

Dynamic Managerial Compensation: On the Optimality of Seniority-based Schemes*

Daniel Garrett

Alessandro Pavan

Northwestern University

Northwestern University

October 2011

Abstract

We characterize the firm's optimal contract for a manager who faces costly effort decisions and whose ability to generate profits for the firm changes stochastically over time. The optimal contract is obtained as the solution to a dynamic mechanism design problem with hidden actions and persistent private shocks to the manager's productivity. When the manager is risk-neutral, the optimal contract often entails a simple pay package that is linear in the firm's cash flows. Furthermore, the power of incentives (i.e., the sensitivity of pay to performance) typically increases over time, thus providing a possible justification for the practice of putting more stocks and options in the pay packages of managers with longer tenure in the firm. Building on the insights from the risk-neutral case, we then explore the properties of optimal contracts for risk-averse managers. We find that risk-aversion reduces, and in some cases can even reverse, the profitability of compensation schemes whose power of incentives increases with tenure.

JEL classification: D82

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

*This version supersedes an older version circulated under the title "Dynamic Managerial Compensation: A Mechanism Design Approach." For useful comments and suggestions, we thank Mike Fishman, Paul Grieco, Igal Hendel, Bill Rogerson and seminar participants at Hong-Kong University, London Business School, Northwestern University, the 2009 SED Meetings (Istanbul), the 2009 Canadian Conference on Economic Theory (Toronto), the 2009 Econometric Society European Meetings (Barcelona) and the 2011 SAET meeting in Portugal. A special thanks is to Yuliy Sannikov for suggestions which helped us improve the paper and to Meysam Zare for excellent research assistance. Email addresses: dfgarrett@gmail.com (Garrett); alepavan@northwestern.edu (Pavan).

1 Introduction

The contracts that the most successful firms offer to their top employees are designed taking into account that managerial ability to generate profits is bound to change over time. Shocks to managerial productivity are expected to originate from variations in the business environment that are to a large extent anticipated at the time of contracting but whose ultimate effect on managerial productivity is the managers' private information.

The questions that this paper addresses are the following: (i) How should firms respond to such shocks to managerial productivity? In particular should they induce their managers to work harder when their productivity increases or rather when they experience a negative productivity shock? and (ii) What type of compensation schemes induce the managers to respond to the shocks the way it is optimal for the firm, when both managerial effort and the shocks to managerial productivity are the managers' private information?

We assume that the firm perfectly understands the value of not renegeing on its promises and thus commits to the compensation scheme that it offers to its managers at the time they are hired. We then use a mechanism design approach to solve for the dynamic contract that maximizes the expected sum of the firm's cash flows, net of managerial compensation.

Our first result provides a characterization of a class of contracts that permit the firm to sustain *any* implementable effort policy at minimum cost. These compensation schemes are the dynamic analog of those proposed by Laffont and Tirole (1986) for a static setting. They specify a *bonus* for each period that is paid conditional upon the firm's cash flows exceeding a critical target that typically depends on past cash-flows, as well as on possible messages sent by the manager over time. The role of these messages is to permit the firm to adjust its compensation scheme in response to variations in managerial productivity.

Our second and third results identify conditions for a given effort policy to be implementable, as well as for the possibility of implementing a given effort policy at minimal cost with a *linear* or a *pseudo-linear* compensation scheme. By a linear scheme, we mean one whose compensation is linear in performance, i.e., in cash flows; by a pseudo-linear scheme, we mean one that gives the manager an utility that is linear in performance. While such schemes are less high-powered than the bonus schemes described above, there are interesting cases where the optimal effort policy for the firm can indeed be implemented at minimum cost with such schemes.

Equipped with the aforementioned results, we then investigate the properties of profit-maximizing contracts, and in particular the optimality of *seniority-based* compensation schemes. The latter are schemes that provide managers with a longer tenure in the firm with more high-powered incentives, thus inducing them to exert higher effort. To this purpose, we start by considering the case where the manager is risk-neutral.¹ We show that, for many processes of interest, the power of incentives

¹Because the firm contracts with the manager at the time the latter is already privately informed about his initial productivity, interesting dynamics emerge even without risk aversion, and even without imposing limited liability on the manager's side. Private information at the contracting stage in fact implies that the firm cannot "sell out" its business

under the optimal contract indeed increases over time, thus providing a possible explanation for the frequent practice of putting more stocks and options in the pay packages of managers with longer tenure in the firm (see, e.g., Gibbons and Murphy, 1991, Lippert and Porter, 1997, and Cremers and Palia, 2010). Contrary to other explanations proposed in the literature (e.g., declining disutility of effort or career concerns²), the optimality of seniority-based schemes is not driven by variations in the managers' preferences, or by variations in their outside option. It is a fairly natural property of optimal contracts for stationary environments where managers possess private information about their time-varying productivity. It results from an optimal allocation of informational rents over time. In other words, it originates in the firm's desire to minimize the managers' compensation while preserving their incentives for both effort and information revelation.

The driving force behind this result is the assumption that the effect of a manager's initial productivity on his future productivity (formally captured by the *impulse response* of future productivity to initial productivity, as explained in detail below) is expected to decline over time. For example, when the manager's private information evolves according to an AR(1) process, this assumption is satisfied as long as the coefficient of linear dependence is less than one. The property of declining impulse responses implies that, to minimize a manager's information rent, as computed at the time of hiring, the firm finds it optimal to induce low effort in the early periods of the relationship and higher effort in the subsequent periods. The reason is that the surplus that a manager who is highly productive at the contracting stage can obtain by mimicking a less productive type increases in the effort that the firm asks the latter to exert, as shown first in Laffont and Tirole (1986). Indeed, this surplus originates in the possibility of generating, on average, the same cash flows as the less productive type by working less, thus economizing on the disutility of effort. It follows that, by reducing the effort of those types who are less productive at the contracting stage, the firm reduces the rent that it must provide to the most productive types. The benefit of reducing the effort of the least productive types is, however, higher when done in the early stages of the relationship, when the effect of the initial productivity is still pronounced, than in the distant future, when such an effect is weak. It follows that effort (and by implication the power of incentives in the compensation scheme that sustains it) typically increases with tenure and gradually approaches the first-best level in the long run.³

We also show that, when the manager is risk neutral, the optimal effort policy can often be sustained by paying the manager according to a simple (state-contingent) linear scheme according

to the manager, for there does not exist a uniform price that extracts all surplus from all types of the manager.

²For a detailed analysis of career concerns incentives, see e.g., Dewatripont, Jewitt and Tirole (1999).

³This property of declining distortions is quite common in the literature on dynamic contracting with adverse selection (see, e.g., Baron and Besanko, 1984, Besanko, 1986, Courty and Li 2000, Battaglini, 2005, among others). As explained in Pavan, Segal, and Toikka (2011), the precise statistical property that is responsible for this result is the assumption that the impulse responses of future productivity θ_t to initial productivity θ_1 , decline over time, as opposed to other properties of the stochastic process such as the degree of correlation between θ_1 and θ_t , or the ability of θ_1 to forecast θ_t , as measured by the volatility of the forecast error of θ_t given θ_1 . One of the key contributions of the current paper is to study the extent to which such predictions are robust to the possibility that the agent is risk averse.

to which, in each period, the firm pays the manager a fixed salary plus a bonus that is linear in the firm's cash flows (or, equivalently, in the firm's stock price, assuming that the latter also depends on managerial effort). When the manager's productivity evolves according to an ARIMA process (more generally, according to any process where the impulse responses exhibit a certain separability with respect to the initial productivity), the slope of the linear scheme increases deterministically over time, i.e., it depends only on the manager's initial productivity and on the number of periods he has been working for the firm. More generally, though, the optimal contract requires that the manager be given the possibility of proposing changes to his pay over time in response to the shocks to his productivity.⁴

Building on the insights from the risk-neutral case, we then explore the properties of optimal compensation schemes for risk-averse managers. We find that, other things equal, risk-aversion reduces (and in some cases can even reverse) the profitability of seniority-based compensation schemes whose power of incentives increases, on average, over time. The reason is that these schemes entail a high sensitivity of compensation to performance precisely in those periods in which the manager faces high uncertainty about his ability to generate cash flows for the firm. Indeed, while a manager's productivity at the time he is hired is a fairly good predictor of his productivity in the near future, it is a fairly poor predictor of his productivity in the distant future. Increasing the sensitivity of compensation to performance over time thus means exposing the manager to a great deal of risk. Whether risk-averse managers with a longer tenure in the firm receive more or less high-powered incentives than younger ones then depends on the interaction between their degree of risk-aversion and the dynamics of the impulse responses of the process governing the evolution of managerial productivity.

We also find that, contrary to the risk-neutral case, the effort that the firm asks a manager to exert after the first period need not be monotone in his productivity, nor need it converge to the first-best level in the long run.

The rest of the paper is organized as follows. Section 2 describes the model. Section 3 characterizes a compensation scheme that sustains all implementable effort policy at minimum cost. Section 4 provides conditions for an effort policy to be implementable, and for the possibility of implementing it at minimum cost with linear and pseudo-linear schemes. Section 5 characterizes the properties of optimal contracts and discusses the effect of risk aversion on the optimality of seniority-based schemes. Section 6 discusses the related literature. Section 7 concludes with a few final remarks. All proofs are in the Appendix at the end of the manuscript.

⁴The idea that a manager must be given the possibility of proposing changes to his reward package is consistent with the recent empirical literature on managerial compensation where it is found that this practice has become frequent in the last decade (see, among others, Kuhnen and Zwiebel, 2008, and Bebchuck and Fried, 2004).

2 The Model

2.1 The environment

Players, actions and information. 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 can be either finite or infinite. In each period $t \in \mathbb{N}_+$ the agent receives some private information $\theta_t \in \Theta_t \subset \mathbb{R}$ about his ability to generate profits for the firm (his type), and then chooses effort $e_t \in E = \mathbb{R}$.⁵

Payoffs. The principal's payoff is the discounted sum of the firm's cash flows, net of the agent's compensation:

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

where $\delta < 1$ is the common discount factor, $\pi_t = \theta_t + e_t$ is the period- t cash flow, c_t is the period- t compensation to the agent, $\pi^T \equiv (\pi_t)_{t=1}^T$, and $c^T \equiv (c_t)_{t=1}^T$.⁶ The function U^P also corresponds to the principal's Bernoulli utility function used to evaluate lotteries over (π^T, c^T) . Both θ^T and e^T are the agent's private information. On the contrary, the stream of cash flows π^T is verifiable, which implies that the agent can be rewarded as a function of the firm's cash flows. By choosing effort e_t in period t , the agent suffers a disutility $\psi(e_t)$. The agent's preferences over (lotteries over) streams of consumption levels c^T and streams of effort choices $e^T \equiv (e_t)_{t=1}^T$ are described by an expected utility function with Bernoulli utility given by⁷

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

where $\mathcal{V} : \mathbb{R} \rightarrow \mathbb{R}$ and all $v_t : \mathbb{R} \rightarrow \mathbb{R}$ functions are strictly increasing, weakly concave, surjective (i.e., onto) and differentiable. This representation covers as special cases the situation where the agent has time-additive-separable preferences for consumption smoothing (v_t strictly concave, \mathcal{V} linear) as well as the case where the agent is risk-averse but cares only about his total compensation (\mathcal{V} strictly concave, v_t linear). The case of risk neutrality corresponds to both v and \mathcal{V} linear. As standard, the aforementioned specification presumes time consistency. In what follows, we will thus assume that, at 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 at history h_t .

⁵That effort can take negative values should not raise concerns: effort simply stands for the effect of the manager's activity on the firm's performance, which can be either positive or negative.

⁶Note that, because θ_t is *not* restricted to be independent of 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 $\theta_t^{new} = f_t(\theta^t)$. The assumption that effort in period t affects the cash flows only in period t is more restrictive, although common in the literature. This assumption is not essential for our results. We refer the reader to the earlier version of the paper (Garrett and Pavan, 2009) for the analysis of the case where π_t is a function of the entire history e^t of past effort choices.

⁷As is common in the literature, we equate the agent's period- t consumption c_t with the compensation from the principal. In other words, we assume that the agent cannot secretly save.

Stochastic process. In each period t , θ_t is drawn from a cumulative distribution function $F_t(\cdot|\theta^{t-1})$, where $\theta^{t-1} \equiv (\theta_1, \dots, \theta_{t-1}) \in \Theta^{t-1} \equiv \times_{s=1}^{t-1} \Theta_s$ with $\Theta_s = [\underline{\theta}_s, \bar{\theta}_s]$ is the history of past productivity levels. We assume that, for any t any θ^{t-1} , $F_t(\cdot|\theta^{t-1})$ is absolutely continuous over \mathbb{R} and strictly increasing with density $f_t(\theta_t|\theta^{t-1}) > 0$ over a connected set⁸ $Supp[F_t(\cdot|\theta^{t-1})] \subset \Theta_t$, and then denote by $\underline{\theta}_t(\theta^{t-1}) \equiv \inf\{Supp[F_t(\cdot|\theta^{t-1})]\}$ and $\bar{\theta}_t(\theta^{t-1}) \equiv \sup\{Supp[F_t(\cdot|\theta^{t-1})]\}$ the infimum and supremum of the support. We also assume that, for each t , $F_t(\theta_t|\theta^{t-1})$ is continuously differentiable in θ^t . Hereafter, we identify the process governing the evolution of the agent's productivity with the collection of kernels $F \equiv \langle F_t \rangle_{t=1}^T$. For each t , we then let

$$R^t \equiv \left\{ \theta^t \in \Theta^t : \theta_1 \in \Theta_1 \text{ and } \theta_l \in Supp[F_l(\cdot|\theta^{l-1})], \text{ all } l = 2, \dots, t \right\}$$

denote the set of productivity histories that are consistent with the process F . We write $\Theta = \Theta^T$ and, similarly, $R = R^T$.

We now define the impulse responses corresponding to the process F as follows. For any t any $\theta^t \in \Theta^t$, we let $J_t^t(\theta^t) \equiv 1$, whereas, for any t any $\tau > t$, any θ^τ such that $\theta^t \in \Theta^t$ and $\theta_s \in Supp[F_s(\cdot|\theta^{s-1})]$ any $s > t$, we let

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

with

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

These J functions are the nonlinear analogs of the familiar constant linear impulse responses for autoregressive processes. For example, when productivity evolves according to an AR(1) process, i.e., when $\theta_t = \gamma\theta_{t-1} + \varepsilon_t$, when γ is a scalar and $(\tilde{\varepsilon}_s)_{s=1}^T$ is a sequence of independent random variables, the impulse response of θ_τ on θ_t , $\tau > t$, is simply $J_t^\tau = \gamma^{\tau-t}$. More generally, the impulse response $J_t^\tau(\theta^\tau)$ of θ_τ to θ_t captures the total effect of an infinitesimal variation of θ_t on θ_τ , taking into account all effects on intermediate types $\theta_{>t}^{\tau-1} = (\theta_{t+1}, \dots, \theta_{\tau-1})$.⁹ In the special 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 $(\theta_t, \dots, \theta_\tau)$ and can be conveniently written as $J_t^\tau(\theta_t, \theta_{t+1}, \dots, \theta_\tau) = \prod_{k=t+1}^\tau I_{k-1}^k(\theta_k, \theta_{k-1})$, with each I_{k-1}^k given by

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

⁸The notation $Supp[F_t(\cdot|\theta^{t-1})]$ stands for the support of the distribution $F_t(\cdot|\theta^{t-1})$, defined to be the smallest compact subset of \mathbb{R} whose complement has probability zero. Note that the kernels $F_t(\cdot|\theta^{t-1})$ are defined for all possible histories $\theta^{t-1} \in \Theta^{t-1}$ including those that have zero measure from an ex ante perspective.

⁹Throughout, we use the notation $\theta_{\geq t}^s = \theta_{>t-1}^s = (\theta_t, \dots, \theta_s)$ for any periods s and t , $s \geq t$.

As anticipated in the Introduction, these impulse response functions play a key role in determining the dynamics of the profit-maximizing contract. Throughout, we will assume that the process F satisfies the property of “*first-order stochastic dominance in types*”: for all $t \geq 2$, $\theta^{t-1} > \theta^{t-1'}$ implies $F_t(\theta_t|\theta^{t-1}) \leq F_t(\theta_t|\theta^{t-1'})$ for all θ_t . This assumption in turn implies that $J_t^\tau(\theta^\tau) \geq 0$ for any t and τ , $\tau \geq t$, any $\theta^\tau \in R^\tau$.

Technical assumptions. To validate a certain dynamic envelope theorem (see Pavan, Segal, and Toikka 2011 for details), and guarantee interior solutions, we will also make the following technical assumptions. We will assume that the sets Θ_t and the functions $J_t^\tau(\cdot)$ are uniformly bounded, in the following sense: there exist nonnegative scalars $B_1, B_2 < +\infty$ such that $|\theta_t| < B_1$ and $J_t^\tau(\theta^\tau) < B_2$, all t and τ , $\tau \geq t$, all $\theta^\tau \in R^\tau$.¹⁰ Lastly, we will impose the following conditions on the disutility function ψ . Firstly, $\psi(e) = 0$ for all $e \leq 0$. Secondly, ψ is continuously differentiable over \mathbb{R} . Thirdly, there exists a scalar $\bar{e} > 0$ such that (i) $\psi'(\bar{e}) > 1$, and (ii) ψ is thrice continuously differentiable over $(0, \bar{e})$ with $\psi''(e) > 0$ and $\psi'''(e) \geq 0$ for all $e \in (0, \bar{e})$. Finally, there exist scalars $C > 0$ and $L > 1$ such that $\psi(e) = Le - C$ for all $e > \bar{e}$. These conditions are satisfied, for example, when $\bar{e} > 1$, $\psi(e) = (1/2)e^2$ for all $e \in (0, \bar{e})$, and $\psi(e) = \bar{e}e - \bar{e}^2/2$ for all $e > \bar{e}$. As mentioned above, these conditions (which are stronger than needed but easy to verify) guarantee that the agent’s payoff satisfies a certain equi-Lipschitz continuity condition which in turn permits us to conveniently express the value function through a differentiable envelope formula.

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 along with a level of consumption where both effort and consumption possibly depend on the history of observed cash flows π^t and on messages sent by the agent over time.

By the Revelation Principle, we restrict attention to direct revelation mechanisms for which a *truthful* and *obedient* strategy is optimal for each type θ_1 of the agent. Hereafter, we refer to any such mechanism as *incentive compatible*.¹¹ Letting $\Pi^t = \times_{\tau=1}^t \Pi_\tau$, with $\Pi_\tau = \{\pi_\tau \in \mathbb{R} : \pi_\tau = \theta_\tau + e, \theta_\tau \in \Theta_\tau, e \in E\}$ denoting the set of possible period- t feasible cash flows, a direct mechanism $\Omega = \langle \xi_t, s_t \rangle_{t=1}^T$ thus 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 history of cash flows π^{t-1} , while $s_t(\theta^t, \pi^{t-1}, \pi_t)$ is the principal’s payment (equivalently, the agent’s consumption) at the end of period t given the reports θ^t and the cash flows $\pi^t = (\pi^{t-1}, \pi_t)$.

¹⁰Throughout, given any function $g : \mathbb{R}^T \rightarrow \mathbb{R}$ the notation $\mathbb{E}^{\tilde{\theta}^T | \theta^t} [g(\tilde{\theta}^T)]$ will denote the expected value of the function $g(\tilde{\theta}^T)$ given the unique probability measure over Θ^T that corresponds to the kernels F starting from history θ^t .

¹¹A truthful and obedient strategy prescribes that the agent truthfully reports his private information and follows the principal’s effort recommendation at each history. Such a strategy is said to be optimal for type θ_1 if it maximizes type θ_1 ’s expected payoff among all possible (measurable) strategies. By the principle of optimality, in any incentive-compatible mechanism, with F -probability one, truthful and obedient behavior must remain optimal also at future histories that are consistent with truthful and obedient behavior in the past.

Note that $s_t(\theta^t, \pi^{t-1}, \pi_t)$ depends on both past and current cash flows. With a slight abuse of notation, for any $\theta^t \in R^t$, we then denote by $e_t(\theta^t) \equiv \xi(\theta^t, \pi^t(\theta^t))$ the equilibrium effort choice for period t and by $c_t(\theta^t) = s_t(\theta^t, \pi^t(\theta^t))$ 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 $\pi_s(\theta^s) = \theta_s + e_s(\theta^s)$ all $s \leq t$.

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 $\hat{\theta}_t \in \Theta_t$;
- The mechanism reacts by prescribing an effort choice $\xi_t(\hat{\theta}_t, \pi^{t-1})$ and a contingent reward scheme $s_t(\hat{\theta}_t, \pi^{t-1}, \cdot) : \Pi_t \rightarrow \mathbb{R}$;
- The agent privately chooses effort e_t , the cash flows $\pi_t = \theta_t + e_t$ are then observed, and the compensation $c_t = s_t(\theta^t, \pi^t)$ is paid and consumed.

