Review of Economic Studies (2021) **88**, 1253–1278 doi:10.1093/restud/rdaa043

Advance access publication 3 August 2020

Only Time Will Tell: A Theory of Deferred Compensation

FLORIAN HOFFMANN

Erasmus University Rotterdam and KU Leuven

ROMAN INDERST

Johann Wolfgang Goethe University Frankfurt and CEPR

and

MARCUS OPP

Stockholm School of Economics and CEPR

First version received February 2019*; Editorial decision June* 2020*; Accepted July* 2020 (*Eds.*)

This article characterizes optimal compensation contracts in principal-agent settings in which the consequences of the agent's action are only observed over time. The optimal timing of pay trades off the costs of deferred compensation arising from the agent's relative impatience and potential consumption smoothing needs against the benefit of exploiting additional informative signals. By capturing this information benefit of deferral in terms of the likelihood ratio dynamics, our characterization covers general signal processes in a unified setting. With bilateral risk neutrality and agent limited liability, optimal contracts are high-powered and stipulate at most two payout dates. If the agent is additionally risk-averse, payouts are contingent on performance exceeding a hurdle that is increasing over time. We obtain clear-cut predictions on how the duration of optimal compensation depends on the nature of information arrival as well as agent characteristics and derive implications for the maturity structure of securities in financial contracting settings.

Key words: Compensation design, Duration of pay, Moral hazard, Persistence, Principal-agent models, Informativeness principle.

JEL Codes: D86.

1. INTRODUCTION

In many real-life principal–agent relationships, informative signals about the agent's action are not all immediately available, but only arrive over time. For example, within the financial sector, investments by private equity or venture capital fund managers only produce verifiable returns to investors upon an exit, a credit rating issued by a credit rating agency or a loan officer's loan decision can be evaluated more accurately over the lifetime of a loan, and a bank's risk management is only stress-tested in times of crisis. Outside the financial sector, innovation activities by researchers, be it in academia or in industry, typically produce signals such as patents or citations only with considerable delay. Similarly, the quality of a CEO's strategic decisions has

The editor in charge of this paper was Veronica Guerreri.

persistent, long-lasting effects on outcomes and may not be assessed until well into the future. The list is certainly not exclusive and, yet, it suggests that the staggered arrival of performance signals over time is an important—if not defining feature—of many moral hazard environments. This article studies the basic question of how to optimally structure the *intertemporal* provision of incentives in such "*only time will tell*" information environments when the agent's relative impatience and consumption smoothing needs make it costly to defer compensation. We allow for general information systems and, yet, obtain a simple and intuitive characterization of optimal contracts, in particular of the *optimal duration of pay*. With bilateral risk neutrality and agent limited liability, optimal contracts are high-powered in that they only reward those outcomes that are most indicative of the agent having taken the desired action, and stipulate at most two payout dates. Our precise and tractable characterization allows for clear-cut predictions on the comparative statics of the optimal timing of pay. Consistent with empirical evidence on the determinants of (executive) pay duration across firms and industries (see *e.g.* Gopalan *et al.*, 2014) we find, for instance, that the optimal duration of pay is higher for firms with higher growth opportunities and more severe agency problems, but decreasing in the agent's outside option. Once we incorporate agent risk aversion, we obtain additional predictions, as contracts optimally smooth out payouts over a larger selection of dates and contingencies. The interaction of relative impatience and risk-aversion then implies an increasing performance hurdle over time and a decreasing pay-performance sensitivity (under decreasing absolute risk aversion).

To focus on the optimal *intertemporal* provision of incentives for general signal processes, we consider an otherwise parsimonious principal-agent setting with a one-time action in the spirit of Holmstrom (1979). The agent chooses an unobservable binary action that affects the distribution of a process of contractible signals, such as output realizations, defaults, annual performance reviews, etc. A compensation contract stipulates payments to the agent, conditioning on all information available at a particular date—the history of signals—and must satisfy both the agent's incentive (IC) and participation constraint (PC). We show that it is without loss of generality to specify compensation contingent on the likelihood ratio process, as induced by the underlying signal process. Our characterization of optimal contracts then only relies on the martingale property of this performance metric, which allows us to accommodate good news, bad news, discrete, and continuous signal processes within a unified framework. Note, while it may be natural that additional signals, such as, *e.g.*, stock prices, become noisier over time, the information contained in the *history* of signals, *e.g.*, the entire path of stock prices, must (weakly) improve due to the inclusive information structure.

We initially assume bilateral risk neutrality with agent limited liability, as is standard in workhorse models of financial contracting such as Innes (1990), DeMarzo and Fishman (2007), or DeMarzo and Sannikov (2006). This allows us to obtain clear-cut timing implications before analysing the additional effects of intertemporal consumption smoothing motives. The key simplification of the compensation design problem in this setting results from the fact that, in order to quantify the information benefit of deferral, we only need to keep track of the upper support of the likelihood ratio distribution, which is a deterministic, increasing function of time. The intuition for this result draws on insights from static models (*e.g.* Innes, 1990) and adapts them to our dynamic environment: if an optimal contract stipulates a bonus at some date t, then this bonus is only paid for an outcome, here a *history* of realized signals up to date t, that maximizes the likelihood ratio across all possible date-t realizations.[1] Then, by tracing out this maximal date-t likelihood ratio over time, we obtain a uni-dimensional metric that quantifies by how much

1. Without a risk-sharing motive, it is optimal to "punish" the agent for all other outcomes, *i.e.*, pay out zero due to limited liability.

the principal is able to reduce incentive pay by deferring longer.[2] Since the maximal likelihood ratio, thus, captures the value of the sequence of informative signals up to date-t to the principal, we refer to it, in a slight abuse of terminology, as the *informativeness* of the date-t information system under bilateral risk neutrality.

The timing of payouts trades off the information benefit of deferral—quantified by the increase in informativeness—with the deadweight costs resulting from the agent's relative impatience. When the agent's outside option is low, higher informativeness allows the principal to reduce the agency rent, and all pay is optimally concentrated at a single date that maximizes the impatience-discounted likelihood ratio. In turn, when the outside option is high enough so that the agent's participation constraint binds, the principal's rent-extraction motive is absent. Then, optimal contracts minimize deadweight impatience costs subject to incentive compatibility. Using simple convexification arguments, we show that optimal contracts *may* now require (at most) two payout dates — regardless of the underlying signal process. To build intuition for this key result, initially assume that signals arrive gradually over time with decreasing incremental informativeness (*i.e.* the maximal likelihood ratio is a strictly concave function of time). Then, as incremental (impatience) costs are increasing faster than the incremental (informativeness) benefit, it is optimal to pay at the first date at which informativeness is high enough to jointly satisfy both the incentive and participation constraint. Now, consider a perturbation of this information environment and assume that "a lot of" information, in the sense of a jump in informativeness, arrives immediately after the previously optimal payout date, *e.g.*, due to an additional performance review. Then, it is optimal to shift part of the incentive pay to a later date to tap the new informative signals of the review, which allows to simultaneously shift the remaining pay necessary to satisfy (PC) to an earlier date, thereby reducing impatience costs. Formally, the principal now benefits from using two payout dates as the cost of having access to time-t informativeness is no longer globally convex.[3]

While our simple model is not designed to match particular institutional details, its comparative statics implications are consistent with various stylized facts. For instance, it predicts that the duration of pay, the weighted average payout time, is shorter if the agent's outside option is higher. Thus, in industries with a high degree of competition for the agent's talent or services we would expect, ceteris paribus, shorter deferral periods. In contrast, the duration is predicted to be longer if the agency problem, as measured by the agent's benefits from shirking, is more severe. To the extent that one interprets better corporate governance as a way to reduce such shirking benefits (*e.g.* via improved monitoring as in Holmstrom and Tirole (1997)), we, thus, obtain the implication that corporate governance and optimal pay duration are substitutes. This is consistent with empirical evidence documenting that pay duration for executives is higher in companies with a higher entrenchment index (see Gopalan *et al.*, 2014). Further, our results shed light on how variation in the nature of information arrival across industries and tasks can help to explain observed patterns in the duration of executive pay. Gopalan *et al.* (2014) find that executive pay duration is longer in firms with more growth opportunities and a higher R&D intensity (see also Baranchuk *et al.*, 2014). This is consistent with our model as firms in such industries typically receive relatively little information in early stages and exhibit high informativeness growth in later development stages.

2. Due to the martingale property of the likelihood ratio process this metric is increasing in time.

3. As convexification of this cost requires at most two payout dates, the principal does not benefit from paying at additional dates. One may now conjecture that it is always optimal to shift the pay that primarily targets the participation constraint to date 0, but this conjecture is wrong. For example, suppose that "a lot" of valuable information is released at date ϵ, immediately after date 0. Then it is optimal to make all pay contingent on this almost costless information, which allows to reduce the size of more expensive long-term bonus payments.

We then augment our modelling framework, adding financial contracting frictions in the form of upper bounds on payouts to the agent. We show how persistent effects in moral hazard settings generate novel predictions for optimal security design (building on Innes (1990) and Hébert (2018)), in particular the optimal *maturity structure* of an entrepreneur's financing decisions. The entrepreneur optimally receives positive payouts if and only if performance (measured in likelihood ratio units) is above a cut-off that is increasing in time. This increasing performance hurdle implies a rich dynamic payoff structure for the entrepreneur's and investors' claims, resulting from the trade-off between the entrepreneur's liquidity needs and the information gain associated with additional performance signals available to investors.

Once we allow for risk aversion, the benefits of deferral, in terms of a more precise performance measurement, are unaffected. However, due to the agent's desire to smooth consumption across time and states, it is no longer optimal to concentrate pay at a single (or two) payout dates following the realization of the signal history with the maximal likelihood ratio. Intuitively, starting from the "risk-neutral" payout time and contingency, it is now optimal to gradually spread out consumption across time and states. As is standard, for a given t, the optimal contract stipulates higher rewards for higher performance. However, due to the costs associated with relative impatience, the likelihood ratio needed to maintain the same level of pay at future dates is increasing with time. Next to this increasing performance hurdle, we find that the pay-performance sensitivity is decreasing over time if the agent's preferences exhibit decreasing absolute risk aversion, while it is increasing with time for the case of increasing absolute risk-aversion.

