanisms,’ *Economics Letters*, 88(1), 21–26.
- [15] Gibbons, Robert, and Kevin J. Murphy (1991) ‘Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence’, *NBER Working Paper*, No. 3792.
- [16] Hill, C.W.L. and P. Phan (1991), ‘CEO Tenure as a Determinant of CEO Pay’, *Academy of Management Journal* 34(3), 707-717.
- [17] Holmstrom, Bengt and Paul Milgrom (1987), ‘Aggregation and Linearity in the Provision of Intertemporal Incentives’, *Econometrica* 55(2), 303-328.
- [18] Kuhnen, Camelia and Jeffrey Zwiebel (2008), ‘Executive Pay, Hidden Compensation and Managerial Entrenchment’, *Rock Center for Corporate Governance Working Paper* No. 16.
- [19] Laffont, Jean-Jacques and Jean Tirole (1986), ‘Using Cost Observation to Regulate Firms’, *Journal of Political Economy*, 94, 614-641.
- [20] Laffont, Jean-Jacques and Jean Tirole (1988), ‘The Dynamics of Incentive Contracts’, *Econometrica*, 56, 1153-1175.
- [21] Laffont, Jean-Jacques and Jean Tirole (1990), ‘Adverse Selection and Renegotiation in Procurement’, *Review of Economic Studies*, 57(4), 597-625.
- [22] Lambert, Richard A. (1983), ‘Long-term contracts and moral hazard’, *Bell Journal of Economics*, 14(2), 441-452.
- [23] Lippert, Robert L. and Gayle Porter (1997), ‘Understanding CEO Pay: A Test of Two Pay-to-Performance Sensitivity Measures with Alternative Measures of Alignment and Influence’, *Journal of Business Research*, 40, 127-138.
- [24] Milgrom, Paul and Chris Shannon (1994), ‘Monotone Comparative Statics’, *Econometrica*, 62(1), 157-180.

- [25] Pavan, Alessandro, Ilya Segal and Juuso Toikka, (2009,a) ‘Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization and Information Disclosure’, mimeo Northwestern University and Stanford University.
- [26] Pavan, Alessandro, Ilya Segal and Juuso Toikka, (2009,b) ‘Infinite-Horizon Mechanism Design: The Independent-Shock Approach,’ mimeo Northwestern University and Stanford University.
- [27] Phelan, Christopher and Robert M. Townsend (1991), ‘Computing Multi-Period, Information-Constrained Optima’, *Review of Economic Studies* 58, 853-881.
- [28] Rogerson, William P. (1985), ‘Repeated Moral Hazard’, *Econometrica* 53(1), 69-76.
- [29] Salanie, Bernard (1990), ‘Selection adverse et aversion pour le risque’, *Annales d’Economie et de Statistique*, No. 18, pp 131-149.
- [30] Sannikov, Yuliy (2007), ‘Agency Problems, Screening and Increasing Credit Lines,’ mimeo Berkeley University.
- [31] Sannikov, Yuliy (2008), ‘A Continuous-Time Version of the Principal-Agent Problem’, *Review of Economic Studies*, 75(3), 957-984.
- [32] Spear, Stephen E. and Sanjay Srivastava (1987) ‘On Repeated Moral Hazard with Discounting’, *Review of Economic Studies*, 54(4), 599-617.
- [33] Tchisty, Alexei (2006), ‘Security Design with Correlated Hidden Cash Flows: The Optimality of Performance Pricing’, mimeo Stern, NYU.
- [34] Topkis, Donald (1998), *Supermodularity and Complementarity*, Princeton University Press, Frontiers of Economic Research.
- [35] Zhang, Yuzhe (2009), ‘Dynamic contracting with persistent shocks’, *Journal of Economic Theory* 144, 635-675.