The principal offers the mechanism Ω at date 1, after the agent has observed the first realization θ_1 of the process F .¹² If the agent refuses to participate in the mechanism Ω , both the agent and the principal receive their outside options, which we assume to be equal to zero. If, instead, the agent accepts Ω , he then stays in the relationship in all subsequent periods.¹³ In the terminology of mechanism design, we refer to a mechanism for which participation is optimal for each type θ_1 as *individually rational*.

Definition 1 *We say that the compensation scheme $s \equiv \langle s_t(\cdot) \rangle_{t=1}^{t=T}$ implements the effort policy $e \equiv \langle e_t(\cdot) \rangle_{t=1}^{t=T}$ if any mechanism $\Omega = \langle \xi_t, s_t \rangle_{t=1}^T$ in which the compensation scheme is s and the recommendation policy ξ satisfies $\xi_t(\theta^t, \pi^{t-1}) = e_t(\theta^t)$ for any (θ^t, π^{t-1}) such that $\theta^t \in R^t$ and $\pi_s = \theta_s + e_s(\theta^s)$ all $s \leq t-1$, is incentive-compatible and individually-rational for the agent. We say that the policy e is implementable if there exists a compensation scheme that implements it.*

3 Cost-minimizing compensation schemes

We start by characterizing a compensation scheme that implements all implementable effort policies at minimal cost for the principal.¹⁴

¹²As explained above, this assumption implies that the principal cannot simply sell out the firm to the agent even in the absence of limited liability constraints and in the presence of risk neutrality.

¹³That 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 situations/implementations where, even in the absence of bonding, participation can be guaranteed after any history consistent with truthful and obedient play in past periods.

¹⁴Throughout, for any pair of period $t, s \geq t$, any history $\theta^t \in \Theta^t$, any measurable function $g : \Theta^s \rightarrow \mathbb{R}$, we denote by $\mathbb{E}^{\tilde{\theta}^s | \theta^t} [g(\tilde{\theta}^s)]$ the expectation of $g(\tilde{\theta}^s)$ given the unique measure over Θ^s that corresponds to the process defined by the kernels F and that starts in period t with history θ^t .

Proposition 1 Suppose that the effort policy e is implementable. Let $c^{opt}(\cdot; e) \equiv \left\langle c_t^{opt}(\cdot; e) \right\rangle_{t=1}^{t=T}$ denote the (R -measurable) solution to the following problem:

$$\begin{aligned} & \min_c \mathbb{E}^{\tilde{\theta}^T} \left[\sum_{t=1}^T \delta^{t-1} c_t(\tilde{\theta}^t) \right] \\ & \text{s.t. } \mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t(c_t(\theta^t)) \right) = W(\theta^T; e) \text{ for } F\text{-almost all } \theta^T \end{aligned} \quad (2)$$

where, for each $\theta^T \in R^T$,

$$W(\theta^T; e) \equiv \sum_{t=1}^T \delta^{t-1} \psi(e_t(\theta^t)) + \int_{\underline{\theta}_1}^{\theta_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(e_t(\tilde{\theta}^t)) \right] ds + \sum_{t=2}^T \delta^{t-1} H_t(\theta^t; e)$$

with

$$\begin{aligned} H_t(\theta^t; e) & \equiv \int_{\underline{\theta}_t(\theta^{t-1})}^{\theta_t} \mathbb{E}^{\tilde{\theta}^T | \theta^{t-1}, s} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_t^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] ds \\ & \quad - \mathbb{E}^{\tilde{\theta}^T | \theta^{t-1}} \left[\int_{\underline{\theta}_t(\theta^{t-1})}^{\tilde{\theta}_t} \mathbb{E}^{\tilde{\theta}^T | \theta^{t-1}, s} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_t^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] ds \right] \end{aligned} \quad (3)$$

for all $t \geq 2$, all $\theta^t \in R^t$. Consider the compensation scheme s defined as follows: in each period t , given the reports $\theta^t \in \Theta^t$ and the observed cash flows $\pi^t \in \mathbb{R}^t$, the principal pays the agent a compensation $s_t(\theta^t, \pi^t) = c_t^{opt}(\theta^t; e)$ if $\theta^t \in R^t$ and $\pi_t \geq \pi_t(\theta^t; e) \equiv \theta_t + e_t(\theta^t)$ and charges the agent a sufficiently low penalty otherwise. The compensation scheme s described above implements the policy e at minimum cost for the principal.

By definition, if a mechanism Ω exists which implements the policy e , then at any period t and for F -almost all truthful histories¹⁵ θ^t , the agent's expected payoff in Ω under a truthful and obedient strategy must coincide with the value function $V^\Omega(\theta^t)$ defined to be the supremum of the agent's expected payoff in Ω , given the truthful history θ^t , among all possible strategies. The proof in the Appendix then uses the envelope theorem for dynamic stochastic problems of Pavan, Segal and Toikka (2011) to establish that, after any truthful history θ^{t-1} , $V^\Omega(\theta^{t-1}, \cdot)$ must be Lipschitz continuous in θ_t and satisfy the following envelope condition

$$\frac{\partial V^\Omega(\theta^t)}{\partial \theta_t} = \mathbb{E}^{\tilde{\theta}^T | \theta^t} \left[\sum_{\tau=t}^T \delta^{\tau-1} J_t^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] \quad (4)$$

¹⁵ A truthful history is one that is reached by reporting truthfully and following the principal's effort recommendations in each previous period. Without loss, these histories can be denoted by the realized sequence of productivities θ^t .

for F -almost every truthful history θ^t . Using the fact that, in any incentive-compatible mechanism Ω , at F -almost every truthful history θ^t

$$V^\Omega(\theta^t) = \mathbb{E}^{\tilde{\theta}^{t+1}|\theta^t} \left[V^\Omega(\tilde{\theta}^{t+1}) \right]$$

and iterating across periods then permits us to establish that, for F -almost all histories θ^T , the agent's utility from the compensation he receives from the principal must satisfy

$$\mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t(c_t(\theta^t)) \right) = W(\theta^T; e) + V^\Omega(\underline{\theta}_1) \quad (5)$$

with $V^\Omega(\underline{\theta}_1)$ nonnegative so as to guarantee participation by all period-1 types. This means that the utility that the agent derives from his compensation is *essentially* uniquely determined by the effort policy e , up to a nonnegative constant.

Given the property above, if the principal replaces the original mechanism Ω implementing the policy e with one where the compensation is as in the proposition then (i) the agent continues to have the right incentives to report truthfully and follow the principal's effort recommendations (this step follows from replication arguments similar to those that establish the Revelation Principle, adapted to the environment considered here) and (ii) the cost to the principal is no greater than under the original scheme Ω , thus establishing the result.

The value of Proposition 1 is twofold: it characterizes the cost to the principal of sustaining any implementable policy e , and it shows that all implementable effort policies can be sustained by offering the agent a compensation scheme that, in each period t , pays a *bonus* conditional upon the period- t cash flows exceeding a target that depends on the history of reports about the manager's productivity.

Interestingly, note that many effort policies can also be implemented with compensation schemes that are less severe than the bonus schemes of Proposition 1. For example, in the following sections we will consider implementation by schemes in which compensation is linear in the cash flows. More generally, most effort policy that can be implemented by a bonus scheme can also be implemented by a scheme that is smooth in the cash flows—in other words, the discontinuity in the bonus scheme is not an essential property of the implementation.¹⁶

It is also worthwhile noticing that communication between the agent and the principal, while a natural property of many organizations, is not essential. As long as equilibrium profits $\pi_t(\theta^{t-1}, \cdot; e)$ are monotone in current productivity θ_t (a property that we will show to hold under the optimal effort policies) then the agent's productivity can be “inferred” by the observed cash flows, in which case the same effort policy can be sustained without explicit communication.

¹⁶This follows from the fact that any step function can be approximated arbitrarily well in the L_1 norm by a smooth function preserving the agent's incentives.

Now note that a necessary condition for the compensation policy $c^{opt}(\cdot; e)$ to solve program (2) in Proposition 1 is that it satisfies the inverse Euler equation in Corollary 1 below. (The proof in the Appendix adapts arguments similar to those in Rogerson (1985) to the present environment.¹⁷)

Corollary 1 *Given any policy e , the compensation $c^{opt}(\cdot; e)$ solves program (2) in Proposition 1 only if, with F -probability one, the following condition holds for any two adjacent periods $1 \leq t, t+1 \leq T$:*

$$\frac{1}{v'_t(c_t(\theta^t))} = \mathbb{E}^{\tilde{\theta}^{t+1}|\theta^t} \left[\frac{\delta}{v'_{t+1}(c_{t+1}(\tilde{\theta}^{t+1}))} \right]. \quad (6)$$

4 Implementable policies

Building on the results in the previous section we now provide conditions for a given effort policy to be implementable. We start with the following definition.

Definition 2 *The recommendation policy ξ (defined over $\Theta^T \times \mathbb{R}^T$) is an extension of the effort policy e if (i) for any t , any $\theta^t \in \Theta^t$, $\xi_t(\theta^t, \pi^{t-1})$ is independent of past cash flows π^{t-1} (and hence denoted by $\xi_t(\theta^t)$), and (ii) for any t any $\theta^t \in R^t$, $\xi_t(\theta^t) = e_t(\theta^t)$.*

The extension ξ is bounded if there exists a constant $M > 0$ such that $\xi_t(\theta^t) < M$, all t , all $\theta^t \in \Theta^t$.

An extension is thus a recommendation policy that (i) is insensitive to the observed cash flows and (ii) is defined for all possible histories $\theta^T \in \Theta^T$, including those that have zero measure under the process F , and coincides with the policy e with F -probability one. The role of such extensions is to specify a behavior for the agent at all histories, including those that can be reached only by departing from truth-telling. We then have the following result.

Proposition 2 *Suppose that there exists a bounded¹⁸ extension ξ of the policy e satisfying the following single-crossing conditions*

$$\mathbb{E}^{\tilde{\theta}^\tau|\theta^t} \left[\begin{aligned} & \psi'(\xi_t(\hat{\theta}^{t-1}, \tilde{\theta}_t)) - \psi'(\xi_t(\hat{\theta}^{t-1}, \hat{\theta}_t) + \hat{\theta}_t - \tilde{\theta}_t) \\ & + \sum_{\tau=t+1}^T \delta^{\tau-t} J_t^\tau(\tilde{\theta}^\tau) \left[\psi'(\xi_\tau(\hat{\theta}^{t-1}, \tilde{\theta}_t, \tilde{\theta}_{>t}^\tau)) - \psi'(\xi_\tau(\hat{\theta}^{t-1}, \hat{\theta}_t, \tilde{\theta}_{>t}^\tau)) \right] \end{aligned} \right] [\theta_t - \hat{\theta}_t] \geq 0 \quad (7)$$

for any t , any $\theta^t \in R^t$, any $\hat{\theta}^t \in \Theta^t$. In addition, assume that one of the following conditions holds: (i) the process F is Markov; (ii) each $\xi_t(\theta^t)$ depends on θ^t only through θ_1 . In each of these cases, the policy e is implementable.

¹⁷Note that Condition (6) holds irrespective of whether the agent's preferences are time-additively separable.

¹⁸The only role played by the boundedness of the extension ξ is to validate a certain one-stage-deviation principle in case $T = \infty$. Note that, because flow payoffs are potentially unbounded, the standard version of the one-stage-deviation principle does not apply here.

Note that the single-crossing conditions in the propositions are trivially satisfied when, in each period, the effort that the firm recommends is nondecreasing in the history of reported productivity. More generally, these conditions are also satisfied when, in the continuation game starting with an arbitrary history $(\theta^t, \hat{\theta}^{t-1}, e^{t-1})$, effort increases, on average, with the agent's report about his current productivity, where the average (both across time and productivity sequences) is computed using the impulse response functions.¹⁹

To appreciate the idea behind the result in the proposition, consider first the case where the process is Markov. In this case, whether or not the agent has been truthful in the past is irrelevant for the incentives he faces in the continuation game that starts after any history $(\theta^t, \hat{\theta}^{t-1}, e^{t-1})$. Suppose, then, that the principal offers a scheme of the type described in Proposition 1; that is, in each period t , given the reports $\hat{\theta}^t \in \Theta^t$, the principal pays a bonus $c_t(\hat{\theta}^t)$ if the current cash flow exceeds the target $\pi_t = \hat{\theta}_t + \xi_t(\hat{\theta}^t)$ and charges the agent a sufficiently large penalty otherwise. Faced with this scheme, the effort that the agent exerts in each period t is determined entirely by the history of reports $\hat{\theta}^t$ made over time, and by the agent's current type θ_t and is given by $\hat{e}_t = \xi_t(\hat{\theta}^t) + \hat{\theta}_t - \theta_t$, all t . Then let $U^\#(\theta^t; \hat{\theta}^{t-1}, \hat{\theta}_t)$ denote the agent's expected payoff in the continuation game that starts in period t with history $(\theta^t, \hat{\theta}^{t-1})$ when, in period t , the agent sends the report $\hat{\theta}_t$, he then chooses effort $\hat{e}_t = \xi_t(\hat{\theta}^t) + \hat{\theta}_t - \theta_t$, and then starting from period $t + 1$ onward he follows a truthful and obedient strategy. The proof in the Appendix verifies that, when the single-crossing conditions in the proposition are satisfied, there exists a compensation scheme of the type described above such that

$$\left[\frac{dU^\#(\theta^t; \hat{\theta}^{t-1}, \theta_t)}{d\theta_t} - \frac{\partial U^\#(\theta^t; \hat{\theta}^{t-1}, \hat{\theta}_t)}{\partial \theta_t} \right] [\theta_t - \hat{\theta}_t] \geq 0,$$

all $\theta^t \in R^t$, $\hat{\theta}^t \in \Theta^t$. As is well known, this endogenous single-crossing condition on derivatives implies that $U^\#(\theta^t; \hat{\theta}^{t-1}, \theta_t) \geq U^\#(\theta^t; \hat{\theta}^{t-1}, \hat{\theta}_t)$ all $\theta^t \in R^t$, $\hat{\theta}^t \in \Theta^t$, meaning that no one-shot deviation from the truthful and obedient strategy is ever profitable for the agent. Starting from this result, one can then verify that all other deviations are also unprofitable, thus establishing the implementability of the proposed effort policy.²⁰

Next consider the case where the process is possibly non-Markov, but where the effort policy depends only on θ_1 . The proof in the Appendix shows that one can construct compensation schemes in which the agent is asked to report only his initial type and such that the utility that he derives from the payments he receives over time is linear in the cash flows π_t , $t > 1$. The idea here is the

¹⁹Note that the single-crossing conditions in (7) are trivially satisfied when, given each productivity sequence θ^T , with $\theta_\tau \in \text{Supp}[F_\tau(\cdot|\theta^{\tau-1})]$ all $\tau > t$, the intertemporal sum $\sum_{\tau=t}^T \delta^{\tau-t} J_t^\tau(\theta^\tau) \psi'(\xi_\tau(\theta^{t-1}, \cdot, \theta_{t+1} \dots, \theta_\tau))$ is nondecreasing in the period- t report $\hat{\theta}_t$ (ex-post monotonicity). More generally, the single-crossing conditions in (7) only require that the expectation (over θ^T) of the above intertemporal average changes signs only once at $\hat{\theta}_t = \theta_t$.

²⁰Note that the scheme described above, contrary to that in Proposition 1, does not constrain the agent to report only feasible sequences of types $\theta^T \in R$. Allowing the agent to report also types that are inconsistent with the process F is just a trick that permits us to establish the suboptimality of one-stage deviations from the truthful and obedient strategy after *any* history, including those that entailed a deviation in the past. Together with continuity a infinity, this in turn establishes the optimality of the truthful and obedient strategy, and hence the implementability of the desired effort policy.

following. First, because the policy depends only on θ_1 , the agent does not need to report anything to the principal after the first period. This eliminates the complication stemming from the fact that, with non-Markov processes, verifying the agent's incentives for truthtelling is typically hard due to the fact that his continuation payoff may depend on the entire history of past types. Second, by using a scheme that makes the agent's payoff, gross of his disutility of effort, linear in the cash flows π_t , $t > 1$, the principal can control the agent's incentives for effort without the need to set bonuses that depend on past types (i.e., without communicating with the agent after the first period). It suffices to set the sensitivity of the the payoff to each cash flow π_t , $t > 1$, equal to $\psi'(\xi_t(\theta_1))$. It is then easy to see that, conditional on having reported truthfully in the first period, and irrespective of the effort exerted in the preceding periods, at any period $t > 1$, the agent finds it optimal to choose a level of effort $e_t = \xi_t(\theta_1)$ for which the marginal disutility is equal to the marginal increase in compensation. The only difficulty for the principal is then to induce the agent to report truthfully in the first period. When the effort policy satisfies the single-crossing condition in the proposition²¹, this can be achieved by committing to pay, in addition to the compensation described above, a bonus in period one conditional upon the period-1 cash flow exceeding a target $\pi_1 = \theta_1 + \xi_1(\theta_1)$ that depends only on the first period announcement. It is important to recognize that, while the schemes described above do facilitate the verification of the implementability of a policy, they are not essential for its sustainability: for example, any effort policy that can be sustained with a scheme of the type described above can also be sustained with a bonus scheme of the type described in Proposition 1, at a weakly lower cost for the principal.

Lastly, note that the case where the effort policy is independent of the agent's reports after the first period is of particular interest. As we will show below, the optimal policy takes this form for example when the agent is risk neutral and the process for the agent's productivity is separable in the first component, in the following sense.

Definition 3 *The process F is separable in the first component (SFC) if each impulse response function $J_1^t(\theta^t)$ depends on θ^t only through θ_1 , which in turn is the case when the process F admits the following representation $\theta_t = \kappa_t(\theta_1) + \phi_t(\varepsilon_2, \dots, \varepsilon_t)$ for some functions $\kappa_t : \Theta_1 \rightarrow \mathbb{R}$ and $\phi_t : \mathcal{E}^t \rightarrow \mathbb{R}$, with the shock ε drawn independently from θ_1 .*

Note that the family of SFC processes includes all ARIMA processes: in this case, each impulse response J_1^t is simply a scalar.

4.1 Implementability (at minimum cost) with linear and pseudo-linear schemes

Before turning to the properties of optimal policies, we conclude this section with a result concerning the possibility of implementing certain effort policies with linear and pseudo-linear schemes. By the former, we mean compensation schemes in which the total compensation that the agent receives is linear in the cash flows. By the latter, we mean compensation schemes in which the *utility* that the

²¹Note that, when the policy depends only on θ_1 , the only relevant single-crossing condition is the one for $t = 1$.

agent derives from his compensation is linear in the cash flows.²² The reason for looking at these schemes is twofold. Many compensation packages used in the real world based on a time-varying combination of stocks and options are known to give rise to linear sensitivities of compensation to firm's performance, as measured for example by cash flows and stock prices. Second, pseudo-linear schemes (which of course coincide with linear schemes when the agent is risk neutral) provide a desirable form of robustness. They guarantee that the agent finds it optimal to report truthfully and choose the desired level of effort even if he is concerned about the possibility that the firm's cash flows may depend not only on his productivity and effort but also of noise that is beyond his control. For example, such schemes continue to implement the desired effort policy even if the agent believes the cash flows to be given by $\pi_t = \theta_t + e_t + \eta_t$ where η_t is transitory noise, independent of the process F and of the agent's effort decisions. Motivated by the aforementioned considerations, we then have the following result.

Proposition 3 *Let ξ be any bounded extension of the policy e satisfying the following single-crossing conditions*

$$\mathbb{E}^{\hat{\theta}^T | \theta^t} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_t^\tau(\tilde{\theta}^\tau) \left[\psi' \left(\xi_\tau(\hat{\theta}^{t-1}, \tilde{\theta}_t, \tilde{\theta}_{>t}^\tau) \right) - \psi' \left(\xi_\tau(\hat{\theta}^{t-1}, \hat{\theta}_t, \tilde{\theta}_{>t}^\tau) \right) \right] \right] [\theta_t - \hat{\theta}_t] \geq 0 \quad (8)$$

for any t , any $\theta^t \in R^t$, any $\hat{\theta}^t \in \Theta^t$. In addition, assume that one of the following conditions holds: (i) the process F is Markov; (ii) each $\xi_t(\theta^t)$ depends on θ^t only through θ_1 . In each of these cases, the policy e can be implemented at minimal cost for the principal with an appropriately designed pseudo-linear scheme.

First note that the single crossing conditions in the proposition are stronger than those in Proposition 2. These schemes thus permit the principal to implement fewer effort policies than the corresponding bonus schemes of Proposition 2. Nonetheless, there are interesting cases where the optimal effort policies can indeed be implemented with such schemes, as will be shown below.