Literature. The premise of our article is that the timing of pay determines the information about the agent's hidden action that the principal can use for incentive compensation. This relates our analysis to the broader literature on comparing information systems in agency problems, which derives sufficient conditions for information to have value for the principal (Holmstrom, 1979; Gjesdal, 1982; Grossman and Hart, 1983; Kim, 1995). Time generates a family of inclusive information systems via the arrival of additional signals, so that these can be ranked both in the sense of Holmstrom (1979) and Kim (1995). The key difference of our article relative to this classical strand of the literature is that having access to a better information system generates endogenous costs due to the agent's relative impatience and consumption smoothing concerns. This trade-off determines the information system in equilibrium, which relates our paper to the literature on information design (see Kamenica and Gentzkow, 2011; Bergemann and Morris, 2016), and, in particular, its applications to moral hazard settings, Georgiadis and Szentes (2020) and Li and Yang (2020).

More concretely, our paper belongs to a small, but growing literature that analyses moral hazard setups in which the agent's action has persistent effects. Hopenhayn and Jarque (2010) analyse optimal contracts in a discrete-time setting with a risk-averse and equally patient agent. While they obtain some characterization for an example with i.i.d. binary signals, their model does not generate concrete implications for the *timing of pay*, which is the focus of our article. In particular, we obtain a precise characterization of the optimal timing of pay for general signal processes. Our analysis, thus, nests the important finance application of moral hazard by a securitizer of defaultable assets, as studied in Hartman-Glaser *et al.* (2012) and Malamud *et al.* (2013).

In our setup, deferral improves the information available to the principal about the agent's initial action. In standard repeated-action settings with instantaneously observable performance signals backloading of rewards instead has the benefit of relaxing (future) incentive and limited liability constraints (see *e.g.* DeMarzo and Sannikov, 2006; Biais *et al.*, 2007; DeMarzo and Fishman, 2007; Sannikov, 2008). Work by Jarque (2010), Edmans *et al.* (2012), Sannikov (2014), or Zhu (2018) combines the effects of repeated private actions and persistence.

The additional complexity, however, requires special assumptions on the signal process. Instead, our setup tries to isolate one effect, the idea that information gets better over time, and studies it in full generality.

2. MODEL SETUP

We consider a principal-agent problem in which the principal observes informative signals about the agent's action over time. Time is continuous $t \in [0, \bar{T}]$.[4] At time 0, the agent A takes an unobservable action $a \in \mathscr{A} = \{a_L, a_H\}$. We denote by a_H the high-cost action which comes at cost k_H, and by a_L the low-cost action with respective cost $k_L = k_H - \Delta k \geq 0$, where $\Delta k > 0$ (see Supplementary Appendix and the application in Hoffmann *et al.* (2020) for a continuous action set). As is standard, we suppose that the principal P wants to implement the high action.

Signals and information. The one-time action a affects the distribution of a stochastic process of verifiable signals X_t that may arrive continuously or at discrete points in time.[5] These abstract signals may correspond to output realizations, annual performance reviews by the principal, or, more generally, any multidimensional combination of informative signals. Formally, we consider a family of filtered probability spaces $(\Omega, \mathcal{F}^X, (\mathcal{F}_t^X)_{0 \leq t \leq \bar{T}}, \mathbb{P}^a)$ indexed by the agent's action a and satisfying the usual conditions (see *e.g.* Jacod and Shiryaev (2003)). Here, $\mathcal{F}_t^X$ refers to the filtration generated by X_t and $\mathbb{P}^a$ denotes the probability measure induced by action a, where we assume $\mathbb{P}^L$ to be absolutely continuous with respect to $\mathbb{P}^H$, which is denoted by $\mathbb{P}^L \ll \mathbb{P}^H$. The following three illustrative examples represent information environments covered by our framework.[6]

Example 1 *At each $t \in \{1, 2, ..., \bar{T}\}$ there is a binary signal $x_t \in \{s, f\}$ that is drawn independently over time. The date-t probability of success "s" equals $\frac{1}{2}$ under a_L and $1 - \rho^t$ under a_H where $\rho \in \left(\frac{1}{2}, 1\right)$.*

Example 2 *X_t is a multivariate counting process where $x_t^{(j)} = 1$ indicates that failure on element j has occurred before time t ($x_t^{(j)} = 0$ otherwise). The action $a \in \mathscr{A}$ affects the joint distribution $G(x|t, a)$.*

Example 3 *The agent's action determines the drift of an arithmetic Brownian motion,*

$$dx_t = a\rho^t dt + \sigma dZ_t,$$

where $\sigma > 0$, $\rho \in (0, 1]$ and dZ_t is a standard Wiener process.

The examples illustrate various manifestations of persistence, both in technical (discrete versus continuous signal processes), as well as in economic terms (learning from "good news" versus "bad news"). The discrete Example 1 captures the idea that short-run successes may be indicative

4. The assumption of a finite horizon $\bar{T}$ is not crucial. One can think of $\bar{T}$ as the last date at which informative signals arrive, or, alternatively, the last possible date the agent can be compensated.

5. Formally, the index set of the stochastic process X_t can be any measurable subset of $[0, \bar{T}]$.

6. Example 1 captures in reduced-form central features of Manso (2011) and Zhu (2018). Example 2 is a generalization of Hartman-Glaser *et al.* (2012) and Malamud *et al.* (2013) allowing for arbitrary correlation structures between failure events. Example 3 is a one-time action version of Sannikov (2014).

of the agent having chosen the short-term action a_L rather than the desired long-run action a_H, as, for $t < -\ln 2/\ln(\rho)$, the probability of success is higher under a_L. In Example 2, one may think of a loan officer granting loans, whose defaults are correlated via macroeconomic conditions. Here, learning takes place via (the absence of) failures. Finally, Example 3 considers the canonical setting where the agent's action determines the drift of an arithmetic Brownian motion, albeit with decaying impact.

Each signal process X_t induces a corresponding likelihood ratio process L_t which captures all the information about the agent's action contained in the entire *history* of realized signals $h^t = \{x_j\}_{0\le j\le t}$ up to t. It will, thus, be the key date-t performance measure in our subsequent analysis of optimal compensation design.[7] Let $\mathbb{P}^a_t$ denote the restriction of the probability measure $\mathbb{P}^a$ to date-t information $\mathcal{F}^X_t$, then the likelihood ratio process is defined as:

$$L_t := 1 - \frac{d\mathbb{P}^L_t}{d\mathbb{P}^H_t}, \tag{1}$$

where existence of the Radon–Nikodym derivative $\frac{d\mathbb{P}^L_t}{d\mathbb{P}^H_t}$ follows from the Radon–Nikodym Theorem given $\mathbb{P}^L \ll \mathbb{P}^H$. To ensure statistical identifiability of the actions a_L and a_H, we assume that some information arrives by date $\bar{T}$, *i.e.*, $\bar{L}_{\bar{T}} > 0$, where $\bar{L}_t := \sup_{\omega\in\Omega} L_t(\omega) \le 1$ refers to the supremum of the date-t likelihood ratio distribution.

For discrete processes, for which each outcome has positive probability mass, the definition in (1) corresponds to the standard, "text-book" likelihood ratio (see *e.g.* Tirole, 2006): For instance, in Example 1, the date-1 likelihood ratio takes on the value of $1-\frac{1/2}{1-\rho}$ upon realization of a success signal "s" and a value of $1-\frac{1/2}{\rho}$ in the event of a failure "f" with respective probabilities $1-\rho$ and ρ given action a_H. In contrast, for the continuous Brownian-motion Example 3 with $\rho = 1$, Girsanov's theorem implies that the likelihood ratio can be described by the relative density process $L_t = 1 - \exp\left[-\frac{a_H - a_L}{\sigma^2}\left(x_t - \frac{a_H + a_L}{2}t\right)\right]$, where $\frac{a_H - a_L}{\sigma^2}$ may be interpreted as the signal-to-noise ratio and $x_t \sim N\left(a_H t, \sigma^2 t\right)$ given a_H.

Contracts and preferences. In order to implement the high action, the principal designs a compensation contract $\mathscr{C}$ that stipulates transfers to the agent as a function of the information available at the time of payout. Formally, a contract can then be represented by a cumulative compensation process b_t that is càdlàg and adapted to $\mathcal{F}^X_t$. We restrict attention to payout processes that satisfy agent limited liability, *i.e.*, $db_t \ge 0$ and assume that the principal is able to commit to any such contract.[8] As is common in dynamic principal-agent models, the principal's and agent's valuation of these transfers differ in two ways: first, we stipulate that the agent is relatively impatient, which makes it costly to defer pay.[9] That is, the discount rates of the agent, r_A, and the principal, r_P, satisfy

$$\Delta r := r_A - r_P > 0.$$

Second, while the principal is risk-neutral, we consider both the case of a risk-neutral as well as the one of a risk-averse agent. With bilateral risk neutrality (see Section 3), the instantaneous

7. The importance of likelihood ratios for optimal compensation design is well established in the literature at least since Holmstrom (1979).

8. Since in our model the preferences of the principal and the agent differ with regards to the valuation of payouts across time and states, there else would be scope for ex post renegotiation of the ex ante optimal contract (see *e.g.* Fudenberg and Tirole (1990) or Hermalin and Katz (1991)).

9. See, e.g., DeMarzo and Duffie (1999); DeMarzo and Sannikov (2006), or Opp and Zhu (2015).

utility transfer to the agent, denoted by dv_t, satisfies $dv_t := db_t$. When the agent is risk-averse with a strictly increasing and strictly concave utility function u defined over consumption flows (see Section 4), the contract specifies the consumption flow to the agent, *i.e.*, $db_t = c_t dt$ so that $dv_t := u(c_t)dt$ with c_t finite.