Next note that these schemes are designed so as to achieve multiple goals at once. First, they undo the agent's risk-aversion by making his payoff, gross of the cost of effort, linear in the cash-flows π_t , $t \geq 1$. By setting the sensitivity to each cash flow equal to $\alpha_t(\theta^t) = \psi'(\xi_t(\theta^t))$ such linearity guarantees that, conditional on having reported truthfully, and irrespective of past effort choices, the agent finds it optimal to exert the right level of effort, thus taking care of the moral-hazard part of the problem. Second, these schemes are designed so as to induce truthtelling. This is obtained through a careful design of the part of the compensation that (in terms of the agent's payoff) is independent of the realized cash flow but which depends on the reported productivity. This takes care of the adverse selection part of the problem. Lastly, these schemes are designed so that the compensation

²²The discussion after Proposition 2 mentioned already some of the advantages of linear and pseudo-linear schemes. However, there we restricted attention to policies that depend only on θ_1 . Furthermore, we did not address the question of what conditions guarantee that such schemes implement the desired effort policies *at minimal cost for the principal*.

the agent receives over time is the one that minimizes the cost to the principal. This is obtained by distributing the compensation over time according to Rogerson's inverse Euler condition, as defined in Corollary 1, thus taking care of the optimal consumption smoothing part of the problem.

To illustrate, consider for simplicity the case where \mathcal{V} is linear but where each v_t is possibly strictly concave, thus maintaining a preference for consumption smoothing, while making the payoff time-additively separable. Let ξ be any extension of the desired effort policy e from R to Θ . Let $c^{opt}(\cdot; \xi)$ be the cost-minimizing allocation of consumption over time, as defined in the program of Proposition 1, with all policies and corresponding constraints naturally extended from R to Θ . Then consider the following pseudo-linear payment scheme s^{PL} . In each period t , and for all (θ^t, π^t) , the principal pays the agent a compensation

$$s_t^{PL}(\theta^t, \pi^t) = v_t^{-1} \left(v_t \left(c_t^{opt}(\theta^t; \xi) \right) + \psi'(\xi_t(\theta^t)) [\pi_t - \theta_t - \xi_t(\theta^t)] \right). \quad (9)$$

It is easy to see that, under such a scheme, in each period t , conditional on having reported truthfully, the agent has the right incentives for choosing the desired level of effort $\xi_t(\theta^t)$. It is also easy to see that, on path, the compensation that the agent receives over time is given by $s_t^{PL}(\theta^t, \pi^t) = c_t^{opt}(\theta^t; \xi)$ which is exactly the compensation that minimizes the cost to the principal among all those that are consistent with the agent's incentive compatibility, as indicated in Proposition 1. As shown in the proof in the Appendix, that the effort policy satisfies the single crossing conditions in the proposition in turn guarantees that, when either the process is Markov, or the policy depends only on θ_1 , (i) the agent has the right incentives to report all his types truthfully, and (ii) each type θ_1 has the incentive to participate, thus establishing implementability of the desired effort policy.

Next, consider the special case where the agent is risk neutral (both \mathcal{V} and each v_t linear). It is worthwhile noticing that, while in this case the particular way the compensation is distributed over time is irrelevant, certain choices have the advantage of guaranteeing that, even if the agent had the option to leave the relationship after the first period, he would never find it optimal to do so. To see this, consider for example the case where F is Markov, $R = \Theta$, and T is finite. By committing to pay the agent in each period a compensation

$$s_t(\theta^t, \pi^t) = S_t(\theta^t) + \alpha_t(\theta^t) \pi_t \quad (10)$$

with²³

$$S_t(\theta^t) \equiv \psi(\xi_t(\theta^t)) - \alpha_t(\theta^t) [\theta_t + \xi_t(\theta^t)] + H_t(\theta^t; \xi)$$

²³The functions $H_t(\cdot; \xi)$ (with domain Θ^t) are defined analogously to the functions $H_t(\theta^t; e)$ in (3) by replacing the equilibrium policy e with the extension ξ and by letting the lower bounds of integration be $\underline{\theta}_t$ as opposed to $\underline{\theta}_t(\theta^{t-1})$.

if $t > 1$ and

$$S_1(\theta_1) \equiv \psi(\xi_1(\theta_1)) - \alpha_1(\theta_1) [\theta_1 + \xi_1(\theta_1)] + \int_{\theta_1}^{\theta_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(\xi_t(\tilde{\theta}^t)) \right] ds$$

if $t = 1$, the principal guarantees that, after each period t , the agent's continuation payoff under a truthful and obedient strategy from period t onwards is equal to

$$\int_{\theta_t}^{\theta_t} \mathbb{E}^{\tilde{\theta}_{\geq t}^T | s} \left[\sum_{\tau=t}^T \delta^{\tau-1} J_t^\tau(\tilde{\theta}_{\geq t}^\tau) \psi'(\xi_\tau(\tilde{\theta}_{\geq t}^{\tau-1}, \tilde{\theta}_{\geq t}^\tau)) \right] ds$$

irrespective of whether or not the agent has been truthful and obedient in the past. Because the latter is clearly positive, this implies that the agent never finds it optimal to leave the firm.

When the agent is risk-averse, the pseudo-linear schemes give the agent a payoff which is linear in the cash flows but typically entail a non-linear compensation. However, when the single crossing conditions of Proposition 3 hold, it may be possible to implement the desired effort policy with a payment scheme where the compensation itself is linear in the cash flows. To illustrate, consider again the case where \mathcal{V} is linear but where each v_t is strictly concave and let $w_t(\cdot) \equiv v_t^{-1}(\cdot)$. Then let s^L be the linear scheme defined for all t , all (θ^t, π^t) , by

$$s_t^L(\theta^t, \pi^t) \equiv c_t^{opt}(\theta^t; \xi) + w_t' \left(v_t \left(c_t^{opt}(\theta^t; \xi) \right) \right) \psi'(\xi_t(\theta^t)) [\pi_t - \theta_t - \xi_t(\theta^t)].$$

By the convexity of each w_t and Jensen's inequality, we then have that

$$s_t^L(\theta^t, \pi^t) \leq s_t^{PL}(\theta^t, \pi^t),$$

with equality if and only if $\pi_t = \theta_t + \xi_t(\theta^t)$, where s^{PL} is the pseudo-linear scheme defined in (9). Faced with the above linear scheme, the agent's payoff under a truthful and obedient strategy is thus the same as when he faces the corresponding pseudo-linear scheme, whereas his payoff under any other plan is weakly lower. This means that, whenever the desired effort policy is implementable at minimal cost with a pseudo-linear scheme it is also implementable at minimal cost with a linear scheme. Given the popularity of linear schemes in the real world and in the managerial compensation literature, the above result is somewhat reassuring.

5 Profit-maximizing effort policies

5.1 Risk-neutrality

Equipped with the results in the previous section, we now characterize the effort policies that maximize the firm's expected cash flows, net of the cost of managerial compensation. We start by characterizing optimal policies for risk-neutral agents. Recall that, by definition, in any mechanism

that is incentive-compatible for the agent, the payoff that each type θ_1 expects under a truthful and obedient strategy must coincide with the value function. Applying (4) to $t = 1$, the latter in turn must satisfy

$$V^\Omega(\theta_1) = \int_{\underline{\theta}_1}^{\theta_1} \mathbb{E}^{\tilde{\theta}^T|s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(s, \tilde{\theta}_{>1}^t) \psi'(e_t(s, \tilde{\theta}_{>1}^t)) \right] ds + V^\Omega(\underline{\theta}_1) \quad (11)$$

for all θ_1 . Integrating (11) by parts, we then have that the expected surplus that the principal must leave to the agent, as computed at the time of hiring, is given by

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

where $\eta(\theta_1) \equiv [1 - F_1(\theta_1)]/f_1(\theta_1)$ denotes the inverse hazard rate of the first-period distribution. Using the fact that the principal's expected payoff is equal to

$$\mathbb{E}[U^P] = \mathbb{E}^{\tilde{\theta}^T} \left[\sum_{t=1}^T \delta^{t-1} \left\{ \tilde{\theta}_t + e_t(\tilde{\theta}^t) - \psi(e_t(\tilde{\theta}^t)) \right\} \right] - \mathbb{E} [V^\Omega(\tilde{\theta}_1)]$$

we then have that

$$\mathbb{E}[U^P] = \mathbb{E}^{\tilde{\theta}^T} \left[\sum_{t=1}^T \delta^{t-1} \left\{ \tilde{\theta}_t + e_t(\tilde{\theta}^t) - \psi(e_t(\tilde{\theta}^t)) - \eta(\tilde{\theta}_1) J_1^t(\tilde{\theta}^t) \psi'(e_t(\tilde{\theta}^t)) \right\} \right] - V^\Omega(\underline{\theta}_1). \quad (12)$$

Given that $J_1^t(\theta^t) \geq 0$ for all t all θ^t under FOSD, we then have that the effort policy that maximizes (12) can be obtained by pointwise maximization of $\mathbb{E}[U^P]$ and is given in the next proposition.

Proposition 4 *Suppose that the agent is risk neutral. (i) Let e^* be the effort policy implicitly defined, for all t all $\theta^t \in R^t$, by²⁴*

$$\psi'(e_t^*(\theta^t)) = 1 - \eta(\theta_1) J_1^t(\theta^t) \psi''(e_t^*(\theta^t)). \quad (13)$$

Suppose that the policy e^ is implementable. Then, under any optimal contract for the principal, e^* is sustained with F -probability one.*

(ii) The following conditions guarantee that the policy e^ is implementable: (a) the process F is either Markov, or separable in the first component (SFC); and (b) each function $\eta(\theta_1) J_1^t(\theta^t)$ is nonincreasing. Under these conditions, the policy e^* can also be implemented by a linear scheme.²⁵ Furthermore, by committing to pay the agent in each period according to (10) the principal guarantees that the agent never finds it optimal to leave the relationship.*

First, consider part (i). It is easy to see that the policy e^* maximizes (12). That, when it is implementable, e^* is sustained under any optimal contract for the principal then follows from the fact that the principal's payoff in any mechanism that is incentive-compatible for the agent is given

²⁴The formula in (13) presumes that the effort $e^*(\theta^t)$ implicitly defined by (13) is strictly positive, which is the case if and only if $\psi''(0) < 1/[\eta(\theta_1) J_1^t(\theta^t)]$. When this is not the case then $e_t^*(\theta^t) = 0$.

²⁵The extension ξ of the policy e^* from R to Θ is given in the Appendix.

by (12), together with the fact that, from Proposition 1, there exists a compensation scheme s^* that implements e^* and that gives the lowest period-1 type θ_1 an expected payoff equal to his outside option, in which case $V^\Omega(\theta_1) = 0$.

Next consider part (ii). The result follows from Propositions 2 and 3. When each function $\eta(\theta_1)J_1^t(\theta^t)$ is nonincreasing, then each function $e_t^*(\theta^t)$ is nondecreasing. In the Appendix, we show that, when this is the case, there always exists an extension ξ of e^* from R to Θ that satisfies not only the single-crossing conditions (7) of Proposition 2 but also the more stringent conditions (8) of Proposition 3. The results in those propositions then imply that when F is Markov, or it is separable in the first component (SFC)—in which case each $e_t^*(\theta^t)$ and hence also each corresponding $\xi_t(\theta^t)$ depends on θ^t only through θ_1 —then the policy e^* can be implemented both by the bonus schemes of Proposition 1 and by the linear schemes described in (10). That, in this case, participation can be guaranteed not only in the first period but also in any subsequent period then follows from what discussed after Proposition 3.

Note that the assumption that each function $\eta(\theta_1)J_1^t(\theta^t)$ is nondecreasing in θ^t is trivially satisfied, for example, when the density of the first period distribution F_1 is log-concave and the process F is ARIMA. As anticipated above, in this case, the impulse response functions J_1^t are scalars and the dynamics of effort under the optimal contract is entirely deterministic. The implementation of the optimal contract is then particularly simple. In period one, the principal offers the agent a menu of contracts, indexed by θ_1 . Each contract specifies for each period t a fixed payment $S_t(\theta_1)$, along with a variable pay $\alpha_t(\theta_1)\pi_t$ that is linear in the observed cash flow π_t , as given by (10). Both the fixed component $S_t(\theta_1)$ and the variable component $\alpha_t(\theta_1)$ vary over time, but in the case of a process that is separable in the first component, do not depend on the productivity shocks experienced after the first period.

A few more observations are in order. First, note that, when the optimal effort policy is the one in Proposition 4, then in each period t , the agent's effort $e_t^*(\theta^t)$ is downward distorted with respect to the first-best level, which is implicitly defined by $\psi'(e^{FB}) = 1$ for all t all θ^t . This property originates in the principal's desire to limit the surplus (equivalently, the rent) that she must leave to the agent to induce him to truthfully reveal his productivity level. By offering to those types who are less productive at the hiring stage (for which $\eta(\theta_1)$ is high) a contract with low-powered incentives, the principal makes it less attractive for the most productive types to mimic. As anticipated in the Introduction, this is because the surplus that a more productive type obtains by mimicking a less productive one is the disutility of effort that he can save by generating the same cash flows as the less productive type by working less. This means that the lower the effort that the least productive types are induced to exert, the lower the surplus that the principal must leave to the most productive ones.

Second, note that the dynamics of effort under the policy e^* is entirely driven by the dynamics of the impulse responses $J_1^t(\theta^t)$ of future types to the initial types. The reason is that these functions capture the extent to which the agent's initial private information θ_1 has persistent effects on the

surplus that the principal must leave to the most productive types to induce them to reveal their private information. To illustrate, consider the following examples.

Example 1 Suppose that θ_t evolves according to an AR(1) process $\theta_t = \gamma\theta_{t-1} + \varepsilon_t$, for some $\gamma \in (0, 1)$. Then $J_1^t(\theta^t) = \gamma^{t-1}$ for all t all θ^t . It follows that e_t^* increases over time and, for any θ_1 , $\lim_{t \rightarrow \infty} e_t^*(\theta_1) = e^{FB}$. \\

Example 2 Assume that each θ_t is i.i.d., so that $J_1^t(\theta^t) = 0$ for all $t \geq 2$ all θ^t . Then effort is distorted only in the first period, i.e. $e_1^*(\theta_1) < e_1^{FB}$ (unless $\theta_1 = \bar{\theta}_1$, in which case $e_1^*(\theta_1) = e_1^{FB}$) and $e_t^*(\theta^t) = e^{FB}$ for all $t \geq 2$ all θ^t . \\

Example 3 Suppose that θ_t follows a random walk, i.e. $\theta_t = \theta_{t-1} + \varepsilon_t$. Then $e_t^*(\theta^t)$ is constant over time and depends only on θ_1 . \\

The property that effort increases over time and gradually converges to the first-best level under the AR(1) process of Example 1 is actually more general; it also applies for example to any ARIMA processes for which J_1^t decreases with t with $\lim_{t \rightarrow \infty} J_1^t = 0$, where J_1^t are nonnegative scalars that depend on the various parameters of the ARIMA process.

Note that, as the examples make clear, the statistical property of the process F that is responsible for the dynamics of effort is the dynamics of the impulse responses $J_1^t(\theta^t)$ and not the dynamics of the correlation between $\tilde{\theta}_1$ and $\tilde{\theta}_t$, or the extent to which $\tilde{\theta}_1$ is a good predictor of future productivity $\tilde{\theta}_t$, as measured by the variance of the forecast error of $\tilde{\theta}_t$ given $\tilde{\theta}_1$. As one can see for example by looking at the random walk case of Example 3, $Corr(\tilde{\theta}_1, \tilde{\theta}_t)$ decreases over time and $Var([\tilde{\theta}_t - \mathbb{E}(\tilde{\theta}_t|\theta_1)]^2)$ increases over time. Nonetheless, $e_t^*(\theta^t) - e^{FB}$ is constant over time due to the fact that the impulse response J_1^t of θ_t to θ_1 are constant in the random walk case.

Also note that, when productivity evolves according to a random walk, then 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 use the same compensation scheme that the agent selects in period one also in each subsequent period. Each linear scheme (indexed by θ_1) has a fixed compensation of

$$S(\theta_1) \equiv \psi(e^*(\theta_1)) + \int_{\underline{\theta}_1}^{\theta_1} \psi'(e^*(s))ds - \alpha(\theta_1)[\theta_1 + e^*(\theta_1)]$$

together with a piece-rate $\alpha(\theta_1)\pi_1$ with $\alpha(\theta_1) = \psi'(e^*(\theta_1))$ that is constant over time. These contracts are reminiscent of those derived in Laffont and Tirole (1986) in a static setting. Contrary to the static case, the entire linear scheme $S(\theta_1) + \alpha(\theta_1)\pi_t$ — as opposed to the point $S(\theta_1) + \alpha(\theta_1)[\theta_1 + e^*(\theta_1)]$ — is now used over time. This is a direct consequence of the fact that the firm's performance π_t now changes stochastically over time in response to the shocks affecting the agent's productivity. 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). 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 implies is that, despite the lack of stationarity, it is optimal for the principal to commit to use the same scheme selected in period one also in each subsequent period, thus inducing the same effort as if the agent's type were constant over time. Also note that the optimal reward scheme (and the corresponding effort dynamics) when θ_t follows a random walk coincides 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 agent's productivity. Indeed, it is easy to verify that the dynamics of effort under the optimal contract is the same when $\theta_t = \theta_1 + \sum_{s=2}^t \varepsilon_s$ as when $\theta_t = \theta_1 + \varepsilon_t$. This property holds, more generally, for any process F that is separable in the first component, i.e., such that $\theta_t = \kappa_t(\theta_1) + \phi_t(\varepsilon_2, \dots, \varepsilon_t)$ for some functions κ and ϕ : the dynamics of effort is then completely independent of the functions ϕ_t that determine the effect of shocks experienced after period one on the agent's current productivity.

Seniority. We now turn to the key property discussed in the Introduction that the power of incentives increases with tenure (i.e., with the length of the employment relationship). What the examples above have in common is the property that the effect of the agent's initial productivity on his future productivity, as captured by the impulse response J_1^t is either constant (in the random walk-case) or declines over time (gradually, in the AR(1) case, and with a single jump to zero at $t = 2$ in the case of independent types). More generally, all these examples share the property of “declining impulse responses” defined as follows.

Definition 4 *The process F satisfies the property of “declining impulse responses” if, for any $s > t \geq 1$, any $(\theta^t, \theta_{>t}^s)$, $\theta_s \geq \theta_t$ implies that $J_1^s(\theta^t, \theta_{>t}^s) \leq J_1^t(\theta^t)$.*

Inspecting the formula for the principal's expected payoff as given by (12), together with standard comparative statics arguments, then immediately reveals that, when this property holds, then as long as the agent's productivity does not fall, the effort that he is asked to exert under the optimal policy does not fall either. Declining impulse responses thus provide a condition under which the agent's effort increases monotonically over time.

An alternative definition of the diminishing effect of the agent's initial productivity on his future productivity can be obtained by taking an ex-ante view and assuming that

$$\frac{\partial \mathbb{E}^{\tilde{\theta}^T | \theta_1}[\tilde{\theta}_t]}{\partial \theta_1} = \mathbb{E}^{\tilde{\theta}^T | \theta_1} \left[J_1^t(\tilde{\theta}^t) \right]$$

decreases with t . In case the agent's disutility of effort is quadratic, one can indeed verify that, under this condition, the agent's expected effort under the optimal policy, $\mathbb{E}^{\tilde{\theta}^T | \theta_1}[e_t^*(\tilde{\theta}^t)]$ (as well as the expected slope of the incentive scheme that sustains it) increases over time. This property, however, does not necessarily extend to more general disutility functions, for it relies on a certain convexity property of the principal's payoff as given by (12).

What the aforementioned definitions of declining impulse responses both capture is the idea that the effect of a manager's initial productivity on his subsequent productivity declines over time, a property that seems appropriate for many situations of interest. This property implies that, as a manager's tenure in the firm grows, the firm finds it optimal to induce him to exert higher effort. This requires a higher sensitivity of his compensation to performance, which can be obtained, for example, by putting more stocks and options in the compensation package.²⁶

Interestingly, note that the reason why, with a risk-neutral agent, the power of incentives increases over time has nothing to do with variations in the manager's preferences or in his outside option (as, for example, in the career-concerns or in the learning-by-doing literature). It is entirely driven by the fact that, when hired, the manager is expected to possess private information about his ability to generate profits for the firm, along with the fact that the effects of such initial private information on future productivity is expected to decline over time. Under these assumptions, to reduce the incentives of those managers who are most productive at the hiring stage from mimicking the less productive ones, it is more effective for the firm to distort effort in the latter's contracts more in the first few periods, when the effect of the initial productivity is still pronounced, than in the distant future when such an effect is weak. Our theory thus complements the aforementioned two theories by offering an alternative justification for the optimality of seniority-based schemes whose power of incentives increases over time.