Hence, the principal's compensation design problem reads as follows:

Problem 1

$$W := \min_{b_t} \mathbb{E}^H \left[\int_0^{\bar{T}} e^{-r_P t} db_t \right] \quad s.t. \tag{W}$$

$$\mathbb{E}^H \left[\int_0^{\bar{T}} e^{-r_A t} dv_t \right] - k_H \geq R, \tag{PC}$$

$$\mathbb{E}^H \left[\int_0^{\bar{T}} e^{-r_A t} dv_t \right] - \mathbb{E}^L \left[\int_0^{\bar{T}} e^{-r_A t} dv_t \right] \geq \Delta k, \tag{IC}$$

$$db_t \geq 0 \quad \forall t, \tag{LL}$$

where $\mathbb{E}^a$ *denotes the expectation under probability measure* $\mathbb{P}^a$.

The principal chooses a compensation contract satisfying agent limited liability (LL) to minimize the present value of wage cost W (discounted at the principal's rate). The first constraint is the agent's time-0 participation constraint (PC): the present value of utility transfers discounted *at the agent's rate* net of the cost of the action must at least match her reservation utility R.[10] Second, incentive compatibility (IC) requires that it is optimal for the agent to choose action a_H given $\mathscr{C}$. A standard change of measure together with the law of iterated expectations (using (1)) then allows us to rewrite (IC) succinctly as:

$$\mathbb{E}^H \left[\int_0^{\bar{T}} e^{-r_A t} L_t dv_t \right] \geq \Delta k. \tag{IC*}$$

Since the only reason for making *contingent* payments is to provide incentives and preferences are time-separable, (IC*) highlights that it is without loss of generality to restrict attention to compensation processes that are adapted to the filtration generated by the likelihood ratio process.[11]

We note that similar to standard static principal-agent settings, such as the special case of our model with $\bar{T} = 0$, general existence of a solution to Problem 1 requires additional assumptions. In the subsequent analysis, we provide *sufficient* conditions on either the information process (Condition 1), bounds on transfers (Condition 2) or the utility function (Condition 3) that guarantee existence of a solution to Problem 1 with a risk-neutral or risk-averse agent, respectively.

10. Since the agent in our model chooses an action once at time 0 and is protected by limited liability, the participation constraint of the agent only needs to be satisfied at $t = 0$.

11. See the proof of Lemma A.1 in the Appendix for details including the steps in deriving (IC*).

Remark 1 *Consider the following adverse selection problem: Agents differ in their ex ante type $\theta \in \{\theta_H, \theta_L\}$, inducing different probability measures $\mathbb{P}^{\theta}$ over signals X_t, and have type-dependent outside options $R_H := R + k_H > R_L := R + k_L > 0$. Then, if the principal only wants high-type agents to participate and R is sufficiently high, the optimal screening contract is characterized by the solution to Problem 1.*[12]

3. CONTRACTING WITH RISK-NEUTRAL AGENT

3.1. *Baseline model*

We initially consider a baseline setting with bilateral risk neutrality. In this setting, the principal's optimal choice of payout times results from a basic trade-off between access to a better information system and the deadweight costs resulting from differential discounting. In particular, the costs of deferring pay are measured by *impatience costs*, $e^{\Delta rt}$, corresponding to the ratio of the principal's and the agent's respective valuations of any date-t transfer. In turn, the key step in quantifying the information benefit of deferral is to note that—with bilateral risk neutrality and limited liability—the value of the date-t information system to the principal can be conveniently summarized by $\bar{L}_t$, the maximal likelihood ratio across all date-t histories. The intuition is straightforward from (IC*): as date-t rewards contingent on histories with higher likelihood ratios provide stronger incentives and risk-sharing is irrelevant, $\bar{L}_t$ is the statistic of the likelihood ratio process that determines minimal compensation costs and we accordingly refer to it as date-t "*informativeness.*" We can, thus, quantify the benefit of deferral by the increase in "*informativeness*" over time—thereby formalizing the notion of "time will tell."

Observation 1 *Informativeness $\bar{L}_t$ is a weakly increasing function of time.*

Intuitively, observing additional signals over time cannot reduce the precision of performance measurement. Formally, the result follows from the fact that L_t is a martingale under $\mathbb{P}^H$ (with an unconditional expectation $\mathbb{E}^H[L_t] = 0$ for all t), so that its maximum must be an increasing function of time. It is easiest to illustrate the well-known martingale property by rewriting (1) as $L_t = 1 - \mathbb{E}^H\left[\frac{d\mathbb{P}^L}{d\mathbb{P}^H}\middle|\mathcal{F}_t^X\right]$ and applying the law of iterated expectations.

As one would expect, "time will tell" can manifest itself economically in many different ways, resulting in smooth informativeness functions (as in Example 2 where good news arrives gradually over time) or step functions (as in settings with lumpy information such as Example 1) or any combination thereof. In fact, for any weakly increasing function $I:[0,\bar{T}] \to [0,1]$ one can construct a signal process X_t such that the associated informativeness function satisfies $\bar{L}_t = I(t)$.[13] This suggests that we can think of $\bar{L}_t$ and its distribution—rather than the signal process X_t—as the primitive of the information environment. To ensure existence of an optimal contract in this setting, we now require

Condition 1 *The maximal-likelihood-ratio event for each date t, $L_t = \bar{L}_t$, occurs with strictly positive probability under action a^H, i.e., $\mathbb{P}^H\left(L_t = \bar{L}_t\right) > 0$.*

Condition 1 is an essential ingredient for the subsequent Theorem 1 as it rules out Mirrleesian existence problems, that, for example, would arise in the Brownian motion setting of Example 3.

12. We thank an anonymous referee for this insight.

13. See the proof of Lemma A.2 in the Appendix for a construction.

(Example 3 will still be relevant in subsequent sections where we introduce bounds on transfers and agent risk aversion).

Theorem 1 *Suppose, $\bar{L}_0 \leq \frac{\Delta k}{R+k_H}$ and Condition 1 holds. Then, an optimal contract exists and is high-powered. The agent is paid at time t, $db_t > 0$, if and only if (i) $L_t = \bar{L}_t$ and (ii) t belongs to the set of optimal payout dates which is characterized as follows: Denote by $\hat{T}$ the (earliest) date that maximizes the discounted likelihood ratio, i.e.,*

$$\hat{T} := \min\{\arg\max_t e^{-\Delta r t} \bar{L}_t\}. \tag{2}$$

1. If $R \leq \bar{R} = \Delta k / \bar{L}_{\hat{T}} - k_H$, (PC) is slack and it is optimal to pay only at date $\hat{T}$ with $db_{\hat{T}} = \frac{e^{r_A \hat{T}}}{\mathbb{P}^H(L_{\hat{T}} = \bar{L}_{\hat{T}})} \frac{\Delta k}{\bar{L}_{\hat{T}}}$ contingent on $L_{\hat{T}} = \bar{L}_{\hat{T}}$.
2. Else, (PC) binds and an optimal contract requires at most two distinct payout dates, that are bounded above by $\hat{T}$.

Regardless of whether the agent's participation constraint binds or not, the contingency of pay in optimal contracts is, thus, chosen to provide maximal incentives: Due to the absence of risk-sharing considerations, the agent is only rewarded for histories that achieve the highest likelihood ratio given this action, $\bar{L}_t$ and obtains zero for all other outcomes due to limited liability. In contrast, the optimal timing of pay depends crucially on whether the principal has a rent-extraction motive ((PC) slack) or not. Before describing this optimal timing choice in more detail, it is useful to discuss the two qualifying conditions in Theorem 1. First, we have imposed an upper bound on $\bar{L}_0$ solely for expositional reasons to focus on the relevant case when the agent's (IC) constraint binds.[14] Second, to illustrate why an optimal contract might not exist when Condition 1 is violated, consider the case with (PC) slack: Then, if $\mathbb{P}^H(L_t = \bar{L}_t) = 0$, the optimal bonus $db_t = \frac{e^{r_A \hat{T}}}{\mathbb{P}^H(L_{\hat{T}} = \bar{L}_{\hat{T}})} \frac{\Delta k}{\bar{L}_{\hat{T}}}$ would be infinite (see Section 3.2 on how appropriate payment bounds can restore existence of a solution).

Discussion of timing PC slack. When (PC) is slack the optimal timing of pay reflects the principal's *rent-extraction* motive: the principal can reduce the utility transfer to the agent, by deferring longer and, hence, using more informative performance signals. However, deferral does not imply a zero-sum transfer of surplus to the principal, but instead involves deadweight costs due to relative impatience. The optimal payout time resolves this trade-off by maximizing the "*discounted informativeness*," where the discount rate Δr reflects the effective cost of deferral per unit of time. If the maximum likelihood ratio $\bar{L}_t$ is differentiable at $\hat{T}$, we obtain the intuitive characterization

$$\left.\frac{d \log \bar{L}_t}{dt}\right|_{t=\hat{T}} = \Delta r. \tag{3}$$

That is, the principal defers until the growth rate of informativeness, $\frac{d \log \bar{L}_t}{dt}$, equals the growth rate of impatience costs, Δr, resulting in wage costs of $W = \frac{e^{\Delta r \hat{T}}}{\bar{L}_{\hat{T}}} \Delta k$.

14. If $\bar{L}_0 > \frac{\Delta k}{R+k_H}$, (IC) is slack. Then, there are a continuum of optimal contracts all of which pay the agent at $t=0$ only and generate wage costs of $W = R + k_H$ for the principal. For example, one optimal contract is to pay the agent $\frac{R+k_H}{\mathbb{P}^H(L_0 = \bar{L}_0)}$ at date 0 whenever $L_0 = \bar{L}_0$.

Of course, this characterization only applies as long as (PC) is satisfied, *i.e.*, $\Delta k/\bar{L}_{\hat{T}} > R+k_H$. Notably, this is a necessary and sufficient condition for (PC) to be slack also in the knife-edge case with multiple global maximizers of $e^{-\Delta rt}\bar{L}_t$. Intuitively, while the principal is indifferent among all compensation packages that minimize wage costs, the agent strictly prefers the one with the highest rent, *i.e.*, the one associated with the least informative performance signal and, hence, earliest payout time (see (2)).

Discussion PC binds. When the outside option R exceeds the threshold $\bar{R} := \Delta k/\bar{L}_{\hat{T}} - k_H$, the binding participation constraint implies that the principal has to offer a compensation package of value $R+k_H$ to the agent. Since the principal's rent-extraction motive is, thus, absent, his objective is now to maximize joint surplus, *i.e.*, to minimize deadweight impatience costs, subject to ensuring incentive compatibility. To understand the optimal timing of pay in this case in more detail, it is useful to express the principal's wage cost in (W) as the product of the compensation value to be provided to the agent, $(R+k_H)$, and weighted average impatience costs $\int_0^{\bar{T}} e^{\Delta rt} dw_t$:

$$W = (R+k_H) \min_{w_t} \int_0^{\bar{T}} e^{\Delta rt} dw_t. \tag{4}$$

Here, the new choice variable $dw_t := \frac{\mathbb{E}^H[e^{-r_A t} db_t]}{R+k_H} \geq 0$ can be economically interpreted as the fraction of the total compensation value that the agent derives from date-t payouts, implying that $\int_0^{\bar{T}} dw_t = 1$. Moreover, (IC*) requires that contractual incentives—the product of the compensation value to the agent, $(R+k_H)$, and the weighted average informativeness $\int_0^{\bar{T}} \bar{L}_t dw_t$—satisfy

$$(R+k_H) \int_0^{\bar{T}} \bar{L}_t dw_t = \Delta k.^{15} \tag{5}$$

This change of variables now clearly highlights that the optimal calibration of bonus payout times achieves a given *weighted average* informativeness of $\int_0^{\bar{T}} \bar{L}_t dw_t = \frac{\Delta k}{R+k_H}$ at lowest *weighted average* impatience costs, $\int_0^{\bar{T}} e^{\Delta rt} dw_t$.[16]

Using elementary tools of convex analysis, we can now provide an intuitive graphical proof that illustrates why any optimal contract solving Problem 1 requires at most two payout dates. Starting from informativeness $\bar{L}_t$ and impatience costs $e^{\Delta rt}$ as functions of time (see left panel of Figure 1), in a first step, plot for each t the implied impatience costs $e^{\Delta rt}$ (on the vertical axis) against $\bar{L}_t$ (on the horizontal axis), see red dotted line in the right panel. The (grey-shaded) convex hull of this parametric curve traces out all possible combinations of weighted average informativeness $\int_0^{\bar{T}} \bar{L}_t dw_t$ and weighted average impatience costs $\int_0^{\bar{T}} e^{\Delta rt} dw_t$ that can be achieved by distributing incentive payments over time. Since the objective in (4) is to *minimize* weighted average impatience costs, only the *lower hull* is relevant (see solid black line in right panel). This lower hull traces out a strictly increasing function which we denote by $C: [\bar{L}_0, \bar{L}_{\bar{T}}] \to [1, e^{\Delta r\bar{T}}]$.[17] Economically, this function can be interpreted as the *cost of informativeness*, the minimum

15. This follows directly from (IC*) and the definition of dw_t, exploiting the fact that if a payment is made at t then only for $L_t = \bar{L}_t$.

16. Note, since we assume that date-0 informativeness is bounded above, $\bar{L}_0 < \frac{\Delta k}{R+k_H}$, it is not feasible for the principal to avoid deadweight costs altogether ($dw_0 = 1$) and satisfy (5) at the same time.

17. This follows from the fact that both $\bar{L}_t$ and $e^{\Delta rt}$ are increasing in t, the latter always strictly.

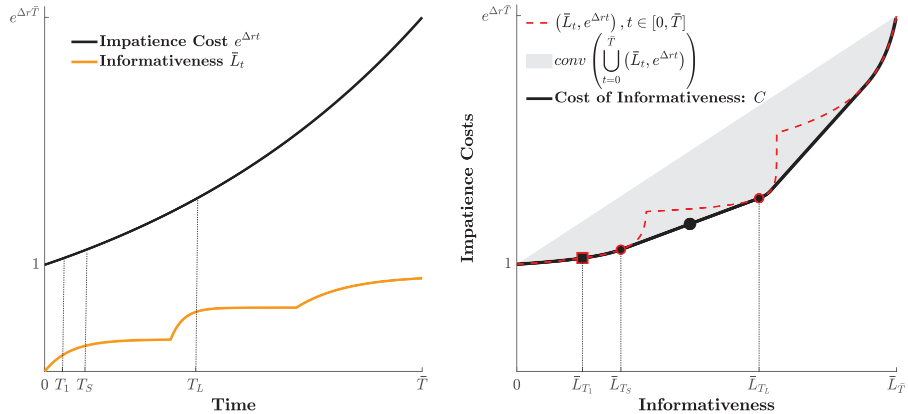

FIGURE 1

Graphical proof for optimality of at most 2 payout dates

Notes: The left panel plots informativeness, $\bar{L}_t$, and impatience costs $e^{\Delta rt}$ as functions of time. The right panel plots $e^{\Delta rt}$ against $\bar{L}_t$ as a parametric curve. The grey shaded region represents the convex hull of these points. Finally, the black solid line refers to the lower hull, C.

impatience cost required to achieve a given level of weighted average informativeness. As incentive compatibility in (5) requires a weighted average informativeness of $\frac{\Delta k}{R+k_H}$, minimum wage costs in (4) are, thus, given by $W=(R+k_H)C\left(\frac{\Delta k}{R+k_H}\right)$.

The reason for why two payout dates are sufficient is then immediate: any point on the lower hull C either corresponds to $\left(\bar{L}_t, e^{\Delta rt}\right)$ for some t (such as the point marked by the square in the right panel of Figure 1), or can be expressed as a convex combination of two such points corresponding to distinct dates t (such as the point marked by the black circle).[18] In the plotted example the number of payout dates, thus, depends on the value of $\frac{\Delta k}{R+k_H}$ as, *e.g.*, governed by the outside option R. The following lemma summarizes these insights.

Lemma 1 *If $R>\bar{R}$, the principal faces wage costs of $W=(R+k_H)C\left(\frac{\Delta k}{R+k_H}\right)$. The associated optimal contract can be characterized as follows:*
(1) ***Single-date:*** *The optimal contract can be implemented with a single payout date T_1 if and only if there exists a date T_1 such that $\bar{L}_{T_1}=\frac{\Delta k}{R+k_H}$ and $C\left(\frac{\Delta k}{R+k_H}\right)=e^{\Delta rT_1}$. The contract specifies $db_{T_1}=(R+k_H)\frac{e^{r_A T_1}}{\mathbb{P}^H(L_{T_1}=\bar{L}_{T_1})}$ if $L_{T_1}=\bar{L}_{T_1}$ and zero otherwise.*
(2) ***Two-dates:*** *Otherwise, the contract requires a short-term payout date T_S and a long-term date T_L such that $T_L>T_S\geq 0$. Let $w_S:=\frac{\bar{L}_{T_L}-\frac{\Delta k}{R+k_H}}{\bar{L}_{T_L}-\bar{L}_{T_S}}\in(0,1)$ denote the value-weight of the payment at T_S, then the respective payout dates, T_S and T_L, solve $w_S e^{\Delta rT_S}+(1-w_S)e^{\Delta rT_L}=C\left(\frac{\Delta k}{R+k_H}\right)$. The long-term date T_L is characterized by:*

$$T_L=\underset{t>T_S}{\arg\max}\frac{\bar{L}_t-\bar{L}_{T_S}}{e^{\Delta rt}-e^{\Delta rT_S}}. \tag{6}$$

18. Formally, any point on the lower hull C is either an extreme point itself or can be expressed as a convex combination of two extreme points (cf., Krein and Milman, 1940).

The contractual payments db_{T_S} and db_{T_L} following realization of $L_{T_S}=\bar{L}_{T_S}$ and $L_{T_L}=\bar{L}_{T_L}$ are then obtained from (5).[19]

The optimal timing of pay in the single-date contract corresponds to the earliest payout date that allows to provide sufficient incentives with binding (PC), reflecting the principal's objective to minimize deadweight impatience costs subject to (IC). Under the two-date contract, this trade-off is reflected in the principal's choice of the pair (T_S, T_L) linked by the following intuitive property: The long-term date T_L globally maximizes the ratio of the *incremental* informativeness over the *incremental* impatience costs relative to date T_S as reflected in (6).

Given that we have now formally established why an optimal contract requires at most two payout dates, we now shed further light on how the number of payout dates depends on the economic properties of information arrival.

Corollary 1 *Conditions for one versus two payout dates.*
(1) Suppose that $e^{\Delta rt}$ is strictly convex relative to $\bar{L}_t$, i.e., $\frac{\frac{d^2\bar{L}_t}{dt}}{\frac{d\bar{L}_t}{dt}} < \Delta r$ for all $t < \hat{T}$, then an optimal contract must specify a single payout date.
2) Suppose $R > \bar{R}$ and $\bar{L}_t$ is a step function, then an optimal contract generically specifies two payout dates.