5.2 Risk aversion

We now investigate how the optimality of seniority-based schemes may be affected by the possibility that the agent is risk averse. For simplicity, we assume here that T is finite (in which case discounting can be dropped) and start by considering the case where the agent does not have preferences for consumption smoothing, so that \mathcal{V} is concave but each v_t is linear. Using the result in Proposition 1, the principal's expected payoff under any cost-minimizing compensation scheme can be expressed as

$$\begin{aligned} \mathbb{E}[U^P] &= \mathbb{E}^{\tilde{\theta}^T} \left[\sum_{t=1}^T (\tilde{\theta}_t + e_t(\tilde{\theta}^t)) \right] \\ &\quad - \mathbb{E}^{\tilde{\theta}^T} \left[\mathcal{V}^{-1} \left(\begin{aligned} &\sum_{t=1}^T \psi(e_t(\tilde{\theta}^t)) \\ &+ \int_{\underline{\theta}_1}^{\tilde{\theta}_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{t=1}^T J_1^t(\tilde{\theta}^t) \psi'(e_t(\tilde{\theta}^t)) \right] ds + \sum_{t=2}^T H_t(\tilde{\theta}^t; e) \end{aligned} \right) \right] \end{aligned} \quad (14)$$

²⁶ As explained in the Introduction, more effort does not necessarily mean more labor, for effort in this model simply captures the effect of the agent's various activity on the firm's expected cash flows.

where

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

The expression in (14) is the analog of the *dynamic virtual surplus* formula in (12) for the case of a risk-averse agent.

We now show how the formula in (14) can be used to characterize the profit-maximizing effort policy. We start with an example where $T = 2$ and where the second-period productivity θ_2 depends linearly on θ_1 as in Examples 1-3. We also restrict attention to effort policies that depend on θ_1 only, and are thus independent of the productivity “shock” which occurs at the beginning of the second period. Finally, we assume that the inverse of the agent’s utility function over consumption is quadratic. These assumptions help us illustrate the key trade-offs in the simplest possible way. We will come back to the general problem at the end of the section.

Example 4 Suppose that $T = 2$ and that the initial productivity θ_1 is drawn from an absolutely continuous distribution F_1 with strictly positive density f_1 on an interval $\Theta_1 = [\underline{\theta}_1, \bar{\theta}_1] \subset \mathbb{R}$. Second-period productivity is given by $\theta_2 = \gamma\theta_1 + \varepsilon_2$ where $\gamma \geq 0$ (note that $J_1^2 = \gamma$) and where the “shock” to productivity ε_2 is drawn from an absolutely continuous distribution G_2 with a strictly positive density g_2 on a compact interval $\mathcal{E}_2 \subset \mathbb{R}$, independently from θ_1 . The agent’s utility function over consumption is given by $\mathcal{V}(c) = \frac{1}{\alpha} \sqrt{2\alpha c + \beta^2} - \frac{\beta}{\alpha}$ with $\alpha, \beta > 0$; note that this function is chosen so that $\mathcal{V}^{-1}(u) = \frac{\alpha}{2} u^2 + \beta u$. In this environment, any effort policy e^{ND} that maximizes the principal’s expected payoff (14) among those that depend on θ_1 only is such that, for any t , F -almost any θ_1 ,²⁷

$$\psi'(e_1^{ND}(\theta_1)) \left[\alpha \left(\mathbb{E}^{\tilde{\theta}^2 | \theta_1} \left[W(\tilde{\theta}^2; e^{ND}) \right] \right) + \beta \right] \\ = 1 - \frac{\psi''(e_1^{ND}(\theta_1))}{f_1(\theta_1)} \int_{\theta_1}^{\bar{\theta}_1} \left[\alpha \left(\mathbb{E}^{\tilde{\theta}^2 | r} \left[W(\tilde{\theta}^2; e^{ND}) \right] \right) + \beta \right] f_1(r) dr \quad (15)$$

and

$$\psi'(e_2^{ND}(\theta_1)) \left[\alpha \left(\mathbb{E}^{\tilde{\theta}^2 | \theta_1} \left[W(\tilde{\theta}^2; e^{ND}) \right] \right) + \beta \right] \\ = 1 - \frac{\psi''(e_2^{ND}(\theta_1))}{f(\theta_1)} \gamma \int_{\theta_1}^{\bar{\theta}_1} \left[\alpha \left(\mathbb{E}^{\tilde{\theta}^2 | r} \left[W(\tilde{\theta}^2; e^{ND}) \right] \right) + \beta \right] f_1(r) dr \\ - \alpha \psi'(e_2^{ND}(\theta_1)) \psi''(e_2^{ND}(\theta_1)) \text{Var}(\tilde{\varepsilon}_2), \quad (16)$$

²⁷We can show further, under the assumption that $\log \psi'$ is strictly concave on $(0, \bar{e})$, that any effort policy that satisfies these equations maximizes (14).

where²⁸

$$\begin{aligned} W(\theta^2; e^{ND}) &= W(\theta_1, \gamma\theta_1 + \varepsilon_2; e^{ND}) \\ &= \sum_{s=1}^2 \psi(e_s^{ND}(\theta_1)) + \int_{\theta_1}^{\theta_1} \sum_{s=1}^2 J_1^s \psi'(e_s^{ND}(r)) dr + \psi'(e_2^{ND}(\theta_1))[\varepsilon_2 - \mathbb{E}[\tilde{\varepsilon}_2]]. \quad \backslash \end{aligned}$$

As in the risk-neutral case, the choice of the optimal effort policy trades off two concerns: (1) limiting the agent's expected compensation and (2) maximizing expected cash flows. Contrary to the risk-neutral case, (1) now requires not only reimbursing the agent for his disutility of effort and providing him an informational rent to induce truth-telling, but also reducing the risk that the agent faces in his second-period compensation (note that, because we are restricting attention to effort policies that depend only on θ_1 and because θ_1 is known to the agent at the time of contracting, the agent does not face any risk concerning his effort—the risk he faces is entirely in terms of his compensation which is required to vary with cash flows for incentives reasons). The optimal period-1 effort is then given by (15) and trades off the marginal effect of increasing type θ_1 's effort on his compensation (the left hand side of (15)) with the marginal effect of this increase on the firm's cash flows (the first term in the right hand side of (15)) discounted by the fact that a higher effort for type θ_1 requires increasing the rent of all types above θ_1 , as captured by the second term in the right hand side of (15). These are exactly the same trade-offs as in the risk neutral case, adjusted for the fact that the agent's utility is now a concave function of his compensation (indeed, it is easy to see that Condition (15) reduces to Condition (13) for the risk-neutral case when $\alpha = 0$ and $\beta = 1$, i.e., when marginal utility becomes a constant.)

Also note that, at least when α is not too large compared to β , the policy $e_1^{ND}(\theta_1)$ is increasing in θ_1 (this is borne out in the specific cases we consider below). As in the risk-neutral case, this monotonicity 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 $e_1^{ND}(\theta_1)$ is decreasing in θ_1 . However, contrary to the risk-neutral case, distortions do not vanish “at the top”. In fact, when applied to $\bar{\theta}_1$, (15) becomes

$$\psi'(e_1^{ND}(\bar{\theta}_1)) \left[\alpha \left(\mathbb{E}^{\tilde{\theta}^2 | \bar{\theta}_1} \left[W(\tilde{\theta}^2; e) \right] \right) + \beta \right] = 1. \quad (17)$$

Now let $\mathcal{W} = \mathcal{V}^{-1}$ and note that the efficient level of effort is constant over time (and histories) and is implicitly given by²⁹

$$\psi'(e^{FB}) \mathcal{W}'(T\psi(e^{FB})) = 1 \quad (18)$$

²⁸Note that, while we continue to denote W as a function of the entire history θ^2 to be consistent with the notation introduced in the previous section, in this example where effort is restricted to depend only on θ_1 , W de facto depends only on θ_1 .

²⁹The FB choice of effort equates the marginal cost of maintaining constant the agent's utility, given by the left-hand side of (18), with the marginal effect of effort on cash flows, as given by the right-hand-side.

which, when applied to the example here, becomes

$$\psi'(e^{FB}) [\alpha (2\psi(e^{FB})) + \beta] = 1.$$

Because $\mathbb{E}^{\bar{\theta}^2|\bar{\theta}_1} [W(\tilde{\theta}^2; e^{FB})] > 2\psi(e^{FB})$, optimal effort is distorted below the efficient level also at the top. The reason is that, with risk aversion, the compensation that the principal must pay to type $\bar{\theta}_1$ to discourage him from mimicking lower types reduces $\bar{\theta}_1$'s marginal utility of money. In turn, this increases the cost for the principal of asking for higher effort, which explains why at the optimum the principal finds it optimal to distort type $\bar{\theta}_1$'s effort downwards relative to the efficient level (see also Battaglini and Coate, 2008, for a similar result in a two-type model).

Next consider the optimal period-2 effort, which is given by (16). The key difference between (16) and (15) is the last term on the right-hand side of (16). This term captures the principal's concern about exposing the agent to the risk associated with the uncertainty that the agent faces about his second period's productivity, as perceived at the time of contracting. Other things equal, this term contributes to reducing effort, as anticipated in the Introduction.

It is helpful to illustrate these effects graphically for particular parameters. To this end, suppose that θ_1 is drawn from a uniform distribution over $[\frac{1}{2}, \frac{3}{2}]$ and that ε_2 is drawn from a uniform distribution on $[-\frac{1}{2}, \frac{1}{2}]$. In addition, let $\psi(e) = \frac{e^2}{4}$ for $e \in (0, \bar{e})$ and $\alpha = \beta = 1$. The optimal shock-independent effort policies e^{ND} are calculated using polynomial approximations and depicted in Figure 1. Note that these policies satisfy the single-crossing conditions of Proposition 3 and hence are implementable with pseudo-linear and linear schemes.

Consider the case of imperfect persistence, i.e., $\gamma = \frac{1}{2}$. While a form of seniority continues to hold ($e_2^{ND}(\theta_1)$ is higher than $e_1^{ND}(\theta_1)$ for most values of θ_1), risk aversion tends to depress e_2 , thus reducing the optimality of reward schemes that assign to managers with a longer tenure more high powered incentives. Furthermore, now there exist values of θ_1 for which $e_2^{ND}(\theta_1) < e_1^{ND}(\theta_1)$. To see the reason for this, note that, when evaluated at $\bar{\theta}_1$, equations (15) and (16) are symmetric except for the term $-\alpha\psi''(e_2^{ND}(\bar{\theta}_1))\psi'(e_2^{ND}(\bar{\theta}_1))Var(\tilde{\varepsilon}_2)$. As anticipated above, this term captures the additional cost associated with a high second-period effort, stemming from the fact that a higher effort requires a higher sensitivity of compensation to performance and hence a higher volatility of the agent's compensation. To better appreciate where this term comes from, recall, from Proposition 1, that incentive-compatibility requires that the total compensation that the agent receives in each state (θ_1, θ_2) — equivalently, $(\theta_1, \varepsilon_2)$ — be given by

$$\mathcal{W}(W(\theta_1, \gamma\theta_1 + \varepsilon_2; e)) = \mathcal{W} \left(\begin{array}{c} \sum_{t=1}^2 \psi(e_t^{ND}(\theta_1)) + \int_{\theta_1}^{\theta_1} \left[\sum_{t=1}^2 J_1^t \psi'(e_t^{ND}(r)) \right] dr \\ + \psi'(e_2^{ND}(\theta_1))[\varepsilon_2 - \mathbb{E}[\tilde{\varepsilon}_2]] \end{array} \right). \quad (19)$$

It is then immediate that reducing $e_2^{ND}(\theta_1)$ 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 effect associated with




Figure 1: Shock-independent effort policies for a risk-averse agent when $\gamma = 0, 1/2$ and 1 .

a decline in the impulse response, as documented in the risk neutral case (here captured by the fact that $J_1^1 = 1 > J_1^2 = 1/2$), thus resulting in $e_2^{ND}(\theta_1) < e_1^{ND}(\theta_1)$.

Next, consider the case where first and second-period productivities are independent, i.e. $\gamma = 0$. When the agent is risk neutral (Example 2), second-period effort does not depend on the productivity realization and is equal to the first-best level. As explained in the previous subsection, the reason is that, when types are independent, distorting effort in the second period does not have any effect on rent extraction. Things are different under risk aversion. In this case, the agent's responsiveness to second-period incentives is influenced by the level of compensation he received (or was promised) in period one as an informational rent to induce him to truthfully reveal his productivity (as captured by the integral term in (19)). Because agents with a higher period-1 productivity receive higher information rents, this in turn implies that incentivizing period-2 effort from them is more costly than incentivizing the same level of effort from less productive agents. As a result, the optimal period-2 effort policy is decreasing in θ_1 when types are independent, as shown in the second graph in Figure 1 — This can also be seen from the Euler condition (16) by noticing that expected utility $\mathbb{E}^{\tilde{\theta}^2|\theta_1} \left[W(\tilde{\theta}^2; e) \right]$ is increasing in θ_1 .

Finally, consider the case where first-period productivity is perfectly persistent, i.e. $\gamma = 1$. In this case, the only difference in the trade-offs determining optimal effort in later periods relative to the early periods is that the agent faces more risk with respect to his payments (since future productivity is less certain at the time of contracting). Thus, while under risk neutrality the principal would maintain effort constant over time, for all θ_1 (see Example 3), under risk aversion effort is uniformly lower in later periods than in early ones. That is, we have a complete reversal of the seniority effect identified in the previous subsection.

Another way the principal could mitigate the risk stemming from variations in the agent's compensation due to shocks to productivity is to respond to such shocks by changing the level of effort requested of the agent. This exposes the agent to new risk coming from the volatility of his effort (recall that the disutility of effort ψ is convex, meaning that the agent dislikes such volatility) but it helps reduce the volatility in compensation. For example, by asking the agent to exert lower effort when his period-2 shock to productivity is high, the principal can reduce the volatility in period-2 cash flows and hence in period-2 compensation.

To illustrate such a possibility, we now return to the general environment (still assuming that T is finite, no discounting, and no preferences for consumption smoothing) and consider the fully-optimal effort policies. Using variational arguments, we can establish the following result.

Proposition 5 *Suppose that the agent is risk-averse. Any effort policy e^* that maximizes the principal's payoff (14) satisfies the following Euler equation, for each t and F -almost every history of productivities θ^t .³⁰*

$$\begin{aligned}
& \psi'(e_t^*(\theta^t)) \mathbb{E}^{\tilde{\theta}^T | \theta^t} \left[\mathcal{W}'(W(\tilde{\theta}^T; e^*)) \right] \\
= & 1 - \psi''(e_t^*(\theta^t)) \frac{J_1^t(\theta^t)}{f_1(\theta_1)} \int_{\theta_1}^{\bar{\theta}_1} \mathbb{E}^{\tilde{\theta}^T | r} \left[\mathcal{W}'(W(\tilde{\theta}^T; e^*)) \right] f_1(r) dr \\
& - \psi''(e_t^*(\theta^t)) \sum_{m=2}^t \frac{J_m^t(\theta^t)}{f_m(\theta_m | \theta^{m-1})} \int_{\theta_m}^{\bar{\theta}_m(\theta^{m-1})} \left(\begin{array}{c} \mathbb{E}^{\tilde{\theta}^T | (\theta^{m-1}, r)} \left[\mathcal{W}'(W(\tilde{\theta}^T; e^*)) \right] \\ - \mathbb{E}^{\tilde{\theta}^T | \theta^{m-1}} \left[\mathcal{W}'(W(\tilde{\theta}^T; e^*)) \right] \end{array} \right) f_m(r | \theta^{m-1}) dr.
\end{aligned} \tag{20}$$

The above Euler equation sets the direct cost of compensating the agent for additional period- t effort after the history θ^t (the left hand side) equal to the marginal benefit of the associated effort in terms of cash flows (i.e., 1), reduced by the extra marginal compensation that the principal must provide to the agent to discourage him from misrepresenting the sequence of productivities in question.

Interestingly, note that, contrary to the risk neutral case, optimal effort in period t now depends not only on the impulse response of θ_t to θ_1 , as captured by $J_1^t(\theta^t)$, but also of the impulse responses of θ_t to all intermediate types, $J_m^t(\theta^t)$, $m = 2, \dots, t-1$.

To illustrate, consider again the case where $T = 2$ and where the productivity process is as in Example 4 above (i.e., $J_1^2(\theta^2) = \gamma$). The Euler condition for second-period effort then becomes

$$\begin{aligned}
& \psi'(e_2^*(\theta_1, \gamma\theta_1 + \varepsilon_2)) \mathcal{W}'(W(\theta_1, \gamma\theta_1 + \varepsilon_2; e^*)) \\
= & 1 - \psi''(e_2^*(\theta_1, \gamma\theta_1 + \varepsilon_2)) \frac{\gamma}{f_1(\theta_1)} \int_{\theta_1}^{\bar{\theta}_1} \mathbb{E}^{\tilde{\varepsilon}_2} \left[\mathcal{W}'(W(r, \gamma r + \tilde{\varepsilon}_2; e^*)) \right] f_1(r) dr \\
& - \frac{\psi''(e_2^*(\theta_1, \gamma\theta_1 + \varepsilon_2))}{g_2(\varepsilon_2)} \int_{\varepsilon_2}^{\bar{\varepsilon}_2} \left(\begin{array}{c} \mathcal{W}'(W(\theta_1, \gamma\theta_1 + r; e^*)) \\ - \mathbb{E}^{\tilde{\varepsilon}_2} \left[\mathcal{W}'(W(\theta_1, \gamma\theta_1 + \tilde{\varepsilon}_2; e^*)) \right] \end{array} \right) g_2(r) dr.
\end{aligned} \tag{21}$$

³⁰ As in the risk neutral case, these conditions apply to all points where effort is interior, i.e., $e_t > 0$.

Equation (21) can be used to determine how the fully-optimal second-period effort policy depends on the shock ε_2 . There are three effects worth noticing. First, holding constant the effort levels and inspecting the left-hand side of the equation, one can see that a higher value of the shock ε_2 implies a higher total monetary utility

$$\begin{aligned}
W(\theta_1, \gamma\theta_1 + \varepsilon_2; e^*) &\equiv \psi(e_1^*(\theta_1)) + \psi(e_2^*(\theta_1, \gamma\theta_1 + \varepsilon_2)) \\
&+ \int_{\theta_1}^{\theta_1} \mathbb{E}^{\tilde{\varepsilon}_2} \left[\sum_{t=1}^2 J_1^t \psi'(e_t^*(r, \gamma r + \tilde{\varepsilon}_2)) \right] dr \\
&+ \int_{\varepsilon_2}^{\varepsilon_2} \psi'(e_2^*(\theta_1, \gamma\theta_1 + r)) dr - \mathbb{E}^{\tilde{\varepsilon}_2} \left[\int_{\varepsilon_2}^{\tilde{\varepsilon}_2} \psi'(e_2^*(\theta_1, \gamma\theta_1 + r)) dr \right]
\end{aligned} \tag{22}$$

that must be provided to the agent to induce him to exert effort and reveal his productivity. Because of the concavity of \mathcal{V} , the marginal cost of increasing the agent's utility (in terms of extra compensation) increases with W . In turn, because higher effort requires increasing the agent's monetary utility, this new effect unambiguously contributes to a lower effort for higher values of the shock, that is, to a policy $e_2(\theta_1, \gamma\theta_1 + \varepsilon_2)$ that is decreasing in ε_2 .

A second effect is captured by the second term of the right-hand side of (21). This term is the marginal cost of increasing the information rent that the principal must give to each type $\check{\theta}_1 > \theta_1$ to preserve incentives for truthful reporting in period one. As in the case of risk neutrality, the size of this effect is determined by the persistence of the first-period productivity, as captured by the impulse response $J_1^2 = \gamma$ of θ_2 to θ_1 . As in the risk-neutral case, the size of this effect is decreasing in θ_1 (since the measure of type above θ_1 decreases with θ_1) and independent of the shock realization ε_2 .

The most subtle effect is the one captured by the third term in the right hand side of (21). This term captures the marginal cost (for the principal) of changing the distribution of second-period rents. As in the risk-neutral case, a higher effort at history $(\theta_1, \gamma\theta_1 + \varepsilon_2)$ requires increasing the rent for all productivity shocks $\check{\varepsilon}_2 > \varepsilon_2$, as required to preserve incentives for truthful reporting in period two. Such a change is costless under risk neutrality, since, by appropriately adjusting the payments, the expectation of the agent's total rents need not vary. In fact, the only effect that matters, under risk neutrality, is how the agent's expected rents vary with period-one productivity θ_1 , which is already captured by the second term in (21). In contrast, under risk aversion, increasing the variation in the agent's rents with respect to the shock to his productivity increases the corresponding risk in compensation, whose cost is ultimately born by the principal. This is true even though the agent's expected rents at the time of contracting, conditional on having productivity θ_1 , do not change. The cost for the principal of a marginal increase in effort at history $(\theta_1, \gamma\theta_1 + \varepsilon_2)$, in terms of extra volatility in total rents, is given by

$$\int_{\varepsilon_2}^{\tilde{\varepsilon}_2} [\mathcal{W}'(W(\theta_1, \gamma\theta_1 + r; e^*)) - \mathbb{E}^{\tilde{\varepsilon}_2} [\mathcal{W}'(W(\theta_1, \gamma\theta_1 + \tilde{\varepsilon}_2; e^*))]] g_2(r) dr. \tag{23}$$




Figure 2: Fully-optimal second-period effort policy for $\gamma = 1/2$

As one can see by inspecting (23), this term is always positive provided that $W(\theta_1, \gamma\theta_1 + \varepsilon_2; e^*)$ is increasing in ε_2 (i.e., provided optimal effort does not decrease too fast). The effect is, however, non-monotone in the productivity shock ε_2 and vanishes at both the bottom ($\varepsilon_2 = \underline{\varepsilon}_2$) and at the top ($\varepsilon_2 = \bar{\varepsilon}_2$) of the distribution of the shock. This effect can thus contribute to making the optimal effort policy $e_2^*(\theta_1, \cdot)$ non-monotone in the shock ε_2 , even if the impulse responses are invariant in ε_2 , as in the example considered here (compare this to the case of a risk-neutral agent, where, for constant impulse responses, optimal second-period effort is constant).

As an illustration of how the various effects discussed above combine to shape the optimal effort policy, consider the parameters of the numerical example provided above. The fully-optimal second-period effort policy is given in Figure 2 for $\gamma = 1/2$. In this example, second-period effort is always decreasing in the shock ε_2 , for each θ_1 , although (as noted above) we expect this finding may be quite sensitive to the parameters. Nonetheless, the example serves to document how optimal effort may vary with shocks to the productivity, even if impulse responses are constant (c.f., the case with a risk-neutral agent).

Because the effort policy $e_2^*(\theta_1, \gamma\theta_1 + \varepsilon_2)$ is decreasing in ε_2 , and hence in $\theta_2 = \gamma\theta_1 + \varepsilon_2$, it is easy to see that the single-crossing conditions of Proposition 3 are violated for $t = 2$, making it difficult to sustain the optimal policy with the linear and pseudo-linear schemes of Section 4.1. Nonetheless, the single-crossing conditions of Proposition 2 hold, meaning that the policy can be implemented using bonus schemes as in Proposition 1. (To see this, note that implementability under these schemes requires that second-period cash flows $\pi_2(\theta_1, \gamma\theta_1 + \varepsilon_2)$ be nondecreasing in ε_2 , which is true if $\partial e_2(\theta_1, \theta_1 + \gamma\varepsilon_2)/\partial \varepsilon_2 \geq -1$, a property that holds under the optimal policy).

Next, consider the findings for the optimality of seniority-based schemes identified above by looking at shock-independent policies. Note that the Euler equations (15) and (16) for the restricted policies are identical to the Euler equations (20) for the fully optimal policies after restricting effort to depend only on θ_1 and, for (16), after taking expectations with respect to ε_2 . It is therefore unsurprising that, at least in the examples we consider numerically, the fully-optimal first-period

policy $e_1^*(\theta_1)$ and the expectation of the fully-optimal second-period policy $\mathbb{E}_{\tilde{\varepsilon}_2} [e_2^*(\theta_1, \tilde{\varepsilon}_2)]$ are very close to the restricted policies $e_1^{ND}(\theta_1)$ and $e_2^{ND}(\theta_1)$, exhibiting the same properties as described above.

More generally, we expect the following principles regarding the characteristics of the optimal effort policy to apply to a broad range of settings.

Claim 1 *Optimal effort in all periods is distorted relative to the first-best level, even when the initial productivity level is “at the top”, i.e. $\theta_1 = \bar{\theta}_1$.*

Claim 2 *When impulse responses decline over time, expected effort (given period-one productivity) tends to increase over time when the initial productivity is low but tends to decrease over time when the initial productivity is high.*

Claim 3 *When the process is fully-persistent (i.e., when impulse responses are equal to one), expected effort tends to decrease over time for all values of the initial productivity.*

Claim 4 *Optimal effort, at least after the initial period, may decline in past and current productivity realizations. This effect helps shield the agent from risk and, in particular, does not require variation in the impulse responses.*

While we did not formally prove that these properties hold for arbitrary settings, we verified that they hold in a range of numerical settings and checked the consistency of the results with the analytical functional forms of the Euler conditions.

Perhaps surprisingly, note that, contrary to the risk-neutral case, the distortion relative to the efficient policies need not be downwards for all histories. For example, consider the environment studied above and suppose that $T = 2$, with types independent over time (i.e., $\gamma = 0$). Suppose that the optimal effort policy e^* maximizes (14). Consider the history $(\underline{\theta}_1, \underline{\varepsilon}_2)$ where the agent’s productivity is at its minimum in both periods and suppose, towards a contradiction, that $e_t^*(\theta^t) \leq e^{FB}$ all $t = 1, 2$, all θ^t , where e^{FB} is given by (18). As one can see from (22), in this case, the monetary utility that the agent receives from his compensation is

$$W((\underline{\theta}_1, \underline{\varepsilon}_2); e^*) \equiv \psi(e_1^*(\underline{\theta}_1)) + \psi(e_2^*(\underline{\theta}_1, \underline{\varepsilon}_2)) - \mathbb{E}^{\tilde{\varepsilon}_2} \left[\int_{\underline{\varepsilon}_2}^{\tilde{\varepsilon}_2} \psi'(e_2^*(\underline{\theta}_1, s)) ds \right] ds.$$

This is less than the monetary utility $2\psi(e^{FB})$ that he would receive if the efficient policies could be implemented at first-best cost for the principal. This in turn implies that the marginal cost for the principal of increasing effort at this history is less than at the first-best level, contradicting the optimality of setting $e_2^*(\underline{\theta}_1, \underline{\varepsilon}_2) \leq e^{FB}$. One can see this formally from the necessary condition (21), which becomes

$$\psi'(e_2^*(\underline{\theta}_1, \underline{\varepsilon}_2)) \mathcal{W}'(W((\underline{\theta}_1, \underline{\varepsilon}_2); e^*)) = 1.$$

Assuming that $e_1^*(\underline{\theta}_1) \leq e^{FB}$, the choice of effort at $(\underline{\theta}_1, \underline{\varepsilon}_2)$ must exceed e^{FB} . By continuity, a similar property applies to histories in a neighborhood of $(\underline{\theta}_1, \underline{\varepsilon}_2)$. The intuition for upward distortions at these histories comes from the fact that the compensation promised under the optimal contract is so low that it makes it relatively cheap for the principal to incentivize higher effort.

Before turning to an environment where the agent has preferences for consumption smoothing, we discuss the implications of the results above for what is often informally referred to as the “*power of incentives*”. The latter measures the sensitivity of compensation to performance, here captured by cash flows. The fact that, with risk aversion, the optimal effort policy may require using a compensation scheme where payments are non-linear in cash flows raises the difficulty of how to measure changes in the power of incentives over time. One way to proceed is by considering the local sensitivity of pay to performance, i.e., the sensitivity of pay in a neighborhood of the cash flows that the mechanism asks the agent to deliver. First observe that any policy which is implementable by a bonus scheme of the type defined in Proposition 1 is also implementable by a compensation scheme where the total compensation to the agent is provided entirely in period T , i.e., for all $t < T$ and all (θ^t, π^t) , $s_t(\theta^t, \pi^t) = 0$, with $s_T(\theta^T, \pi^T)$ differentiable in a neighborhood of the equilibrium cash flows $\pi^T = \pi^T(\theta^T)$. We then have that a necessary condition for the agent to choose effort obediently is that, for each t and each history $\theta^t \in R^t$,

$$\psi'(e_t(\theta^t)) = \mathbb{E}^{\tilde{\theta}^T | \theta^t} \left[\mathcal{V}' \left(\mathcal{W} \left(W(\tilde{\theta}^T; e^*) \right) \frac{\partial s_T(\tilde{\theta}^T; \pi^T(\tilde{\theta}^T))}{\partial \pi_t} \right) \right]. \quad (25)$$

That is, the marginal disutility of effort at the desired effort level $e_t(\theta^t)$ must be equal to the expected marginal effect of a higher effort on the utility the agent derives from compensation. Using the fact that $\mathcal{W} = \mathcal{V}^{-1}$, one way Condition (25) can be satisfied is by letting

$$\frac{\partial s_T(\theta^T; \pi^T(\theta^T))}{\partial \pi_t} = \mathcal{W}'(W(\theta^T; e^*)) \psi'(e_t(\theta^t))$$

for each $\theta^T \in R$, each t , which yields

$$\mathbb{E}^{\tilde{\theta}^T | \theta^t} \left[\frac{\partial s_T(\tilde{\theta}^T; \pi^T(\tilde{\theta}^T))}{\partial \pi_t} \right] = \mathbb{E}^{\tilde{\theta}^T | \theta^t} \left[\mathcal{W}'(W(\tilde{\theta}^T; e^*)) \psi'(e_t(\theta^t)) \right]. \quad (26)$$

We interpret the left-hand side of (26) to be the *expected power of incentives* at the truthful and obedient history $(\theta^t, \pi^t(\theta^t))$. That is, conditional on having being truthful and obedient in the past, and on planning to remain truthful and obedient in the future, the left-hand side of (26) measures how the agent’s expected compensation changes (locally) with his period- t performance. Also note that, in the risk neutral case, this measure of the power of incentives coincides with the slope $\alpha_t(\theta^t)$ of the linear scheme, as defined in the previous subsection.

To understand the dynamics of the optimal power of incentives, it is helpful to observe that the

right-hand side of (26) coincides with the left-hand side of the Euler condition (20). Thus, one can see that there are similar seniority effects as for the optimal effort policy. In particular, Principles 2 and 3 also apply to the dynamics of the power of incentives.

Finally, note from (20) that there are various histories at which the expected power of incentives is equal to one. For example, with $T = 2$, $\bar{\theta}_1$, $(\bar{\theta}_1, \underline{\varepsilon}_2)$ and $(\bar{\theta}_1, \bar{\varepsilon}_2)$ are such histories. Whilst with a risk-neutral agent, a power of incentives equal to one (which is the marginal effect of effort on cash flows) implies that the agent chooses the efficient effort level, this is no longer true when the agent is risk averse. The reason, as is clear from the discussion above, is that the total payment the agent receives (or expects to receive) differs from that for the efficient mechanism (i.e., $2\psi(e^{FB})$). Thus one should be careful not to associate (expected) power of incentives equal to one with full efficiency. Similarly, the econometrician should not attempt to measure inefficiencies simply by examining the power of incentives.

Consumption smoothing.

We now turn to the case where the agent has preferences for intertemporal consumption smoothing. For simplicity, and without any serious loss, we assume that each v_t is strictly concave, while \mathcal{V} is linear. To facilitate the comparison with the case without consumption smoothing examined above, we also maintain the assumption that T is finite and drop discounting (neither of these simplifications drives the results). Because the agent's payoff now depends on the timing of payments, the principal must distribute the payments optimally over time, adding an additional dimension to the problem. One way to arrive to the optimal policy e^* is the following. Using Proposition 1 along with Corollary 1, one determines the cost for the principal of sustaining any effort policy e . Once this is accomplished, one then chooses the policy e^* that maximizes the principal's expected cash flows, net of the compensation to the agent.

The Euler condition (20) now becomes

$$\begin{aligned} & \psi'(e_t^*(\theta^t)) w_t'(u_t(\theta^t)) \\ = & 1 - \psi''(e_t^*(\theta^t)) \frac{J_1^t(\theta^t)}{f_1(\theta_1)} \int_{\theta_1}^{\bar{\theta}_1} w_1'(u_1(r)) f_1(r) dr \\ & - \psi''(e_t^*(\theta^t)) \sum_{m=2}^t \frac{J_m^t(\theta^t)}{f_m(\theta_m|\theta^{m-1})} \int_{\theta_m}^{\bar{\theta}_m(\theta^{m-1})} \{w_m'(u_m(\theta^{m-1}, r)) - w_{m-1}'(u_{m-1}(\theta^{m-1}))\} f_m(r|\theta^{m-1}) dr, \end{aligned} \quad (27)$$

where, for each t and each $\theta^t \in R^t$, $u_t(\theta^t) = v_t(c_t(\theta^t))$, and $w_t = v_t^{-1}$. Note that the “inverse Euler equation” of Corollary 1 can be written as

$$w'_{m-1}(u_{m-1}(\theta^{m-1})) = \mathbb{E}^{\tilde{\theta}^m|\theta^{m-1}} \left[w'_m(u_m(\tilde{\theta}^m)) \right]$$

for each $2 \leq m \leq T$ and each θ^{m-1} . This implies, for instance, that each term in the final sum in the right-hand-side of (27) shrinks to zero at $\underline{\theta}_m(\theta^{m-1})$, for each $m = 2, \dots, t$, each $\theta^{m-1} \in R^{m-1}$, which is consistent with the case without consumption smoothing as discussed above.

Given (27), we expect that the trade-offs determining the optimal effort policies are similar to those discussed above for the non-consumption-smoothing case. Indeed, these conjectures are verified in the numerical cases discussed above with respect to Example 4 with $v_t(c) = \frac{1}{\alpha}\sqrt{2\alpha c + \beta^2} - \frac{\beta}{\alpha}$ for $t = 1, 2$ and $\alpha, \beta > 0$, and with $\mathcal{V}(c) = c$. With identical parameter values, we find that the qualitative properties of the optimal effort policies remain essentially the same as in the case without consumption smoothing.

6 Related literature

The literature on managerial compensation is obviously too large to be summarized within the context of this paper. We refer the reader to Prendergast (1999) for an excellent overview and to Edmans and Gabaix (2009) for a survey of some recent developments. Below, we limit our discussion to the work that is mostly related to our paper.

Particularly related is the empirical literature on the use of seniority-based compensation schemes. This literature finds mixed evidence as to the effect of tenure on performance-based pay. Early papers suggested that managers with a longer tenure tend to have weaker incentives and explained this by the possibility that the board of directors tends to become captured by CEOs over time (see, e.g., Hill and Phan, 1991). More recent papers provide evidence for the contrary (e.g., Gibbons and Murphy, 1992, Lippert and Porter, 1997, and Cremers and Palia, 2010). These differences often originate in the choices about which incentive measures are relevant (e.g., whether or not to consider stock options). Our paper contributes to this literature by offering a completely novel explanation for the optimality (or suboptimality) of seniority-based schemes based on the interaction between the persistence of the managers' private information and their attitude towards risk.

The paper is also related to the literature on “dynamic moral hazard” and to its application to managerial compensation. Seminal works in this literature include Lambert (1983), Rogerson (1985) and Spear and Srivastava (1987). These works provide 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) characterizes the optimal policy in a continuous-time setting with Brownian shocks. In contrast to these works, 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 and defined over consumption net of the disutility of effort. Under these assumptions, the optimal contract takes the form of a simple linear aggregator of aggregate profits.

Contrary to the above works, in the current paper we assumed that, in each period, the agent observes the (persistent) shock to his productivity before choosing effort.³¹ In this respect, the

³¹While we abstracted from the possibility that performance is also affected by transitory noise, in many cases of interest, the contracts we characterized continue to implement the desired effort policies even when performance is affected by transitory shocks that are observed by the agent only after committing his effort.

paper is closely related to Laffont and Tirole (1986) who first proposed this alternative timing. This alternative approach permits one to use techniques from the mechanism design literature to solve for the optimal contract. The same approach has been recently applied to dynamic managerial compensation by Edmans and Gabaix (2010). Relative to their work, our contribution is twofold: (i) We provide conditions for the implementability of effort policies that respond to shocks to managerial productivity and (ii) we characterize the contracts that implement such policies at minimum cost for the principal. We then use this characterization to determine the properties of the profit-maximizing effort policy.³² Allowing for general processes and for time-varying effort policies is instrumental to our results about the dynamics of the power of incentives, the optimality of linear and pseudo-linear schemes, and the effect of the agent's risk aversion on the dynamics of the power of incentives.

The paper is also related to our work on managerial turnover in a changing world (Garrett and Pavan, 2010). In that paper we assumed that all agents are risk neutral and focused on the dynamics of retention decisions. In contrast, in the present paper, we abstracted from retention (i.e., assumed a single agent) and focused instead on the effect of risk-aversion on the dynamics of effort and on the optimality of seniority-based compensation schemes.

Related is also the literature on the optimal use of financial instruments in dynamic principal-agent relationships. For instance, DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007), Sannikov (2007),³³ and Biais et al. (2010) 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 compensation 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 over time and output realizations, so is the equity share. In contrast, we provide an explanation for why and how this share may change over time. While these papers suppose that 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 rate that increases with the balance.³⁴

From a methodological standpoint, in this paper we applied recent results from the dynamic mechanism design literature to arrive to the characterization of optimal contracts. In particular, the

³²Two other differences are that we allow for the possibility that (i) the agent is privately informed at the time he signs the contract, and (ii) has preferences for consumption smoothing.

³³As in our work, and contrary to the other papers cited here, Sannikov (2007) allows the agent to possess private information prior to signing the contract. Assuming the agent's initial type can be either "bad" or "good", he characterizes the optimal separating menu where only good types are funded.

³⁴Other recent papers that consider persistent private information are He (2008), Edmans, Gabaix, Sadzik, and Sannikov (2010), Strulovici (2010), and Williams (2011). Contrary to these papers and the current one, DeMarzo and Sannikov (2008), Bergemann and Hege (1998, 2005), Horner and Samuelson (2009), and Garfagnini (2011), consider an environment in which both the investors and the agent learn about the firm's true productivity over time and where the agent's beliefs about the likely success of the project differ from the investors' only in case the agent deviates, in which case the divergence in beliefs may be persistent.

approach here builds on the techniques developed by Pavan, Segal, and Toikka (2011). That paper provides a general treatment of incentive compatibility in dynamic settings. It extends previous work by Besanko (1985), Courty and Li (2000), Battaglini (2005), Eso and Szentes (2007), and Kapicka (2008), among others, by allowing for significantly more general payoffs and stochastic processes and by identifying the role of impulse responses as the key driving force for the dynamics of optimal contracts.³⁵

An important dimension in which the current paper makes progress is the characterization of optimal dynamic mechanisms for risk-averse agents with information that is correlated over time.³⁶ In this respect, the paper is also related to the literature on optimal dynamic taxation (also known as Mirrleesian taxation, or new public finance). Battaglini and Coate (2008) consider a discrete-time-two-type model with Markov transitions and show continuity of the optimal mechanism with respect to the agent’s degree of risk aversion. Zhang (2009) considers a model with finitely many types, but where contracting occurs in continuous time and where the arrival rate of the transitions between types follows a Poisson process. For most of the analysis, he restricts attention to two types and finds that many of the results derived for the i.i.d. case (e.g., 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, a result which is consistent with the findings in the current paper. Consistent with Battaglini and Coate (2008), he also finds that, contrary to the risk-neutral case, distortions do not vanish as soon as the agent becomes a high type. Building on the first-order approach developed in Pavan, Segal, and Toikka (2011), more recently Fahri and Werning (2010) and Golosov, Troshkin, and Tsyvinski (2010) study the dynamics of taxes in a setting with a continuum of types drawn from an AR(1) process. Our results appear broadly consistent with the aforementioned findings; however, by considering more general processes and payoffs, we shed light on properties of optimal contracts not identified before (in particular on the effect of risk-aversion on the optimality of seniority-based schemes).

7 Conclusions

We investigated the properties of optimal compensation schemes in an environment in which managerial ability to generate profits changes over time and is the managers’ private information. We showed how optimal contracts can be obtained as a solution to a dynamic mechanism design problem. We then used the solution (i) to shed light on the properties of compensation schemes that sustain the desired effort policy at minimal cost for the firm, (ii) to identify conditions for a given effort policy to be implementable, and (iii) to derive the properties of profit-maximizing effort policies.

³⁵We refer the reader to that paper for a more extensive review of the dynamic mechanism design literature.

³⁶For static models with risk aversion, see Salanie (1990), and Laffont and Rochet (1998).