The number of payout dates depends on whether the information benefit of deferral, as captured by incremental informativeness, is increasing faster or slower than deadweight impatience costs. To build intuition consider, first, the case of a concave informativeness function $\bar{L}_t$ (see top left panel of Figure 2), which arises in environments where information arrives gradually over time with decreasing incremental informativeness, e.g., because the agent's initial action has a decaying effect. Then as the growth rate of informativeness is decreasing while impatience costs grow at a constant rate of Δr, the first condition in Corollary 1 is immediately satisfied and the cost of informativeness C is a strictly convex function (see top right panel of Figure 2). In such decreasing-returns-to-deferral information environments, strictly increasing *marginal costs* of informativeness C' imply that it is cheapest to concentrate all pay at the earliest point in time for which informativeness is sufficient to satisfy (IC). One may have conjectured that the optimal contract with binding (PC) stipulates an incentive payment at $\hat{T}$—the optimal payout date with slack (PC) maximizing discounted informativeness—and an additional sufficiently high date-0 payment to satisfy (PC) at lowest possible impatience costs. As the informativeness level with slack (PC), $\bar{L}_{\hat{T}}$, can be graphically identified by the point on $\left(\bar{L}_t, e^{\Delta rt}\right)$ that minimizes the slope of a ray through the origin (see green dotted line in the right panels of Figure 2),[20] the graph shows that this particular conjecture is wrong whenever C is strictly convex (compare the point indicated by the "circle" to the "star" in top right panel).

In contrast, two payout dates are required in environments where incremental informativeness increases strictly faster than impatience costs for some t, such that the first condition in Corollary 1 is violated. For instance, in the extreme case where $\frac{d^2\bar{L}_t}{dt}/\frac{d\bar{L}_t}{dt} > \Delta r$ holds for all t, the implied parametric curve (plotting $e^{\Delta rt}$ against $\bar{L}_t$) is globally concave and it is generically optimal to convexify by paying out at dates 0 and $\bar{T}$. Else, if informativeness grows faster than impatience costs over some intervals and slower over others, the number of payout dates depends on the

19. See the proof of Lemma 1 for an explicit characterization.

20. This follows directly from the optimal timing condition with slack (PC) in (2). The minimum slope can be interpreted as the shadow price on (IC), κ_{IC}.

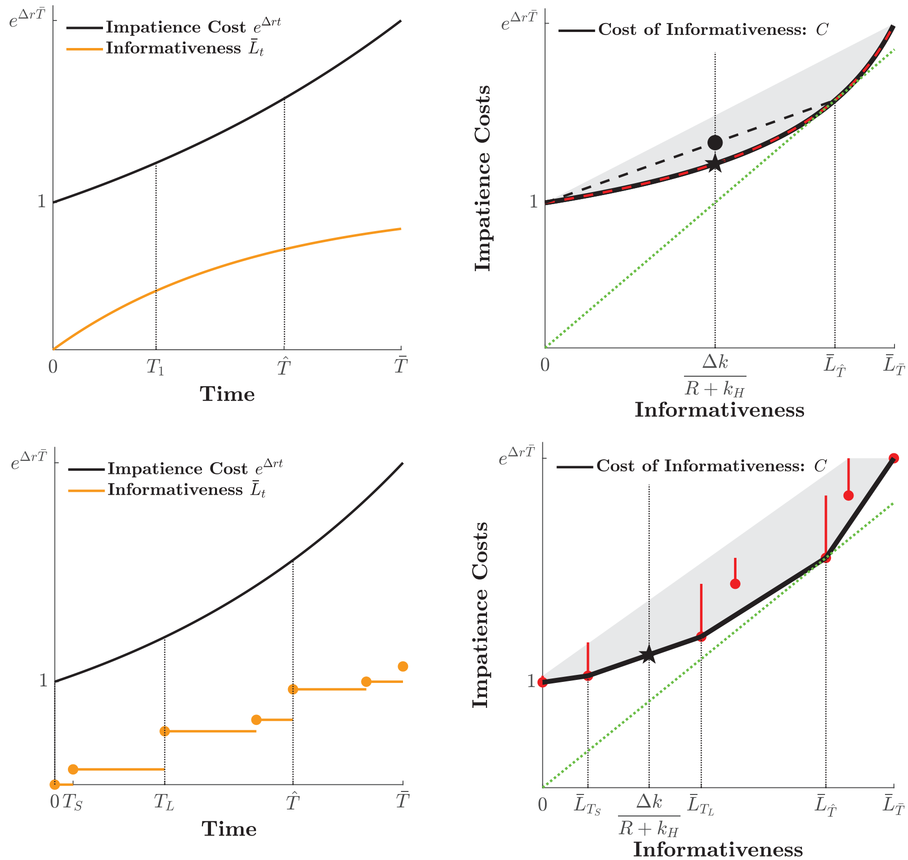

FIGURE 2

Convexification benefits and the number of payout dates

Notes: The upper panels plot a case with concave informativeness implying a strictly convex C (so that any given level of informativeness is optimally achieved with one payout date). In the lower panels, informativeness is a step function such that contracts with two payout dates are strictly optimal almost everywhere.

level of informativeness required by (IC) $\frac{\Delta k}{R+k_H}$ (compare, for instance, the contracts marked by a circle and square in Figure 1). This intuition extends to economic settings where information arrives only at discrete points in time, say via periodical performance reviews, such that the second condition in Corollary 1 applies (see lower panel of Figure 2). Due to the lumpiness of information arrival—$\bar{L}_t$ only taking on a finite number of values (see right panel)—two payout dates are generically needed to achieve the required average level of informativeness $\frac{\Delta k}{R+k_H}$.[21] Here, the potential payout dates T_S and T_L can be graphically identified as vertices of C.

Implications for the duration of pay. Now, regardless of whether there are one or two payout dates, the lack of a rent-extraction motive with binding (PC) intuitively implies that it is optimal to reduce deadweight impatience costs by paying the agent earlier than in the case with slack (PC). More formally, the *duration of the compensation contract*, $\int_0^{\bar{T}} t dw_t$, measuring the weighted

21. A single payout date is only possible in the knife-edge case when $\bar{L}_t = \frac{\Delta k}{R+k_H}$.

average payout time, is shorter with binding than with slack (PC).[22] In fact, Theorem 1 goes one step further by establishing that with binding (PC) all payouts occur weakly before $\hat{T}$, the payout date with slack (PC).[23] These insights can now be readily extended to study the comparative statics of pay duration with regards to the outside option.

Corollary 2 *The duration of an optimal compensation package $\int t dw_t$ is decreasing in R and increasing in Δk.*

The comparative statics in R and Δk follow from the fact that higher pay and more informative signals are substitutes for providing incentives to the agent. When an increase in the agent's binding outside option R exogenously raises the value of the compensation package to the agent, say due to high competition for the agent's talent, this substitutability implies that the principal optimally shortens the duration of the compensation package, such as to reduce contract informativeness (strictly so if (PC) binds). In contrast, if the agency problem gets more severe, *i.e.*, Δk increases, then the principal relies on both more informative performance signals via a longer pay duration and a larger compensation package. To the extent that poorer corporate governance implies a more severe agency problem, *e.g.*, as weaker monitoring increases the benefits of shirking, our model, thus, predicts a substitutability of corporate governance and optimal pay duration.

While the parameters R and Δk primarily characterize the agent's type or the difficulty of the task at hand, comparative statics in the informativeness function $\bar{L}_t$ should primarily capture variation in the nature of information arrival across tasks or industries. To this end, we consider a parametric family of informativeness functions to highlight the distinct effects of *growth* and *level* of informativeness on the optimal timing of pay.

Corollary 3 *Suppose $\bar{L}_t = \iota e^{\phi g(t)}$ where ι, ϕ are strictly positive constants and $e^{\phi g(t)}$ is a concave, differentiable function. Then, as long as the payout date is interior, it is*
(i) strictly increasing in ϕ and independent of ι if (PC) is slack, and
(ii) strictly decreasing in ϕ and ι, otherwise.

Corollary 3 reveals how a binding participation constraint fundamentally alters the trade-offs in setting payout times. When (PC) is slack, it is in the interest of the principal to defer as long as the growth rate of informativeness ϕ exceeds Δr. Thus, a higher ϕ induces the principal to increase the payout date whereas the level parameter ι does not affect the principal's trade-off. In contrast, when (PC) binds, both a higher initial level and a higher growth rate allow the principal to achieve the informativeness *level* of $\frac{\Delta k}{R+k_H}$ required by (IC) at an earlier date. Hence, the principal responds by shortening the duration of the contract in response to increases in ι and ϕ. Corollary 3 can be readily extended to settings in which the growth rate of informativeness changes over time such as, *e.g.*, in R&D intensive industries with long development periods implying relevant informativeness growth only at later dates. Our analysis would then predict firms in such industries to exhibit longer pay duration compared to traditional industries in which

22. This duration measure is analogous to the *Macaulay* duration which is standard in the fixed-income literature; the weights of each payout date are determined by the present value of the associated payment divided by the agent's valuation of the compensation package (see *e.g.* Gopalan *et al.*, 2014).

23. This is easily seen from the graphical construction of T_S, T_L, and $\hat{T}$ above, together with the fact that (PC) is binding if and only if the weighted average informativeness with binding (PC) is lower than with slack (PC), *i.e.*, $R > \bar{R} \Leftrightarrow \frac{\Delta k}{R+k_H} < \bar{L}_{\hat{T}}$ (see also the proof of Theorem 1).

signals with substantial informativeness are available early on (see evidence in Baranchuk *et al.*, 2014; Gopalan *et al.*, 2014).

3.2. *Payment bounds and security design*

So far, the focus of our article was to provide a tractable characterization of the optimal timing of pay. However, in cases where the optimal contract in Theorem 1 prescribes high rewards for low-probability events, resource constraints such as limited liability on the side of the principal or regulatory constraints, such as bonus caps, may become relevant (see also Jewitt *et al.*, 2008). In financial contracting applications, incorporating such resource constraints into our setting can be viewed as a first step towards optimal security design under moral hazard with persistent effects, providing a characterization of the optimal maturity structure of financial claims. Concretely, we next consider optimal compensation design given the following upper bound on payments:

Condition 2 *For each t, payments to the agent are bounded above,* $db_t \leq \bar{b}dt$, *with* $\bar{b} > 0$.