When the manager is risk-neutral, we showed that it is typically optimal for the firm to offer a reward package whose power of incentives increases over time, thus inducing the manager to exert, on average, more effort as his tenure in the firm grows. This result hinges on the joint assumptions that (i) the manager has private information about his productivity at the time he is hired, and (ii) that the effect of the manager’s initial productivity on his productivity in the subsequent periods declines over time. Both assumptions seem reasonable for many cases of interest. The first assumption implies that the firm optimally reduces the power of incentives in the contracts offered to those types who are least productive at the initial contracting stage so as to discourage more productive types from mimicking. The second assumption implies that doing so has more effect on information rents when done in the early periods, when the effect of the initial private information is still pronounced, than in the later periods when such an effect has become small.

Building on the results for the risk-neutral case, we then showed that risk aversion reduces, and in some cases can even reverse, the profitability of such seniority-based schemes. Contracts where the power of incentives increases too much, on average, over time expose the manager to an undesirable level of risk. The reason is that the sensitivity of pay to performance is then highest precisely in those periods in which the risk that the manager faces about his productivity (and hence about his pay), as perceived at the time he is hired, is the highest.

We expect these insights to be useful also for other dynamic contracting problems, such as the design of optimal taxes in environments with persistent shocks to workers’ productivity, a problem that is at the frontier of the new public finance literature.

8 Appendix

Proof of Proposition 1. The proof proceeds in three steps and is based on two lemmas. Lemma 1 establishes that, in any incentive-compatible mechanism implementing the policy e , the agent’s value function at each history must satisfy a certain envelope condition. The second lemma uses this envelope condition iteratively to show that for F -almost all sequences θ^T , the utility that the agent derives from the compensation he receives from the principal must be equal to $W(\theta^T; e) + K$, where W is uniquely pinned down by the policy e and where K is a constant that must be nonnegative if the mechanism is individually rational for the agent. The last step verifies that, if a mechanism Ω exists that implements the policy e , then the one whose compensation is as defined in the proposition also implements the same policy at a (weakly) lower cost for the principal, thus establishing the result.

Step 1. Given any mechanism $\Omega = \langle \xi, s \rangle$, and any history h_t , let $V^\Omega(h_t)$ denote the agent’s value function at history h_t . This function is defined to be the supremum of the agent’s expected payoff over all possible reporting and effort strategies. We say that the mechanism Ω remains incentive compatible at history h_t if $V^\Omega(h_t)$ coincides with the agent’s expected payoff under a truthful and obedient strategy from t onwards. We then have the following result.

Lemma 1 Fix the history $h_{t-1} = (\theta^{t-1}, \hat{\theta}^{t-1}, e^{t-1})$.³⁷ If the mechanism Ω remains incentive compatible at $F_t(\cdot|\theta^{t-1})$ -almost all histories $h_t = (h_{t-1}, \theta_t)$,³⁸ then $V^\Omega(h_{t-1}, \cdot)$ is Lipschitz continuous in θ_t over $\text{Supp}[F_t(\cdot|\theta^{t-1})]$ and at each point $\theta_t \in \text{Supp}[F_t(\cdot|\theta^{t-1})]$ at which $V^\Omega(h_{t-1}, \cdot)$ is differentiable and Ω is incentive compatible

$$\frac{dV^\Omega(h_{t-1}, \theta_t)}{d\theta_t} = \mathbb{E}^{\tilde{\theta}^T|\theta^t} \left[\sum_{\tau=t}^T \delta^{\tau-1} J_t^\tau(\tilde{\theta}^\tau) \psi'(\xi_\tau((\hat{\theta}^{t-1}, \tilde{\theta}_{\geq t}^\tau), \pi^{\#\tau-1}(h_{t-1}, \tilde{\theta}_{\geq t}^{\tau-1}))) \right]$$

where $\pi^{\#\tau-1}(h_{t-1}, \theta_{\geq t}^{\tau-1}) = (\pi_s^\#)_{s=1}^{\tau-1}$ with $\pi_s^\# = \theta_s + e_s$ for all $s < t$ and $\pi_s^\# = \theta_s + \xi_s(\hat{\theta}^{t-1}, \theta_{\geq t}^s, \pi^{\#s-1})$ all $s \geq t$.

Proof of the lemma. Consider the following fictitious environment. At any period $s \geq t$, the agent can misreport his private information θ_s but is then “forced” to choose effort so as to perfectly “hide” his lies. That is, at any period $s \geq t$, and for any given sequence of reports $(\hat{\theta}^s)$, the agent must exert effort e_s so that the observed cash flow $\pi_s = \theta_s + e_s$ equals the one expected by the principal when the agent’s period- s reported type is $\hat{\theta}_s$ and he follows the principal’s recommended effort choice $\xi_s(\hat{\theta}^s, \pi^{s-1})$. This is to say that, at any period $s \geq t$, given the public history $(\hat{\theta}^s, \pi^{s-1})$ and the true productivity θ_s the agent must choose effort

$$e_s^\#(\theta_s; \hat{\theta}^s, \pi^{s-1}) = \hat{\theta}_s + \xi_s(\hat{\theta}^s, \pi^{s-1}) - \theta_s. \quad (28)$$

Now fix the period- $(t-1)$ history h_{t-1} . Given the reports $\hat{\theta}_{\geq t}^T$, let $(\hat{c}_{\geq t}^T, \hat{\pi}_{\geq t}^T)$ be the stream of payments and cash flows in the continuation game that starts at period t with history (h_{t-1}, θ_t) when in the continuation game the agent sends the reports $(\hat{\theta}_{\geq t}^T)$ and then follows the behavior described above (i.e., chooses effort according to (28)). For any sequence of reports $\hat{\theta}_{\geq t}^T$ and any sequence of true types $(\theta_{\geq t}^T)$, the agent’s payoff in this fictitious environment is given by

$$\mathcal{U}(\theta_{\geq t}^T; \hat{\theta}_{\geq t}^T, h_{t-1}) \equiv \sum_{s=t}^T \delta^{s-1} [\hat{c}_s - \psi(\hat{\theta}_s + \xi_s(\hat{\theta}^s, \hat{\pi}^{s-1}) - \theta_s)] + X(h_{t-1})$$

where $X(h_{t-1})$ is a function of the past history h_{t-1} . The assumptions that ψ is continuously differentiable with derivative bounded uniformly over E implies that \mathcal{U} is totally differentiable in $\theta_{\geq t}^s$, any $s \geq t$, and equi-Lipschitz continuous in $\theta_{\geq t}^T$ in the norm:

$$\|\theta_{\geq t}^T\| \equiv \sum_{s \geq t} \delta^{s-t} |\theta_s|.$$

Together with the assumptions that $\|\theta^T\| < k$ for all $\theta^T \in \Theta^T$ (which is implied by the assumption

³⁷The first element of h_{t-1} denotes the true productivity history, whereas the term $\hat{\theta}^{t-1}$ denotes the history of reports. The last term e^{t-1} denotes the history of effort choices.

³⁸Given the history $h_{t-1} = (\theta^{t-1}, \hat{\theta}^{t-1}, e^{t-1})$, (h_{t-1}, θ_t) denotes the history $((\theta^{t-1}, \theta_t), \hat{\theta}^{t-1}, e^{t-1})$.

that the sets Θ_t are bounded uniformly over t) and that the impulse responses $J_s^t(\theta^t)$ are also uniformly bounded, this means that the value function $V^\Omega(h_{t-1}, \cdot)$ is Lipschitz continuous in θ_t over Θ_t and hence also over $Supp[F_t(\cdot|\theta^{t-1})]$. This result follows from Pavan, Segal, and Toikka (2011) — Proposition 4.

Now suppose that, given h_{t-1} , the mechanism Ω remains incentive compatible at $F_t(\cdot|\theta^{t-1})$ -almost any history $h_t = (h_{t-1}, \theta_t)$ in the unrestricted game 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 game where he is forced to choose effort according to (28). The results in Pavan, Segal, and Toikka then imply that at any $\theta_t \in Supp[F_t(\cdot|\theta^{t-1})]$, at which $V^\Omega(h_{t-1}, \cdot)$ is differentiable and Ω is incentive compatible

$$\frac{dV^\Omega(h_{t-1}, \theta_t)}{d\theta_t} = \mathbb{E}^{\tilde{\theta}^T|\theta^t} \left[\sum_{\tau=t}^T \delta^{\tau-1} J_t^\tau(\tilde{\theta}^\tau) \frac{\partial \mathcal{U}(\tilde{\theta}_{\geq t}^T; \tilde{\theta}_{\geq t}^T, h_{t-1})}{\partial \theta_\tau} \right]$$

where $\partial \mathcal{U}(\theta_{\geq t}^T; \theta_{\geq t}^T, h_{t-1})/\partial \theta_\tau$ denotes the partial derivative of $\mathcal{U}(\theta_{\geq t}^T; \theta_{\geq t}^T, h_{t-1})$ with respect to the true type θ_τ under truthtelling. The result then follows from the fact that

$$\frac{\partial \mathcal{U}(\theta_{\geq t}^T; \theta_{\geq t}^T, h_{t-1})}{\partial \theta_\tau} = \psi'(\xi_\tau((\hat{\theta}^{t-1}, \theta_{\geq t}^\tau), \pi^{\#\tau-1}(h_{t-1}, \theta_{\geq t}^{\tau-1}))).$$

■

Step 2. Using the previous lemma inductively, then leads to the following result.

Lemma 2 *Let Ω be any incentive-compatible and individually rational mechanism implementing the policy e . Then for F -almost all $\theta^T \in \Theta^T$, \in*

$$\mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t(c_t(\theta^t)) \right) = W(\theta^T; e) + K$$

with $K \geq 0$.

Proof of the lemma. Using the result in the previous lemma and the fact that any mechanism Ω that is individually-rational and incentive-compatible for all θ_1 must remain incentive-compatible at F -almost all truthful histories³⁹ $\theta^t \in R^t$, any t , we then have that, for F -almost all truthful histories $\theta^t \in R^t$

$$V^\Omega(\theta^t) = V^\Omega(\theta^{t-1}, \underline{\theta}_t(\theta^{t-1})) + \int_{\underline{\theta}_t(\theta^{t-1})}^{\theta^t} \mathbb{E}^{\tilde{\theta}^T|\theta^{t-1}, s} \left[\sum_{\tau=t}^T \delta^{\tau-1} J_t^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] ds. \quad (29)$$

Furthermore, the fact that Ω is incentive compatible for all θ_1 , implies that, for any t , F -almost all

³⁹As explained also in the main text, a truthful history $h_t = (\theta^t, \hat{\theta}^{t-1}, e^{t-1})$ is one that is reached by reporting truthfully and following the principal's effort recommendations in each previous period. Because in any such history $\hat{\theta}_s = \theta_s$ and $e_s = \xi_s(\theta^s, \pi^{s-1}(\theta^{s-1}))$ all $s < t$, without risk of confusion, these histories can be conveniently denoted by the realized sequence of productivities θ^t .

truthful histories $\theta^{t-1} \in R^t$,

$$V^\Omega(\theta^{t-1}) = \mathbb{E}^{\tilde{\theta}_t | \theta^{t-1}} \left[V^\Omega(\theta^{t-1}, \tilde{\theta}_t) \right].$$

Combining, we have that

$$V^\Omega(\theta^{t-1}, \underline{\theta}_t(\theta^{t-1})) = V^\Omega(\theta^{t-1}) - \mathbb{E}^{\tilde{\theta}_t | \theta^{t-1}} \left[\int_{\underline{\theta}_t(\theta^{t-1})}^{\tilde{\theta}_t} \mathbb{E}^{\tilde{\theta}^T | \theta^{t-1}, s} \left[\sum_{\tau=t}^T \delta^{\tau-1} J_\tau^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] ds \right].$$

Applying (29) to $V^\Omega(\theta^{t-1})$ we then have that

$$\begin{aligned} V^\Omega(\theta^{t-1}, \underline{\theta}_t(\theta^{t-1})) &= V^\Omega(\theta^{t-2}, \underline{\theta}_{t-1}(\theta^{t-2})) \\ &+ \int_{\underline{\theta}_{t-1}(\theta^{t-2})}^{\theta^{t-1}} \mathbb{E}^{\tilde{\theta}^T | \theta^{t-2}, s} \left[\sum_{\tau=t-1}^T \delta^{\tau-1} J_{t-1}^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] ds \\ &- \mathbb{E}^{\tilde{\theta}_t | \theta^{t-1}} \left[\int_{\underline{\theta}_t(\theta^{t-1})}^{\tilde{\theta}_t} \mathbb{E}^{\tilde{\theta}^T | \theta^{t-1}, s} \left[\sum_{\tau=t}^T \delta^{\tau-1} J_t^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] ds \right]. \end{aligned} \quad (30)$$

Combining (29) with (30) we then have that

$$\begin{aligned} V^\Omega(\theta^t) &= V^\Omega(\theta^{t-2}, \underline{\theta}_{t-1}(\theta^{t-2})) + \int_{\underline{\theta}_{t-1}(\theta^{t-2})}^{\theta^{t-1}} \mathbb{E}^{\tilde{\theta}^T | \theta^{t-2}, s} \left[\sum_{\tau=t-1}^T \delta^{\tau-1} J_{t-1}^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] ds \\ &+ \delta^{t-1} H_t(\theta^t; e). \end{aligned}$$

Applying the same steps inductively to each $V^\Omega(\theta^{s-1}, \underline{\theta}_s(\theta^{s-1}))$, $s \leq t$, then leads to the conclusion that

$$\begin{aligned} V^\Omega(\theta^t) &= V^\Omega(\underline{\theta}_1) + \int_{\underline{\theta}_1}^{\theta_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{\tau=1}^T \delta^{\tau-1} J_1^\tau(\tilde{\theta}^\tau) \psi'(e_\tau(\tilde{\theta}^\tau)) \right] ds \\ &+ \sum_{s=2}^t \delta^{s-1} H_s(\theta^s; e). \end{aligned} \quad (31)$$

Using the fact that the agent's payoff under a truthful and obedient strategy must coincide with the value function at F -almost all histories and, in the case $T = \infty$, taking the limit for $t \rightarrow \infty$, then establishes the result, with $K = V^\Omega(\underline{\theta}_1)$. The fact that the mechanism is individually-rational for all θ_1 then implies that $K \geq 0$. ■

Step 3. Now it is easy to see that, if a mechanism Ω exists that implements the policy e , then e can also be implemented by a mechanism Ω' such that $s'_t(\theta^t, \pi^t) = c_t(\theta^t)$ if $\theta^t \in R^t$ and $\pi_t \geq \pi_t(\theta^t) \equiv \theta_t + e_t(\theta^t)$, and $s'_t(\theta^t, \pi^t) = -L_t$ with $L_t > 0$ large enough, otherwise [here $c_t(\theta^t) = s_t(\theta^t, \pi^t(\theta^t))$ is the equilibrium compensation under the original mechanism Ω]. The proof follows from essentially the same replication arguments that establish the Revelation Principle.

Having established that, in any mechanism Ω that implements the policy e , the utility the agent derives from his compensation must satisfy

$$\mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t(c_t(\theta^t)) \right) = W(\theta^T; e) \text{ for } F\text{-almost all } \theta^T,$$

by the definition of $c^{opt}(\cdot; e)$, it is then easy to see that any mechanism Ω^{opt} where the compensation scheme is such that $s_t(\theta^t, \pi^t) = c_t^{opt}(\theta^t; e)$ if $\theta^t \in R^t$ and $\pi_t \geq \pi_t(\theta^t) \equiv \theta_t + e_t(\theta^t)$, and $s_t(\theta^t, \pi^t) = -L_t$ otherwise, implements e at minimum cost. That it induces participation by all period-1 types follows from the fact that the payoff that each type θ_1 expects under a truthful and obedient strategy is given by

$$\int_{\underline{\theta}_1}^{\theta_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(e_t(\tilde{\theta}^t)) \right] ds$$

which is nonnegative under the assumption of first-order-stochastic dominance in types which implies that impulse responses are non-negative. ■

Proof of Corollary 1. The proof is by contradiction. Suppose $c^{opt}(\cdot; e)$ solves program (2) in Proposition 1 and assume that there exists a period t and an F -positive probability set $Q \subset R^t$ such that, for $\theta^t \in Q$,

$$\frac{1}{v_t'(c_t^{opt}(\theta^t; e))} > \mathbb{E}^{\tilde{\theta}^{t+1} | \theta^t} \left[\frac{\delta}{v_{t+1}'(c_{t+1}^{opt}(\tilde{\theta}^{t+1}; e))} \right].$$

The argument for the case where the inequality is reversed is symmetric. Then consider the following alternative policy $c^\#(\cdot; e)$. For any $\tau \neq t, t+1$, any $\theta^\tau \in R^\tau$, $c_\tau^\#(\theta^\tau; e) = c_\tau^{opt}(\theta^\tau; e)$; in period t ,

$$c_t^\#(\theta^t; e) = \begin{cases} c_t^{opt}(\theta^t; e) & \text{if } \theta^t \notin Q \\ v_t^{-1} \left(v_t \left(c_t^{opt}(\theta^t; e) \right) + k \right) & \text{if } \theta^t \in Q, \end{cases}$$

while in period $t+1$,

$$c_{t+1}^\#(\theta^{t+1}; e) = \begin{cases} c_{t+1}^{opt}(\theta^{t+1}; e) & \text{if } \theta^{t+1} \notin Q \\ v_{t+1}^{-1} \left(v_{t+1} \left(c_{t+1}^{opt}(\theta^{t+1}; e) \right) - k/\delta \right) & \text{if } \theta^{t+1} \in Q. \end{cases}$$

Clearly, this scheme also satisfies the constraints in program (2). The difference in terms of ex-ante expected cost to the principal under this scheme and under the original scheme $c^{opt}(\cdot; e)$ is given by

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

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 R^t , given the event that $\tilde{\theta}^t \in Q$.

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

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

The principal can then reduce her expected payment to the agent by switching to the compensation $c^\#(\cdot; e)$ with $k < 0$ arbitrarily small, contradicting the assumption that $c^{opt}(\cdot; e)$ solves program (2).

■

Proof of Proposition 2. We start with (i). Consider the following compensation scheme s . For any t any $(\theta^t, \pi^t) \in \Theta^t \times \Pi^t$, $s_t(\theta^t, \pi^t) = c_t(\theta^t)$ if $\pi_t \geq \theta_t + \xi_t(\theta^t)$ and $s_t(\theta^t, \pi^t) = -L_t$ otherwise, with the functions c_t satisfying for all $\theta^T \in \Theta^T$,⁴⁰

$$\begin{aligned} \mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t(c_t(\theta^t)) \right) &= \sum_{t=1}^T \delta^{t-1} \psi(\xi_t(\theta^t)) \\ &+ \int_{\underline{\theta}_1}^{\theta_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(\xi_t(\tilde{\theta}^t)) \right] ds \\ &+ \sum_{t=2}^T \delta^{t-1} H_t(\theta^t; \xi) \end{aligned} \quad (32)$$

where for each $\theta^t \in \Theta^{t41}$

$$\begin{aligned} H_t(\theta^t; \xi) &\equiv \int_{\underline{\theta}_t}^{\theta_t} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_\tau^t(s, \tilde{\theta}_{>t}^\tau) \psi'(\xi_\tau(\theta^{t-1}, s, \tilde{\theta}_{>t}^\tau)) \right] ds \\ &- \mathbb{E}^{\tilde{\theta}_t | \theta_{t-1}} \left[\int_{\underline{\theta}_t}^{\tilde{\theta}_t} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_\tau^t(s, \tilde{\theta}_{>t}^\tau) \psi'(\xi_\tau(\theta^{t-1}, s, \tilde{\theta}_{>t}^\tau)) \right] ds \right]. \end{aligned}$$

Under the proposed scheme, given the announcement $\hat{\theta}^s$ and the true type profile θ^s , the effort the agent chooses in each period s is given by

$$e_s(\theta_s; \hat{\theta}^s) = \xi_s(\hat{\theta}^s) + \hat{\theta}_s - \theta_s \quad (33)$$

Hence, we can conveniently denote histories by $(\theta^t, \hat{\theta}^t)$. Now let

$$\begin{aligned} U^\#(\theta^t; \hat{\theta}^t) &\equiv \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta_t} \left[\mathcal{V} \left(\sum_{s=1}^t \delta^{s-1} v_s(c_s(\hat{\theta}^s)) + \sum_{s=t+1}^T \delta^{s-1} v_s(c_s(\hat{\theta}^t, \tilde{\theta}_{>t}^s)) \right) \right] \\ &- \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta_t} \left[\sum_{s=1}^t \delta^{s-1} \psi(e_s(\theta_s; \hat{\theta}^s)) + \sum_{s=t+1}^T \delta^{s-1} \psi(\xi_s(\hat{\theta}^t, \tilde{\theta}_{>t}^s)) \right] \end{aligned}$$

⁴⁰The functions c_t can be constructed inductively from $t = 1$ onwards – proof available upon request.