Imposing an upper bound on the payment rate—next to agent limited liability which acts as a lower bound—allows us to drop Condition 1. Hence, we can extend our analysis to information settings where a solution to the original Problem 1 with bilateral risk neutrality does not exist, including, in particular, the Brownian motion Example 3. While the bound $\bar{b}$ may be time and/or state dependent reflecting, *e.g.*, varying financial resources of the principal or concerns about performance manipulation by the agent (see *e.g.* Innes (1990) for a resulting monotonicity constraint), we omit these possibilities for notational convenience.[24] However, we do need to assume that the payment bound is not too strict such that a_H remains implementable.[25]

Proposition 1 *Suppose that Condition 2 holds and the set of contracts strictly satisfying the constraints is non-empty. Then, an optimal contract exists and stipulates pay at the maximum rate* $\bar{b}$ *if and only if the realization of the likelihood ratio,* l_t, *satisfies*

$$\kappa_{IC} l_t + \kappa_{PC} \geq e^{\Delta r t}, \tag{7}$$

where $\kappa_{IC} \geq 0$ *and* $\kappa_{PC} \geq 0$ *denote the shadow cost on (IC) and (PC), respectively.*

The optimal contract with payment bounds is an intuitive extension of the contract obtained in the baseline setting (see Theorem 1). Initially, consider the case with (PC) slack, *i.e.*, $\kappa_{PC} = 0$, which occurs if the outside option is below a threshold $R \leq \bar{R}(\bar{b})$. Since the principal can no longer satisfy (IC) by relying exclusively on a reward for the history with the best impatience—informativeness trade-off, he selects the next best alternatives according to the discounted likelihood ratio $e^{-\Delta r t} l_t$, up to the value of $1/\kappa_{IC}$ that results in satisfying (IC) (this can be readily seen from (7) with $\kappa_{PC} = 0$). This cost–benefit trade-off results both in a wider selection

24. Similarly, the setup can be readily extended to allow for time/state dependency in the lower payment bound (instead of the constant zero lower bound imposed by agent limited liability) or a dependency of either payment bound on endogenously chosen payouts. For instance, in some settings payouts to the agent are downward rigid such that pay has to weakly increase over time. Imposing such a downward rigidity constraint then has the intuitive implication of a higher optimal pay duration. This is because higher pay today requires higher future pay regardless of future performance. We thank an anonymous referee for this insight.

25. Trivially, if $\bar{b} \to 0$, then the principal cannot provide incentives. A sufficient condition for implementability is that (IC) and (PC) are satisfied if the principal pays the maximum rate $\bar{b}dt$ at all dates t for all positive realizations of the likelihood ratio.

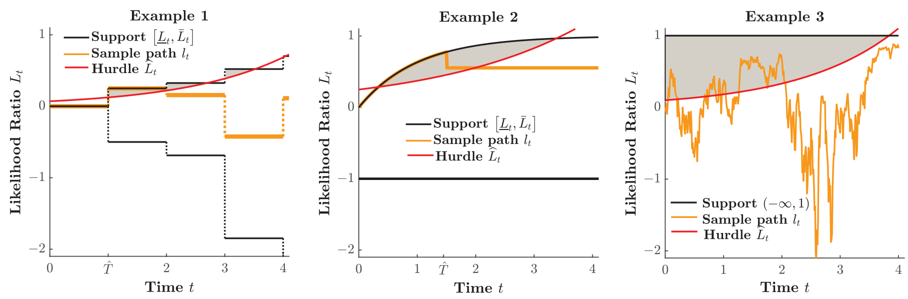

FIGURE 3

Payment bounds and payout dates for Examples 1 to 3

Notes: The graph plots the lower support, $\underline{L}_t$, and upper support, $\bar{L}_t$, of the likelihood ratio distribution as a function of time as well as one sample path l_t for specifications of the three information environments in Examples 1 to 3. In all graphs, the red line refers to the performance hurdle $\hat{L}_t$ for payouts (see Corollary 4). If the realized performance is above the performance hurdle, *i.e.*, in the grey-shaded region, then the agent obtains the maximal payout rate $\bar{b}$. The left panel corresponds to Example 1 with $\rho = 2/3$, the middle panel to a one-dimensional ($j=1$) specification of Example 2, where the arrival time distribution is exponential with parameters $\lambda_H = 1$ and $\lambda_L = 2$ given a_H and a_L respectively, while the right panel depicts Example 3 with $\rho = 1$, $\sigma = 2$, $a_H = 3/2$, and $a_L = -3/2$.

of payout dates and payout states compared to the case without payment bounds (see grey-shaded region in Figure 3 for our three examples with $\kappa_{PC} = 0$). When (PC) binds so that $\kappa_{PC} > 0$, this payout region is extended even further. For both cases, we obtain the following robust prediction for the pay-performance relation over time:

Corollary 4 *The minimal likelihood ratio required for a strictly positive payout at date t is given by* $\hat{L}_t = \frac{e^{\Delta r t} - \kappa_{PC}}{\kappa_{IC}}$ *and strictly increasing in t.*

The increase in the performance hurdle over time (indicated by the red line in Figure 3) results from the fact that the benefit of further deferral must outweigh the additional costs resulting from relative impatience. To interpret the prediction of an increasing performance hurdle within various economic applications, it is important to note that it applies to performance measured in likelihood ratios units, not necessarily in terms of profits. Concretely, consider the Brownian motion Example 3 with $\rho = 1$ and suppose that a fund manager's hidden action affects the drift rate of the instantaneous fund return with $a_H = \alpha + r_P > a_L = r_P$. That is, only under high effort does the fund manager deliver an expected excess return (alpha) α above the market rate r_P. Now, as, for given t, the monotone likelihood ratio property (MLRP) is satisfied, a higher average *excess* return, $\bar{x}_t - r_P$, maps directly into a higher likelihood ratio, $L_t = 1 - \exp\left(-[\bar{x}_t - r_P - \alpha/2]\frac{\alpha}{\sigma^2/t}\right)$, such that the agent gets paid if and only if $\bar{x}_t - r_P$ is sufficiently large.[26] Furthermore, while in this example the likelihood ratio associated with a given average excess return may both increase as well as decrease over time,[27] the hurdle that this performance measure has to clear in order to trigger payments to the agent becomes from Corollary 4 unambiguously more demanding.

So far, we have framed the problem as one of optimal compensation design maximizing the principal's payoff. However, the results derived above directly extend to a standard security

26. Of course, MLRP is not satisfied in various other settings such as, *e.g.*, Example 1.

27. Intuitively, since performance measurement based on more data is more precise, as captured by a higher signal-to-noise ratio $\alpha/(\sigma^2/t)$, observing the same average excess return at a later date results in a more favourable assessment of the agent's performance in likelihood ratio units whenever the outcome is sufficiently good, $\bar{x}_t - r_P > \alpha/2$, and a less favourable one if the outcome is sufficiently bad, $\bar{x}_t - r_P < \alpha/2$. Here the cut-off $\alpha/2$ corresponds to the average excess return that is equally likely under a_H and a_L.

design setting where an entrepreneur subject to moral hazard seeks financing from a competitive financial market (investors). The entrepreneur then chooses the dynamic structure of her and the investors' cash flows to maximize her expected payoff, where the optimal mix of securities and their maturity structure can be obtained from the characterization of the optimal contract given above by increasing the agent's outside option R up to the point where the principal breaks even.[28]

4. CONTRACTING WITH RISK-AVERSE AGENT

While in the previous application to financial contracting payment bounds are exogenously given by the resources available for distribution, they may also arise endogenously from the agent's risk aversion. As in Sannikov (2008), our analysis with a risk-averse agent assumes that the agent cannot save privately and considers the following class of agent utility functions:

Condition 3 *The agent's utility over flow consumption $u:[0,\infty)\rightarrow[0,\infty)$ is a strictly increasing, strictly concave C^2 function that satisfies $u(0)=0$ and $u'(c)\rightarrow 0$ as $c\rightarrow\infty$.*

Condition 3 implies that optimal contracts will never specify lump-sum transfers, so that it is without loss of generality to stipulate $db_t=c_t dt$, where c_t is the agent's date-t consumption flow. It is now also easy to see that the payment bound specification with $db_t\leq\bar{b}dt$ can already be viewed as an extreme case of risk aversion on the side of the agent (as in Plantin and Tirole (2018)): her marginal utility from flow consumption at any given point in time drops from one to zero when it exceeds some satiation point $\bar{b}$. Due to this connection, the characterization of the optimal contract under "standard" risk aversion can be thought of as a generalization of Proposition 1.

Proposition 2 *Suppose the agent's utility function satisfies Condition 3 and the set of contracts strictly satisfying the constraints is non-empty. Then, there exist non-negative shadow prices, κ_{IC} and κ_{PC}, such that $c_t>0$ if $\kappa_{IC}l_t+\kappa_{PC}>\frac{e^{\Delta rt}}{u'(0)}$ and $c_t=0$, otherwise. Optimal interior consumption solves*

$$\frac{e^{\Delta rt}}{u'(c_t)}=\kappa_{IC}l_t+\kappa_{PC}. \tag{8}$$

Compared to Proposition 1, the optimality condition in (8) now also reflects the inverse marginal utility term $\frac{1}{u'(c_t)}$ (in addition to impatience costs $e^{\Delta rt}$).[29] Here, $\frac{1}{u'(c_t)}$ can be interpreted as the marginal cost of transferring utility to the agent at date t. Since the cost of transferring utility is strictly convex under risk aversion, it is no longer optimal to pay at only a single (or two) payout dates following the realization of the maximum likelihood ratio history. Instead, starting from the maximal reward, which is specified for maximal performance, $\bar{L}_t$, at the "risk-neutral" payout time of $T^*=\arg\max_t e^{-\Delta rt}\left(\kappa_{IC}\bar{L}_t+\kappa_{PC}\right)$, it is optimal for the principal to smooth agent consumption across states and payout dates (see contour plot in Figure 4). Here, T^* corresponds to an optimal payout time in the risk-neutral setting for given shadow prices.[30]

If (PC) is slack, the concordant rent-extraction motive of the principal implies that smoothing only occurs across states with positive likelihood ratios. Moreover, the size of rewards is calibrated

28. It would be interesting to enrich this setting by allowing payment bounds at a particular point in time t to also depend on the endogenously chosen payouts (dividends, sale of equity stakes) at earlier points in time (cf. footnote 24).

29. As long as limited liability does not interfere, (8) implies that the inverse of the *discounted* marginal utility, $1/\left(e^{-\Delta rt}u'(c_t)\right)$, is a martingale (since L_t is a martingale). Except for the additional impatience term, our principal-agent model with persistence, thus, shares a key property with repeated-action models (see *e.g.* Rogerson, 1985).

30. Of course, the values of the shadow prices are themselves affected by risk aversion. Interestingly, when (PC) is slack, then T^* is given by the optimality condition in (2) regardless of the value of κ_{IC}.

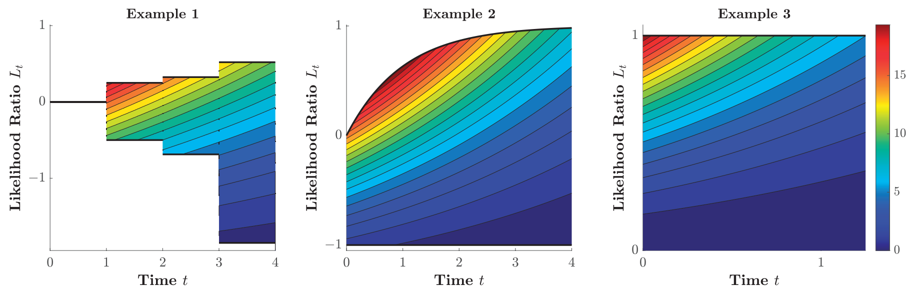

FIGURE 4

Contour plot of optimal contract under agent risk aversion

Notes: The graphs plots optimal compensation to the agent for the respective information environments of Figure 3. For Examples 1 and 2, the highest payment is made for some $t>0$, as sufficiently precise signal histories only become available after some time (see dark red contour line). In contrast, for Example 3, the highest possible reward is stipulated instantaneously for $l_{0+}\to 1$ as the likelihood ratio distribution has full support for all $t>0$.

such that the marginal cost of transferring utility is proportional to the discounted likelihood ratio. In contrast, if (PC) binds, the risk-sharing motive implies that the agent may even be rewarded for negative likelihood ratios, and the more so the higher the relevance of the participation constraint as reflected in κ_{PC}.[31]

We are now ready to state the main result of this section, characterizing the general economic features of the optimal contract arising from the agent's risk aversion and relative impatience as well as their interplay. For this, it is convenient to note that (8) defines for a given optimal contract, *i.e.*, fixed shadow values, a differentiable function $c(t,l)$, mapping $[0,\bar{T}]\times(-\infty,1]$ into the reals such that $c_t=\max\{0,c(t,l)\}$ for any realization of l feasible at a given t.[32]

Proposition 3 *Let $ARA(c)=-\frac{u''(c)}{u'(c)}$ denote the absolute risk aversion coefficient, then the sensitivity of payments at date-t with respect to performance, $\frac{\partial c(t,l)}{\partial l}$, and time, $\frac{\partial c(t,l)}{\partial t}$, satisfy within the payment region:*

$$\frac{\partial c(t,l)}{\partial l}=\frac{1}{ARA(c)}\frac{1}{l+\frac{\kappa_{PC}}{\kappa_{IC}}}>0, \tag{9}$$

$$\frac{\partial c(t,l)}{\partial t}=-\frac{\Delta r}{ARA(c)}<0. \tag{10}$$

The pay-performance sensitivity declines (increases) over time if the agent's preferences exhibit decreasing (increasing) absolute risk aversion, i.e., $\operatorname{sgn}\left(\frac{\partial}{\partial t}\frac{\partial c(t,l)}{\partial l}\right)=\operatorname{sgn}\left(ARA'(c)\right)$.

Proposition 3 highlights the different roles of consumption smoothing (risk aversion) and relative impatience for optimal contract design. *Risk aversion* implies that rewards are smoothed out across likelihood ratio states for a given t with higher rewards for higher performance, $\frac{\partial c}{\partial l}>0$, as in Holmstrom (1979). *Relative impatience* implies that the performance hurdle is increasing over

31. In fact, once κ_{PC} is sufficiently high (and if the likelihood ratio L_t is bounded below for $t=\bar{T}$), the limited liability constraints never bind and are, hence, irrelevant for optimal contract design.

32. For example, in Example 1, the likelihood ratio can take on only two values at $t=1$ (corresponding to success or failure). The contract then can still specify $c(t,l)$ for $l\in(-\infty,1]$, even though this variable is only relevant for the two levels of the likelihood ratio that can actually be reached.

time, *i.e.*, holding the likelihood ratio fixed, rewards are decreasing over time, $\frac{\partial c}{\partial t} < 0$.[33] These comparative statics can be directly inferred from Figure 4 which plots the optimal compensation schedule as a contour plot across likelihood ratio states and time for our three leading Examples 1 to 3. Finally, the *interplay* of impatience and risk aversion implies that the pay-performance sensitivity is decreasing over time (assuming decreasing absolute risk aversion). The intuition for this result is as follows. A fall in consumption due to the passage of time, $\frac{\partial c}{\partial t} < 0$, makes the agent effectively more risk averse (under decreasing absolute risk aversion), which, in turn, makes it optimal to reduce the performance-sensitivity of rewards.

We next investigate further properties of the optimal contract for two commonly used utility functions:

Corollary 5 *If the agent has CARA utility, $u(c) = 1 - e^{-\rho c}$ with $\rho > 0$, then*

$$c_t = \frac{1}{\rho}\left(\ln\rho + \ln(\kappa_{IC} l_t + \kappa_{PC}) - \Delta r t\right).$$

If the agent has CRRA utility, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$,[34] then

$$c_t = e^{-\frac{\Delta r}{\gamma}t}(\kappa_{IC} l_t + \kappa_{PC})^{\frac{1}{\gamma}}.$$

Moreover, if $\gamma = \frac{1}{2}$ and limited liability constraints never bind,[35] we obtain:

$$\kappa_{IC} = \frac{\Delta k}{2\int_0^{\bar{T}} e^{-(\Delta r + r_A)t}\mathbb{E}^H\left[L_t^2\right]dt}, \text{ and, } \kappa_{PC} = (\Delta r + r_A)\frac{R + k_H}{2}.$$

Thus, ceteris paribus, consumption decreases linearly over time for CARA and exponentially under CRRA preferences, and the rate of decay is inversely related to the respective risk-aversion parameter. Moreover, for the special case of square root utility (CRRA utility with $\gamma = \frac{1}{2}$), it is possible to solve explicitly for the shadow values κ_{IC} and κ_{PC}. In this setting, the relevant measure of informativeness is then given by the variance of the likelihood ratio distribution, $\mathbb{E}^H\left[L_t^2\right]$, which is again an increasing function of time (cf. Observation 1) due to the martingale property of L_t.

In sum, the optimal timing of pay with a risk averse agent is also shaped by the trade-off between the information benefit of deferral and the costs of transferring utility to the agent at different points in time. However, due to risk-sharing and consumption smoothing concerns, the value of the date-t information system to the principal may now depend on the entire likelihood ratio distribution, not only on its upper support as in the risk-neutral benchmark model (see Section 3). Despite this additional complexity, we still obtain clear-cut *qualitative* predictions (see Propositions 2 and 3), which can be interpreted as intuitive extensions of Holmstrom (1979) to a setting with dynamic information arrival. More concrete results regarding the "timing" of bonus payouts require a *quantification* of the information benefits of deferral which depends on the concrete form of the agent's consumption risk and smoothing preferences. We obtain such a quantification, *e.g.*, for commonly used CRRA utility specifications in Corollary 5 as well as within our risk-neutral benchmark framework.

33. One may think of the latter comparative static either as the local change in consumption at date t when no further information arrives, or as comparing two signal histories of different length but the same likelihood ratio.

34. Technically, if $\gamma > 1$, CRRA utility violates boundedness from below in Condition 3. However, an optimal contract still exists as long as the date-$\bar{T}$ likelihood ratio is bounded below.

35. Limited liability is slack for any (t, l_t) combination if the likelihood ratio distribution at $\bar{T}$ is bounded below and R is sufficiently high.

5. CONCLUSION

This article studies principal-agent settings in which the agent's action has persistent effects. The key contribution relative to the existing literature is that our approach allows us to obtain a complete, tractable, and intuitive characterization of optimal contracts in general information environments. We are able to accommodate good-news, bad-news, discrete, and continuous signal processes within a unified framework by exploiting the martingale property of the likelihood ratio process. The characterization of optimal contracts can be readily applied to various settings of economic interest.

We initially follow the security design literature by considering a risk-neutral, relatively impatient agent (see *e.g.* DeMarzo and Sannikov, 2006; Biais *et al.*, 2007; DeMarzo and Fishman, 2007). The absence of risk-sharing considerations implies that the maximal likelihood ratio is a sufficient statistic for the likelihood-ratio distribution at each date t, resulting in a precise and simple metric of the information gain over time. Regardless of how information arrives over time, optimal contracts feature a single payout date when the principal has a rent-extraction motive, and at most two payout dates with a binding participation constraint (Theorem 1). The simple form of the optimal contract allows for an intuitive understanding of the forces that determine optimal deferral, broadly consistent with empirical evidence (see *e.g.* Gopalan *et al.*, 2014).

The benefits of deferral, in terms of having access to a better information system, carry over to the case with a risk-averse agent. However, risk aversion implies that optimal contracts stipulate rewards for a larger selection of payout dates and states according to the following principle: first, rewards are increasing in performance measured in likelihood ratio units. Relative impatience then implies an increasing performance hurdle over time, *i.e.*, holding performance fixed, rewards are decreasing in time. Finally, the *interaction* of impatience and risk aversion implies that the pay-performance sensitivity is decreasing over time for any utility function exhibiting decreasing absolute risk aversion.