⁴¹Note that the expressions for the expectations of the future realizations of the process use the property that the process is Markov.

denote the expected payoff when, following $(\theta^t, \hat{\theta}^{t-1})$, in period t the agent sends the message $\hat{\theta}_t$, he then chooses effort $e_t(\theta_t; \hat{\theta}^t) = \xi_t(\hat{\theta}^t) + \hat{\theta}_t - \theta_t$ and then from period $t + 1$ onwards follows a truthful and obedient strategy. Under the proposed scheme,

$$\begin{aligned}
U^\#(\theta^t; \hat{\theta}^t) &= \sum_{s=1}^t \delta^{s-1} \left[\psi(\xi_t(\hat{\theta}^t)) - \psi(e_s(\theta_s; \hat{\theta}^s)) \right] + \int_{\underline{\theta}_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(\xi_t(\tilde{\theta}^t)) \right] ds \\
&\quad + \sum_{s=2}^t \delta^{s-1} H_s(\hat{\theta}^s; \xi) \\
&\quad + \mathbb{E}^{\tilde{\theta}_{t+1} | \theta_t} \left[\int_{\underline{\theta}_{t+1}}^{\tilde{\theta}_{t+1}} \mathbb{E}^{\tilde{\theta}_{>t+1}^T | s} \left[\sum_{\tau=t+1}^T \delta^{\tau-1} J_{t+1}^\tau(s, \tilde{\theta}_{>t+1}^\tau) \psi'(\xi_\tau(\hat{\theta}^t, s, \tilde{\theta}_{>t+1}^\tau)) \right] ds \right] \\
&\quad - \mathbb{E}^{\tilde{\theta}_{t+1} | \hat{\theta}_t} \left[\int_{\underline{\theta}_{t+1}}^{\tilde{\theta}_{t+1}} \mathbb{E}^{\tilde{\theta}_{>t+1}^T | s} \left[\sum_{\tau=t+1}^T \delta^{\tau-1} J_{t+1}^\tau(s, \tilde{\theta}_{>t+1}^\tau) \psi'(\xi_\tau(\hat{\theta}^t, s, \tilde{\theta}_{>t+1}^\tau)) \right] ds \right].
\end{aligned}$$

Integrating by parts and using (33) one can verify that

$$\frac{\partial U^\#(\theta^t; \hat{\theta}^t)}{\partial \theta_t} = \delta^{t-1} \psi'(\xi_t(\hat{\theta}^t) + \hat{\theta}_t - \theta_t) + \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta_t} \left[\sum_{s=t+1}^T \delta^{s-1} J_t^s(\theta_t, \tilde{\theta}_{>t}^s) \psi'(\xi_s(\hat{\theta}^t, \tilde{\theta}_{>t}^s)) \right].$$

One can also verify that

$$\frac{dU^\#(\theta^t; \hat{\theta}^{t-1}, \theta_t)}{d\theta_t} = \delta^{t-1} \psi'(\xi_t(\hat{\theta}^{t-1}, \theta_t)) + \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta_t} \left[\sum_{s=t+1}^T \delta^{s-1} J_t^s(\theta_t, \tilde{\theta}_{>t}^s) \psi'(\xi_s(\hat{\theta}^{t-1}, \theta_t, \tilde{\theta}_{>t}^s)) \right]$$

The single crossing conditions in the proposition imply that

$$\left[\frac{dU^\#(\theta^t; \hat{\theta}^{t-1}, \theta_t)}{d\theta_t} - \frac{\partial U^\#(\theta^t; \hat{\theta}^{t-1}, \hat{\theta}_t)}{\partial \theta_t} \right] [\theta_t - \hat{\theta}_t] \geq 0$$

all $\theta^t \in R^t$, $\hat{\theta}^t \in \Theta^t$, which in turn implies that $U^\#(\theta^t; \hat{\theta}^{t-1}, \theta_t) \geq U^\#(\theta^t; \hat{\theta}^{t-1}, \hat{\theta}_t)$ all $\theta^t \in R^t$, $\hat{\theta}^t \in \Theta^t$. Because these conditions apply to all t , all histories $(\theta^t, \hat{\theta}^t) \in R^t \times \Theta^t$, we then conclude that, after any history $(\theta^t, \hat{\theta}^{t-1})$, a single deviation in period t from the truthful and obedient strategy in the continuation game that starts in period t with history $(\theta^t, \hat{\theta}^{t-1})$ is never optimal for the agent. This result, together with the boundeness of the extension ξ and the assumptions that the sets Θ_t , the impulse response functions $J_s^t(\cdot)$, and the marginal disutility of effort $\psi'(\cdot)$ are all bounded uniformly over t , guarantee that all other deviations are also unprofitable. The proof follows from standard backward-induction reasoning when T is finite. When it is infinite, it follows from arguments similar to those in Proposition 6 in Pavan, Segal, and Toikka (2011), adapted to the setting of this paper. We then conclude that, when the principal offers the mechanism $\Omega = \langle \xi, s \rangle$, with recommendation policy ξ satisfying the single-crossing conditions in the proposition and with compensation scheme as defined above, the agent finds it optimal to report truthfully and follow the recommended effort choice at all histories. This establishes that e is implementable.

Next, consider case (ii) and note that in this case the single-crossing conditions (7) are trivially satisfied at any $t > 1$. Let $w_t \equiv v_t^{-1}$ and $\mathcal{W} \equiv \mathcal{V}^{-1}$. Let s be the compensation scheme defined as follows. For $t = 1$,

$$s_1(\theta_1, \pi_1) = w_1 \left(\mathcal{W} \left(\begin{array}{c} \sum_{t=1}^T \delta^{t-1} \psi(\xi_t(\theta_1)) + \int_{\underline{\theta}_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}_{>1}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(s, \tilde{\theta}_{>1}^t) \psi'(\xi_t(s)) \right] ds \\ - \mathbb{E}^{\tilde{\theta}_{>1}^T | \theta_1} \left[\sum_{t=2}^T \delta^{t-1} \psi'(\xi_t(\theta_1)) (\tilde{\theta}_t + \xi_t(\theta_1)) \right] \end{array} \right) \right)$$

if $\pi_1 \geq \theta_1 + \xi_1(\theta_1)$ and $s_1(\theta_1, \pi_1) = -L_1$ otherwise. For any $t > 1$,

$$s_t(\theta^t, \pi^t) = w_t \left(\begin{array}{c} -\delta^{1-t} \sum_{\tau=1}^{t-1} \delta^{\tau-1} v_\tau(s_\tau(\theta^\tau, \pi^\tau)) \\ + \delta^{1-t} \mathcal{W} \left(\begin{array}{c} \sum_{t=1}^T \delta^{t-1} \psi(\xi_t(\theta_1)) + \int_{\underline{\theta}_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}_{>1}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(s, \tilde{\theta}_{>1}^t) \psi'(\xi_t(s)) \right] ds \\ - \mathbb{E}^{\tilde{\theta}_{>1}^T | \theta_1} \left[\sum_{t=2}^T \delta^{t-1} \psi'(\xi_t(\theta_1)) (\tilde{\theta}_t + \xi_t(\theta_1)) \right] \\ + \sum_{\tau=2}^t \delta^{\tau-1} \psi'(\xi_\tau(\theta_1)) \pi_\tau \end{array} \right) \end{array} \right)$$

if $\pi_1 \geq \xi_1(\theta_1) + \theta_1$, and $s_t(\theta^t, \pi^t) = -L_t$ otherwise. Note that the above scheme is constructed so that, at each period t , if the agent were to receive no further payments from the principal in the future, then upon meeting the first-period target $\pi_1 = \xi_1(\hat{\theta}_1) + \hat{\theta}_1$, his utility from past and current payments would depend only on the first period report $\hat{\theta}_1$ and be linear in the cash flows $(\pi_s)_{s=2}^t$ generated between period 2 and period t . To see that such a scheme induces the desired effort choices, suppose first that T is finite. By construction, given any history of reports $\hat{\theta}^T$ and any history of profit realizations π^T , with $\pi_1 \geq \xi_1(\hat{\theta}_1) + \hat{\theta}_1$, the agent's total utility from money is given by

$$\begin{aligned} \mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t(s_t(\hat{\theta}^t, \pi^t)) \right) &= \sum_{t=1}^T \delta^{t-1} \psi(\xi_t(\hat{\theta}_1)) + \int_{\underline{\theta}_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}_{>1}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(s, \tilde{\theta}_{>1}^t) \psi'(\xi_t(s)) \right] ds \\ &\quad + \sum_{t=2}^T \delta^{t-1} \psi'(\xi_t(\hat{\theta}_1)) (\pi_t - \mathbb{E}^{\tilde{\theta}_{>1}^T | \hat{\theta}_1} [\tilde{\theta}_t] - \xi_t(\hat{\theta}_1)). \end{aligned}$$

Next, consider the case where $T = \infty$. By continuity of \mathcal{V} ,

$$\begin{aligned} \mathcal{V} \left(\sum_{t=1}^{\infty} \delta^{t-1} v_t(s_t(\hat{\theta}^t, \pi^t)) \right) &= \lim_{Z \rightarrow \infty} \mathcal{V} \left(\sum_{t=1}^Z \delta^{t-1} v_t(s_t(\hat{\theta}^t, \pi^t)) \right) \\ &= \sum_{t=1}^{\infty} \delta^{t-1} \psi(\xi_t(\hat{\theta}_1)) + \int_{\underline{\theta}_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}_{>1}^T | s} \left[\sum_{t=1}^{\infty} \delta^{t-1} J_1^t(s, \tilde{\theta}_{>1}^t) \psi'(\xi_t(s)) \right] ds \\ &\quad + \sum_{t=2}^{\infty} \delta^{t-1} \psi'(\xi_t(\hat{\theta}_1)) (\pi_t - \mathbb{E}^{\tilde{\theta}_{>1}^T | \hat{\theta}_1} [\tilde{\theta}_t] - \xi_t(\hat{\theta}_1)). \end{aligned}$$

It is then easy to see that, under the proposed scheme, the agent's decision problem is time-separable after $t = 1$. After having reported $\hat{\theta}_1$ in period one, and having delivered a cash flow $\pi_1 \geq \xi_1(\hat{\theta}_1) + \hat{\theta}_1$, at any period $t > 1$, and for any possible history $(\theta^t, (\hat{\theta}_1, \hat{\theta}_{>1}^t), e^{t-1})$, it is optimal for the agent to choose a level of effort $e_t = \xi_t(\hat{\theta}_1)$. To establish the result it then suffices to verify

that each type θ_1 finds it optimal to report truthfully in period one. The payoff that each type θ_1 obtains by reporting $\hat{\theta}_1$ and then choosing optimally a level of effort $e_1(\theta_1; \hat{\theta}_1) = \xi_1(\hat{\theta}_1) + \hat{\theta}_1 - \theta_1$ for $t = 1$ and $e_t = \xi_t(\hat{\theta}_1)$ for any $t > 1$ is

$$\begin{aligned} U^\#(\theta_1; \hat{\theta}_1) &= \psi(\xi_1(\hat{\theta}_1)) - \psi(\xi_1(\hat{\theta}_1) + \hat{\theta}_1 - \theta_1) \\ &\quad + \int_{\theta_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}_{>1}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(s, \tilde{\theta}_{>1}^t) \psi'(\xi_t(s)) \right] ds \\ &\quad - \mathbb{E}^{\tilde{\theta}_{>1}^T | \hat{\theta}_1} \left[\sum_{t=2}^T \delta^{t-1} \psi'(\xi_t(\hat{\theta}_1)) (\tilde{\theta}_t + \xi_t(\hat{\theta}_1)) \right] + \mathbb{E}^{\tilde{\theta}_{>1}^T | \theta_1} \left[\sum_{t=2}^T \delta^{t-1} \psi'(\xi_t(\hat{\theta}_1)) (\tilde{\theta}_t + \xi_t(\hat{\theta}_1)) \right] \end{aligned}$$

Once again, integrating by parts, one can verify that the single-crossing condition (7) in the proposition implies that

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

which in turn implies that $U^\#(\theta_1; \theta_1) \geq U^\#(\theta_1; \hat{\theta}_1)$ all $\theta_1, \hat{\theta}_1 \in \Theta_1$, thus establishing the result. ■

Proof of Proposition 3.

We start with the following definition. For any given extension ξ of any given policy e , let $c^{opt}(\cdot; \xi) \equiv \langle c_t^{opt}(\cdot; \xi) \rangle_{t=1}^{t=T}$ denote any arbitrary (Θ -measurable) extension of the compensation $c^{opt}(\cdot; e)$ defined in Proposition 1 such that

$$\mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t(c_t^{opt}(\theta^t)) \right) = W^\#(\theta^T; \xi) \quad \text{all } \theta^T \in \Theta^T \quad (34)$$

where

$$\begin{aligned} W^\#(\theta^T; \xi) &\equiv \sum_{t=1}^T \delta^{t-1} \psi(\xi_t(\theta^t)) \\ &\quad + \int_{\theta_1}^{\theta_1} \mathbb{E}^{\tilde{\theta}_{>1}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(\xi_t(\tilde{\theta}^t)) \right] ds + \sum_{t=2}^T \delta^{t-1} H_t^\#(\theta^t; \xi) \end{aligned}$$

with

$$\begin{aligned} H_t^\#(\theta^t; \xi) &\equiv \int_{\underline{\theta}_t}^{\theta_t} \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta^{t-1}, s} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_\tau^T(\theta^{t-1}, s, \tilde{\theta}_{>t}^\tau) \psi'(\xi_\tau(\hat{\theta}^{t-1}, s, \tilde{\theta}_{>t}^\tau)) \right] ds \\ &\quad - \mathbb{E}^{\tilde{\theta}_{>t} | \theta^{t-1}} \left[\int_{\underline{\theta}_t}^{\tilde{\theta}_t} \mathbb{E}^{\tilde{\theta}_{>t}^T | s} \left[\sum_{\tau=t}^T \delta^{\tau-t} J_\tau^T(\theta^{t-1}, s, \tilde{\theta}_{>t}^\tau) \psi'(\xi_\tau(\hat{\theta}^{t-1}, s, \tilde{\theta}_{>t}^\tau)) \right] ds \right] \end{aligned}$$

In other terms, $c^{opt}(\cdot; \xi)$ is any (Θ -measurable) compensation $c \equiv \langle c_t(\cdot) \rangle_{t=1}^{t=T}$ that minimizes

$$\mathbb{E}^{\tilde{\theta}^T} \left[\sum_{t=1}^T \delta^{t-1} c_t(\tilde{\theta}^t) \right]$$

among those that satisfy the constraints given by (34).

Next, consider the following pseudo-linear scheme. For $t = 1$,

$$s_1(\theta_1, \pi_1) = w_t \left(\mathcal{W} \left(\mathcal{V} \left(v_1 \left(c_1^{opt}(\theta_1; \xi) \right) \right) + \psi'(\xi_1(\theta_1)) [\pi_1 - \theta_1 - \xi_1(\theta_1)] \right) \right)$$

while for any $t > 1$ any (θ^t, π^t)

$$s_t(\theta^t, \pi^t) = w_t \left(\begin{array}{c} \frac{1}{\delta^{t-1}} \mathcal{W} \left(\mathcal{V} \left(\sum_{\tau=1}^t \delta^{\tau-1} v_t \left(c_\tau^{opt}(\theta^\tau; \xi) \right) \right) + \sum_{\tau=1}^t \delta^{\tau-1} \psi'(\xi_\tau(\theta^\tau)) [\pi_\tau - \theta_\tau - \xi_\tau(\theta^\tau)] \right) \\ - \frac{1}{\delta^{t-1}} \sum_{\tau=1}^{t-1} \delta^{\tau-1} v_t(s_\tau(\theta^\tau, \pi^\tau)) \end{array} \right)$$

where recall that $w_t \equiv v_t^{-1}$ and $\mathcal{W} \equiv \mathcal{V}^{-1}$. Note that this scheme is constructed so that, when the agent reports his type truthfully and chooses effort obediently, the payments he receives over time coincide with the cost-minimizing compensations given by $c^{opt}(\cdot; \xi)$. From Proposition 1, it then follows that, if this scheme implements the policy e , it then necessarily implements it at minimal cost for the principal. In what follows, it thus suffices to show that, under such a scheme, a truthful and obedient strategy is optimal for the agent. To this purpose, note that this scheme is also chosen so that, given any sequence of reports and cash flows (θ^T, π^T) , the utility that the agent derives from the entire stream of payments is given by

$$\mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t(s_t(\theta^t, \pi^t)) \right) = W^\#(\theta^T; \xi) + \sum_{t=1}^T \delta^{t-1} \psi'(\xi_t(\theta^t)) (\pi_t - \theta_t - \xi_t(\theta^t)).$$

The latter is linear in each cash-flow π_t with coefficient of linear dependence $\alpha_t(\theta^t) = \psi'(\xi_t(\theta^t))$. As mentioned already in the proof of Proposition 2, such linearity guarantees that the decision problem for the agent is time-separable: in each period $t \geq 1$, after having reported $\hat{\theta}^t$ and irrespective of the history (θ^t, e^{t-1}) , the agent finds it optimal to choose a level of effort $e_t = \xi_t(\hat{\theta}^t)$. In other words, such linearity takes care of the moral-hazard part of the problem. To establish the result, it then suffices to verify that, in each period t , the agent finds it optimal to report his type θ_t truthfully. To

establish this, let

$$\mathcal{U}^\#(\theta^t; \hat{\theta}^t, e^{t-1}) \equiv \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta^t} \left[\mathcal{V} \left(\begin{aligned} & \sum_{\tau=1}^{t-1} \delta^{\tau-1} v_\tau \left(s_\tau \left(\hat{\theta}^\tau, (\theta_z + e_z)_{z=1}^\tau \right) \right) + \\ & + \delta^{t-1} v_t \left(s_t \left(\hat{\theta}^t, \left((\theta_z + e_z)_{z=1}^{t-1}, \theta_t + \xi_t(\hat{\theta}^t) \right) \right) \right) + \\ & \sum_{\tau=t+1}^T \delta^{\tau-1} v_\tau \left(s_\tau \left((\hat{\theta}^t, \tilde{\theta}_{>t}^\tau), \left((\theta_z + e_z)_{z=1}^{t-1}, \theta_t + \xi_t(\hat{\theta}^t), (\tilde{\theta}_z + \xi_z(\hat{\theta}^t, \tilde{\theta}_{>t}^z))_{z=t+1}^\tau \right) \right) \right) \end{aligned} \right) \right] \\ - \sum_{\tau=1}^{t-1} \delta^{\tau-1} \psi(e_\tau) + \delta^{t-1} \psi(\xi_t(\hat{\theta}^t)) - \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta^t} \left[\sum_{\tau=t+1}^T \psi(\xi_\tau(\hat{\theta}^t, \tilde{\theta}_{>t}^\tau)) \right]$$

denote the expected payoff when, under the proposed scheme, following the history $(\theta^t, \hat{\theta}^{t-1}, e^{t-1})$, in period t the agent sends the message $\hat{\theta}_t$, he then chooses optimally effort $e_t = \xi_t(\hat{\theta}^t)$ and then from period $t+1$ onwards he follows a truthful and obedient strategy. To establish that truthful reporting in each period is optimal for the agent, we consider the two cases in the Proposition separately.

Case (i): Markov process. Using the definitions of s and $c^{opt}(\cdot; \xi)$ and the fact that the process is Markov, we have that in this case

$$\begin{aligned} \mathcal{U}^\#(\theta^t; \hat{\theta}^t, e^{t-1}) &= \int_{\underline{\theta}_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\tilde{\theta}^t) \psi'(\xi_t(\tilde{\theta}^t)) \right] ds \\ &+ \sum_{\tau=1}^{t-1} \delta^{\tau-1} \left\{ \psi(\xi_\tau(\hat{\theta}^\tau)) - \psi(e_\tau) + \psi'(\xi_\tau(\hat{\theta}^\tau)) \left[\theta_\tau + e_\tau - \hat{\theta}_\tau - \xi_\tau(\hat{\theta}^\tau) \right] \right\} \\ &+ \delta^{t-1} \psi'(\xi_t(\hat{\theta}^t)) \left[\theta_t - \hat{\theta}_t \right] \\ &+ \sum_{\tau=2}^t \delta^{\tau-1} H_\tau^\#(\hat{\theta}^\tau; \xi) + \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta^t} \left[\sum_{s=t+1}^T \delta^{s-1} H_s^\#(\hat{\theta}^t, \tilde{\theta}_{>t}^s; \xi) \right] \end{aligned} \quad (35)$$

from which we obtain that

$$\frac{\partial \mathcal{U}^\#(\theta^t; \hat{\theta}^t, e^{t-1})}{\partial \theta_t} = \delta^{t-1} \psi'(\xi_t(\hat{\theta}^t)) + \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta^t} \left[\sum_{s=t+1}^T \delta^{s-1} J_t^s(\theta_t, \tilde{\theta}_{>t}^s) \psi'(\xi_s(\hat{\theta}^t, \tilde{\theta}_{>t}^s)) \right]$$

and

$$\frac{d\mathcal{U}^\#(\theta^t; \hat{\theta}^{t-1}, \theta_t, e^{t-1})}{d\theta_t} = \delta^{t-1} \psi'(\xi_t(\hat{\theta}^{t-1}, \theta_t)) + \mathbb{E}^{\tilde{\theta}_{>t}^T | \theta^t} \left[\sum_{s=t+1}^T \delta^{s-1} J_t^s(\theta_t, \tilde{\theta}_{>t}^s) \psi'(\xi_s(\hat{\theta}^{t-1}, \theta_t, \tilde{\theta}_{>t}^s)) \right]$$

The single crossing conditions in the proposition imply that

$$\left[\frac{d\mathcal{U}^\#(\theta^t; \hat{\theta}^{t-1}, \theta_t, e^{t-1})}{d\theta_t} - \frac{\partial \mathcal{U}^\#(\theta^t; \hat{\theta}^t, e^{t-1})}{\partial \theta_t} \right] \left[\theta_t - \hat{\theta}_t \right] \geq 0$$

all $\theta^t \in R^t$, $\hat{\theta}^t \in \Theta^t$, which in turn implies that $\mathcal{U}^\#(\theta^t; \hat{\theta}^{t-1}, \theta_t, e^{t-1}) \geq \mathcal{U}^\#(\theta^t; \hat{\theta}^t, e^{t-1})$ all $\theta^t \in R^t$, $\hat{\theta}^t \in \Theta^t$. As argued in the proof of Proposition 2, these conditions guarantee that, after any history $(\theta^t, \hat{\theta}^{t-1}, e^{t-1})$, a single deviation in period t from the truthful and obedient strategy in the continuation game that starts in period t with history $(\theta^t, \hat{\theta}^{t-1}, e^{t-1})$ is never optimal for the agent. The same arguments as in the proof of Proposition 2 then imply that all other deviations are also unprofitable. We conclude that the agent finds it optimal to report truthfully at all histories, which means that the above pseudo-linear scheme s implements the policy e .

Case (ii). Policies depending on θ_1 only. Using the fact that

$$\mathbb{E}^{\tilde{\theta}^T | \theta^{t-1}, \theta_t} \left[J_t^T(\tilde{\theta}^\tau) \right] = \frac{\partial}{\partial \theta_t} \left[\mathbb{E}^{\tilde{\theta}^T | \theta^{t-1}, \theta_t} \left[\tilde{\theta}_\tau \right] \right]$$

all t, τ , $\tau > t$, all θ^t , one can verify that

$$\begin{aligned} \mathcal{V} \left(\sum_{t=1}^T \delta^{t-1} v_t \left(s_t(\hat{\theta}^t, \pi^t) \right) \right) &= \sum_{t=1}^T \delta^{t-1} \psi(\xi_t(\hat{\theta}_1)) + \int_{\underline{\theta}_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}_{>1}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(s, \tilde{\theta}_{>1}^t) \psi'(\xi_t(s)) \right] ds \\ &\quad + \sum_{t=2}^T \delta^{t-1} \psi'(\xi_t(\hat{\theta}_1)) (\pi_t - \mathbb{E}^{\tilde{\theta}_{>1}^T | \hat{\theta}_1} [\tilde{\theta}_t] - \xi_t(\hat{\theta}_1)). \end{aligned}$$

It is therefore easy to see that, under the proposed scheme, after having reported $\hat{\theta}_1$ in period one, at any period $t \geq 1$, and for any possible history $(\theta^t, \hat{\theta}^t, e^{t-1})$, it is optimal for the agent to choose a level of effort $e_t = \xi_t(\hat{\theta}_1)$. It is also easy to see that the payoff that each type θ_1 obtains by reporting $\hat{\theta}_1$ in period one and then choosing effort optimally at any subsequent information set is equal to

$$\begin{aligned} U^\#(\theta_1; \hat{\theta}_1) &= \int_{\underline{\theta}_1}^{\hat{\theta}_1} \mathbb{E}^{\tilde{\theta}_{>1}^T | s} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(s, \tilde{\theta}_{>1}^t) \psi'(\xi_t(s)) \right] ds \\ &\quad + \sum_{t=1}^T \delta^{t-1} \psi'(\xi_t(\hat{\theta}_1)) \left[\mathbb{E}^{\tilde{\theta}_t | \theta_1} [\tilde{\theta}_t] - \mathbb{E}^{\tilde{\theta}_t | \hat{\theta}_1} [\tilde{\theta}_t] \right] \end{aligned}$$

from which we obtain that

$$\frac{\partial U^\#(\theta_1; \hat{\theta}_1)}{\partial \theta_1} = \mathbb{E}^{\tilde{\theta}_{>1}^T | \theta_1} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\theta_1, \tilde{\theta}_{>1}^t) \psi'(\xi_t(\hat{\theta}_1)) \right]$$

and

$$\frac{dU^\#(\theta_1; \theta_1)}{d\theta_1} = \mathbb{E}^{\tilde{\theta}_{>1}^T | \theta_1} \left[\sum_{t=1}^T \delta^{t-1} J_1^t(\theta_1, \tilde{\theta}_{>1}^t) \psi'(\xi_t(\theta_1)) \right]$$

Once again, the single-crossing condition (7) in the proposition implies that, for all $\theta_1, \hat{\theta}_1 \in \Theta_1$,

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

thus implying that $U^\#(\theta_1; \theta_1) \geq U^\#(\theta_1; \hat{\theta}_1)$ all $\theta_1, \hat{\theta}_1 \in \Theta_1$, which establishes the result. ■

Proof of Proposition 4. Part (i). The assumptions that ψ is continuously differentiable 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 $J_1^t(\theta^t) \geq 0$ for all t all θ^t , imply that, for all t all $\theta^t \in R^t$, the principal's payoff, as given by (12), is strictly increasing in e_t for all $e_t < e_t^*(\theta^t)$, and strictly decreasing in e_t for all $e_t > e_t^*(\theta^t)$, where $e_t^*(\theta^t)$ is implicitly given by (13) when $\psi_+''(0) < 1/[\eta(\theta_1)J_1^t(\theta^t)]$ and by $e_t^*(\theta^t) = 0$ otherwise. It is then immediate that when the policy e^* is implementable, then, under any optimal contract for the principal, e^* is implemented in each period t at F -almost all histories θ^t . This follows from the fact that the principal's payoff in any mechanism that is incentive-compatible and individually-rational for the agent is given by (12) together with the fact the policy e^* maximizes (12) and the fact that, if the policy e^* is implementable, then from Proposition 1 there exists a compensation scheme s^* that implements e^* and that gives the lowest period-1 type $\underline{\theta}_1$ an expected payoff equal to his outside option, in which case $V^\Omega(\underline{\theta}_1) = 0$.

Part (ii). The result follows from Propositions 2 and 3. Consider the following bounded extension of the policy e^* from R to Θ . For any $t \geq 2$, any $2 \leq s \leq t$, let $\varphi_s : \Theta^t \rightarrow \Theta^t$ be the function defined, for all $\theta^t \in \Theta^t$, by

$$\varphi_s(\theta^t) \equiv \begin{cases} \theta^t & \text{if } \theta_s \in \text{Supp}[F_s(\cdot|\theta^{s-1})] \\ (\theta^{s-1}, \min\{\text{Supp}F_s(\cdot|\theta^{s-1})\}, \theta_{s+1}, \dots, \theta_t) & \text{if } \theta_s < \min\{\text{Supp}[F_s(\cdot|\theta^{s-1})]\} \\ (\theta^{s-1}, \max\{\text{Supp}[F_s(\cdot|\theta^{s-1})]\}, \theta_{s+1}, \dots, \theta_t) & \text{if } \theta_s > \max\{\text{Supp}[F_s(\cdot|\theta^{s-1})]\} \end{cases} .$$

For all $\theta_1 \in \Theta_1$, then let $\lambda^1(\theta_1) \equiv \theta_1$, while for any $t \geq 2$, let $\lambda^t : \Theta^t \rightarrow R^t$ be the function defined, for all $\theta^t \in \Theta^t$, by $\lambda^t(\theta^t) \equiv \varphi_t \circ \varphi_{t-1} \circ \dots \circ \varphi_2(\theta^t)$. Note that the function λ^t maps each vector of reports $\theta^t \in \Theta^t \setminus R^t$ into a vector of reports $\hat{\theta}^t = (\lambda_s^t(\theta^t)) \in R^t$. This is obtained by replacing recursively any report θ_s that, given the past reports $\hat{\theta}^{s-1} = (\lambda_l^t(\theta^t))_{l=1}^{s-1}$ is smaller than any feasible type with $\hat{\theta}_s = \min\{\text{Supp}[F_s(\cdot|\hat{\theta}^{s-1})]\}$, and, likewise, by replacing any report θ_s that is higher than any feasible type with $\hat{\theta}_s = \max\{\text{Supp}[F_s(\cdot|\hat{\theta}^{s-1})]\}$.

Now, let ξ be the recommendation policy given by $\xi_t(\theta^t) = e_t^*(\lambda^t(\theta^t))$ all t , all $\theta^t \in \Theta^t$. Clearly, when $\Theta^t = R^t$ all t , the policy ξ is simply the policy which recommends the equilibrium effort at all histories irrespective of whether or not the agent has been truthful in the past. When, instead, $R^t \subset \Theta^t$ for some t , the policy ξ is obtained by replacing the reports that indicate a departure from truthtelling in past periods with the reports $\lambda^t(\theta^t)$ and then recommending the equilibrium effort $e_t^*(\lambda^t(\theta^t))$ for the type history $\lambda^t(\theta^t)$.

Assumption (b) in the proposition guarantees that each function $e_t^*(\theta^t)$ is nondecreasing. By construction, each function $\xi_t(\theta^t)$ is also nondecreasing. We conclude that the extension ξ constructed above satisfies not only the single-crossing conditions (7) of Proposition 2 but also the stronger conditions (8) of Proposition 3. From the results in those propositions, we then conclude that when, in addition, F is Markov or each function $J_1^t(\theta^t)$ depends on θ^t only through θ_1 (in which case so

does each function ξ_t), then the policy e^* can be implemented at minimum costs by both the bonus schemes of Proposition 1 and by the linear schemes of Proposition 3. ■

Proof of Example 4. The proof follows from the result in Proposition 5, after restricting effort to depend only on θ_1 and, for (16), after taking expectations with respect to ε_2 . ■

Proof of Proposition 5. Suppose that e^* maximizes the principal's payoff (14). For each t , let $k_t \in \mathbb{R}$ be a scalar and $y_t : R^t \rightarrow \mathbb{R}$ be a measurable function. Consider the expression for the principal's payoff (14), evaluated at the perturbed solution $e^* + ky$ given by $e_t^*(\theta^t) + k_t y_t(\theta^t)$ all t all $\theta^t \in R^t$. Then a necessary condition for the optimality of the policy e^* is that, for all t , the derivative of the principal's payoff (14) with respect to k_t evaluated at $k_s = 0$ all $s = 1, \dots, T$, must vanish. That is,

$$\mathbb{E}^{\tilde{\theta}^T} \left[y_t(\tilde{\theta}^t) - \mathcal{W}' \left(W(\tilde{\theta}^T; e^*) \right) \left(\begin{array}{l} \psi'(e_t(\tilde{\theta}^t)) y_t(\tilde{\theta}^t) \\ + \int_{\theta_1}^{\tilde{\theta}_1} \mathbb{E}^{\tilde{\theta}^T | s} \left[J_1^t(\tilde{\theta}^t) \psi''(e_t(\tilde{\theta}^t)) y_t(\tilde{\theta}^t) \right] ds \\ + \sum_{m=2}^t \int_{\underline{\theta}_m(\tilde{\theta}^{m-1})}^{\tilde{\theta}_m} \mathbb{E}^{\tilde{\theta}^T | \tilde{\theta}^{m-1}, s} \left[J_m^t(\tilde{\theta}^t) \psi''(e_t(\tilde{\theta}^t)) y_t(\tilde{\theta}^t) \right] ds \\ - \sum_{m=2}^t \mathbb{E}^{\tilde{\theta}_m | \tilde{\theta}^{m-1}} \left[\int_{\underline{\theta}_m(\tilde{\theta}^{m-1})}^{\tilde{\theta}_m} \mathbb{E}^{\tilde{\theta}^T | \tilde{\theta}^{m-1}, s} \left[J_m^t(\tilde{\theta}^t) \psi''(e_t(\tilde{\theta}^t)) y_t(\tilde{\theta}^t) \right] ds \right] \end{array} \right) \right] = 0.$$

Integration by parts, together with the law of iterated expectations, gives

$$\mathbb{E}^{\tilde{\theta}^t} \left[y_t(\tilde{\theta}^t) \left(\begin{array}{l} 1 - \psi'(e_t(\tilde{\theta}^t)) \mathbb{E}^{\tilde{\theta}^T | \tilde{\theta}^t} \left[\mathcal{W}' \left(W(\tilde{\theta}^T; e^*) \right) \right] \\ - \psi''(e_t(\tilde{\theta}^t)) \frac{J_1^t(\tilde{\theta}^t)}{f_1(\tilde{\theta}_1)} \int_{\tilde{\theta}_1}^{\tilde{\theta}_1} \mathbb{E}^{\tilde{\theta}^T | r} \left[\mathcal{W}' \left(W(\tilde{\theta}^T; e^*) \right) \right] f_1(r) dr \\ - \psi''(e_t(\tilde{\theta}^t)) \sum_{m=2}^t \frac{J_m^t(\tilde{\theta}^t)}{f_m(\tilde{\theta}_m | \tilde{\theta}^{m-1})} \\ \times \int_{\tilde{\theta}_m}^{\tilde{\theta}_m} \left(\begin{array}{l} \mathbb{E}^{\tilde{\theta}^T | \tilde{\theta}^{m-1}, r} \left[\mathcal{W}' \left(W(\tilde{\theta}^T; e^*) \right) \right] \\ - \mathbb{E}^{\tilde{\theta}^T | \tilde{\theta}^{m-1}} \left[\mathcal{W}' \left(W(\tilde{\theta}^T; e^*) \right) \right] \end{array} \right) f_m(r | \tilde{\theta}^{m-1}) dr \end{array} \right) \right] = 0.$$

Ignoring corner solutions, this condition holds for all measurable functions $y_t(\cdot)$ only if the term in the inner bracket is equal to zero for F -almost all θ^t , which leads to the Euler condition (20). ■

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.

- [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] Bergemann, Dirk and Ulrich Hege (1998), ‘Dynamic Venture Capital Financing, Learning and Moral Hazard,’ *Journal of Banking and Finance*, 22: 703-735.
- [7] Bergemann, Dirk and Ulrich Hege (2005), ‘The Financing of Innovation,’ *RAND Journal of Economics*, 36, 719-752.
- [8] Besanko, David (1985), ‘Multi-Period Contracts between Principal and Agent with Adverse Selection’, *Economics Letters*, 17, 33-37.
- [9] Biais, Bruno, Thomas Mariotti, Jean-Charles Rochet and Stéphane Villeneuve, (2010) ‘Large Risks, Limited Liability, and Dynamic Moral Hazard,’ *Econometrica*, 78(1) 73-118.
- [10] Buttazzo, Giuseppe, Mariano Giaquinta, and Stefan Hildebrandt (1998), ‘*One-dimensional variational problems: an introduction*,’ Volume 15 of Oxford lecture series in mathematics and its applications, Oxford science publications.
- [11] Courty, Pascal. and Li Hao (2000), ‘Sequential Screening,’ *Review of Economic Studies*, 67, 697–717.
- [12] Cremers, Martijn and Darius Palia (2010), ‘Tenure and CEO Pay’, mimeo Yale University.
- [13] DeMarzo, Peter and Michael Fishman (2007), ‘Optimal Long-Term Financial Contracting’, *Review of Financial Studies*, 20, 2079-2127.
- [14] 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.
- [15] DeMarzo, Peter and Yuliy Sannikov (2008), ‘Learning in Dynamic Incentive Contracts’, mimeo, Stanford University.
- [16] 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.
- [17] Edmans, Alex and Xavier Gabaix (2010), ‘Tractability and in Incentive Contracting’, *Review of Financial Studies*, forthcoming.

- [18] Edmans, Alex and Xavier Gabaix (2009), ‘Is CEO Pay Really Inefficient? A Survey of New Optimal Contracting Theories’ *European Financial Management* (forthcoming).
- [19] Edmans, Alex, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov (2010), ‘Dynamic Incentive Accounts,’ mimeo NYU Stern, Princeton and Wharton.
- [20] Eso, P. and B. Szentes, (2007), “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74(3), 705-731.
- [21] Fahri, Emmanuel and Ivan Werning (2010) ‘Insurance and Taxation over the Lyfe Cycle,’ mimeo MIT.
- [22] Garfagnini, Umberto (2011) ‘Delegated Experimentation,’ mimeo Northwestern Univesity.
- [23] Garrett, Daniel, and Pavan Alessandro (2009) ‘Dynamic Managerial Compensation: a Mechanism Design Approach,’ mimeo Northwestern University.
- [24] Garrett, Daniel, and Pavan Alessandro (2010) ‘Managerial Turnover in a Changing World,’ mimeo Northwestern University.
- [25] Gibbons, Robert, and Kevin J. Murphy (1992) ‘Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence’, *Journal of Political Economy* 100, 468-505.
- [26] Golosov, M. M. Troshkin, and A. Tsyvinski, (2010), “Optimal Dynamic Taxes,” Mimeo Yale University.
- [27] Hill, C.W.L. and P. Phan (1991), ‘CEO Tenure as a Determinant of CEO Pay’, *Academy of Management Journal* 34(3), 707-717.
- [28] He, Zhiguo (2008), ‘Optimal Executive Compensation when Firm Size Follows a Geometric Brownian Motion,’ *Review of Financial Studies*, 22, 859-892.
- [29] Holmstrom, Bengt and Paul Milgrom (1987), ‘Aggregation and Linearity in the Provision of Intertemporal Incentives’, *Econometrica* 55(2), 303-328.
- [30] Horner, Johannes and Larry Samuelson (2009), ‘Incentives for Experimenting Agents,’ mimeo, Yale University,
- [31] Kuhnen, Camelia and Jeffrey Zwiebel (2008), ‘Executive Pay, Hidden Compensation and Managerial Entrenchment’, *Rock Center for Corporate Governance Working Paper* No. 16.
- [32] Laffont, Jean-Jacques and Rochet Jean-Charles (1998), ‘Regulation of a Risk Averse Firm,’ *Games and Economic Behavior*, 25(2), 149-173.
- [33] Laffont, Jean-Jacques and Jean Tirole (1986), ‘Using Cost Observation to Regulate Firms’, *Journal of Political Economy*, 94, 614-641.

- [34] Lambert, Richard A. (1983), ‘Long-term contracts and moral hazard’, *Bell Journal of Economics*, 14(2), 441-452.
- [35] 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.
- [36] Pavan, Alessandro, Ilya Segal and Juuso Toikka, (2011) ‘Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization and Information Disclosure’, mimeo Northwestern University and Stanford University.
- [37] Phelan, Christopher and Robert M. Townsend (1991), ‘Computing Multi-Period, Information-Constrained Optima’, *Review of Economic Studies* 58, 853-881.
- [38] Rogerson, William P. (1985), ‘Repeated Moral Hazard’, *Econometrica* 53(1), 69-76.
- [39] Salanie, Bernard (1990), ‘Selection adverse et aversion pour le risque’, *Annales d’Economie et de Statistique*, No. 18, pp 131-149.
- [40] Sannikov, Yuliy (2007), ‘Agency Problems, Screening and Increasing Credit Lines,’ mimeo Berkeley University.
- [41] Sannikov, Yuliy (2008), ‘A Continuous-Time Version of the Principal-Agent Problem’, *Review of Economic Studies*, 75(3), 957-984.
- [42] Spear, Stephen E. and Sanjay Srivastava (1987) ‘On Repeated Moral Hazard with Discounting’, *Review of Economic Studies*, 54(4), 599-617.
- [43] Strulovici, Bruno (2010) ‘Renegotiation-Proof Contracts with Moral Hazard and Persistent Private Information,’ mimeo Northwestern University.
- [44] Tchisty, Alexei (2006), ‘Security Design with Correlated Hidden Cash Flows: The Optimality of Performance Pricing’, mimeo Stern, NYU.
- [45] Zhang, Yuzhe (2009), ‘Dynamic contracting with persistent shocks’, *Journal of Economic Theory* 144, 635-675.
- [46] Williams, N. (2011), ‘Persistent Private Information,’ *Econometrica*, forthcoming.