In future work, it would be interesting to extend our one-time action setup by allowing the agent to take on subsequent actions. One can envision that these follow-up actions do not only have persistent effects themselves, but are also state contingent: For example, the follow-up action "restructuring and downsizing" may only be required after initial failure, or the action "develop product" only available after a patent has been granted (see *e.g.* Green and Taylor (2016) for a related setting). In such settings, for each history h^t, one needs to account for all previously exerted action choices and keep track of a vector of likelihood ratios. Optimal contract design now needs to account for the correlation structure. Are outcomes that are indicative of taking the recommended initial action also indicative of the recommended follow-up action (as in Sannikov (2014)) or conflicting (as in Zhu (2018))? A rigorous analysis for general correlation structures should shed light on possible complementarities and substitutabilities in dynamic incentive provision, and, ultimately, whether the principal has an incentive to hire separate agents for different tasks. Relatedly, it may also be interesting to study situations in which the agent observes performance signals privately (as in Levitt and Snyder (1997)), and may manipulate the information observed by the principal.[36]

Finally, it would be interesting to endogenize the information process by giving the *principal* a more active role. In our model information arrives exogenously and costs arise only indirectly when using (later) information for the purpose of incentive compensation. There are, of course, settings in which information has to be generated by the principal, and acquiring additional

36. To make the problem interesting, manipulation could entail costly destruction of output (as in Innes (1990)), inefficient diversion (as in DeMarzo and Fishman (2007)), or may only be observable with some delay (as in Varas (2018)).

information may be costly (see Plantin and Tirole, 2018). In such settings, next to designing the optimal compensation contingent on the available information, the principal has to choose which information to acquire.

APPENDIX

A. PROOFS

Lemma A.1 *It is without loss of generality to restrict attention to compensation processes adapted to* $(\mathcal{F}_t^L)_{0\le t\le\bar{T}}$, *the filtration generated by* L_t *as defined in (1).*

Proof of Lemma A.1. The result follows directly from an application of the Halmos–Savage theorem showing that for each date t, L_t is a sufficient statistic for $\{x_s\}_{0\le s\le t}$ with respect to a. Then, as the expected cost of transferring utility is time separable, the result follows from a straightforward dynamic extension of the arguments in Holmstrom (1979) (Proposition 3).

For completeness we spell out some more details. Employing a standard change of measure we can rewrite (IC) as

$$\mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_A t}\left(1-\frac{d\mathbb{P}^L}{d\mathbb{P}^H}\right)dv_t\right]\ge\Delta k,$$

which, by iterated expectations, is equivalent to (IC*):

$$\mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_A t}\left(1-\mathbb{E}^H\left[\frac{d\mathbb{P}^L}{d\mathbb{P}^H}\middle|\mathcal{F}_t^X\right]\right)dv_t\right]=\mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_A t}L_t dv_t\right]\ge\Delta k. \tag{A.1}$$

Now consider the compensation design Problem 1 with (IC) replaced by (A.1), first, for the case of a risk neutral agent so that $v_t=b_t$. By inspection, there exists for each $\mathcal{F}_t^X$-adapted contract $\{b_t\}$ a corresponding $\mathcal{F}_t^L$-adapted contract $\left\{\tilde{b}_t\right\}$ with $b_{0^-}=\tilde{b}_{0^-}=0$ and $d\tilde{b}_t:=\mathbb{E}^H\left[db_t|\mathcal{F}_t^L\right]$ leaving the constraints unaffected and resulting in the same wage costs. To establish the result for the case of a risk-averse agent (where we impose Condition 3), it is convenient to rewrite Problem 1 with the contingent discounted rate of agent utility $\nu_t:=e^{-r_A t}u(c_t)$ as the choice variable:

$$W:=\min_{\nu_t}\mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_P t}u^{-1}(e^{r_A t}\nu_t)dt\right]\quad \text{s.t.} \tag{W}$$

$$\mathbb{E}^H\left[\int_0^{\bar{T}}\nu_t dt\right]-k_H\ge R, \tag{PC}$$

$$\mathbb{E}^H\left[\int_0^{\bar{T}}L_t\nu_t dt\right]\ge\Delta k, \tag{IC}$$

$$\nu_t\ge 0\quad\forall t. \tag{LL}$$

Then, there exists for each $\mathcal{F}_t^X$-adapted contract $\{\nu_t\}$ a corresponding $\mathcal{F}_t^L$-adapted contract $\{\tilde{\nu}_t\}$ with $\tilde{\nu}_t:=\mathbb{E}^H\left[\nu_t|\mathcal{F}_t^L\right]$ leaving the constraints unaffected but now even resulting in lower wage costs due to improved risk-sharing. **Q.E.D.**

Lemma A.2 *For any weakly increasing function* $I:[0,\bar{T}]\to[0,1]$, *one can construct a signal process* X_t *such that the associated informativeness function satisfies* $\bar{L}_t=I(t)$.

Proof of Lemma A.2. Note, first, that as $I:[0,\bar{T}]\to[0,1]$ is monotonically increasing, it is differentiable almost everywhere and can have at most countably many discontinuities. We now provide, for given $I(t)$, a construction of an information system with $\bar{L}_t=I(t)$ based on Examples 1 and 2. To do so, denote the set of times for which I has a jump discontinuity by T^d. Then consider the following information environment: the agent's action determines the arrival rate of a bad event, $\lambda_t(a)$, with $\lambda_t(a_H)\le\lambda_t(a_L)$ for all $t\in[0,\bar{T}]$. Further, at each $t\in T^d$ there is a binary signal $x_t\in\{s,f\}$ that is drawn independently over time with the probability of success "s" at a particular $t\in T^d$ depending on the agent's action with $p_t(a_H) > p_t(a_L)$. Then, it is easy to show that for each t, the history corresponding to the maximal likelihood ratio is the one with the maximal possible number of successes and no bad event, *i.e.*,

$$\bar{L}_t = 1 - \frac{\exp\left\{-\int_0^t \lambda_s(a_L)ds\right\} \prod_{s\in\{i\in T^d : i\le t\}} p_s(a_L)}{\exp\left\{-\int_0^t \lambda_s(a_H)ds\right\} \prod_{s\in\{i\in T^d : i\le t\}} p_s(a_H)}.$$

It is now sufficient to note that $\lambda_t(a_H) \le \lambda_t(a_L)$ and $p_t(a_H) > p_t(a_L)$ can be chosen such that $\bar{L}_t = I(t)$ for each t. In particular, for $t \in T^d$ one can choose $\lambda_t(a_H) = \lambda_t(a_L)$ and $p_t(a_H)/p_t(a_L) = 1 + dI(t)/(1-I(t))$. At points of differentiability we choose $(\lambda_t(a_L) - \lambda_t(a_H)) = I'(t)/(1-I(t))$. **Q.E.D.**

Proof of Theorem 1. For convenience we restate the compensation design problem with bilateral risk neutrality using (A.1)

$$W := \min_{b_t} \mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_P t} db_t\right] \quad \text{s.t.} \tag{W}$$

$$\mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_A t} db_t\right] - k_H \ge R, \tag{PC}$$

$$\mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_A t} L_t db_t\right] \ge \Delta k, \tag{IC}$$

$$db_t \ge 0 \quad \forall t. \tag{LL}$$

We now first show that, given the conditions of the theorem, (IC) always binds.

Lemma A.3 *If the agent is risk-neutral, the shadow value on (IC), κ_{IC}, is zero if and only if $\bar{L}_0 \ge \frac{\Delta k}{R+k_H}$.*

Proof of Lemma A.3 From (PC) and (LL) together with differential discounting we have that $W \ge R + k_H$. Hence, we need to show that $W = R + k_H$ if and only if $\bar{L}_0 \ge \frac{\Delta k}{R+k_H}$. To show sufficiency, consider the contract that delivers a compensation value of $\mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_A t} db_t\right] = (R + k_H)$ to the agent and specifies a single contingent payment at $t = 0$ following the realization of $L_0 = \bar{L}_0$. Given Condition 1 such a contract exists with contingent bonus $(R+k_H)/\mathbb{P}^H\left(L_0 = \bar{L}_0\right)$ and it satisfies (PC) and (LL). Further, (IC) is satisfied if $\bar{L}_0 \ge \frac{\Delta k}{R+k_H}$. To show the necessary part, observe that any contract with $W = R + k_H$ cannot feature any deferred pay due to differential discounting. Note further, that the contract that provides strongest incentives with date-0 payments only is, from (IC), the one that makes the entire expected pay $\mathbb{E}^H[db_0]$ contingent on $\bar{L}_0$. Now suppose that $\bar{L}_0 < \frac{\Delta k}{R+k_H}$, then the contract requires $\mathbb{E}^H[db_0] > R + k_H$ in order to satisfy (IC), implying $W > R + k_H$. **Q.E.D.**

The following Lemma characterizes the contingency of pay in optimal contracts:

Lemma A.4 *An optimal contract stipulates rewards at date-t only if a signal history with maximum likelihood ratio realizes, i.e., $\forall t$, $db_t = 0$ for all $l < \bar{L}_t$ a.s..*

Proof of Lemma A.4. When $\bar{L}_0 < \frac{\Delta k}{R+k_H}$ we have from Lemma A.3 that $\kappa_{IC} > 0$, implying $W > R + k_H$. The proof then is by contradiction. So assume that under the optimal contract there exists some t with $db_t > 0$ for some $l < \bar{L}_t$. Then if all such payments occur with probability zero they can be dropped leaving all constraints and compensation costs unaffected. Else, there exists another feasible contract with $db_t = 0$ for all $l < \bar{L}_t$ that yields strictly lower compensation costs, contradicting optimality of the candidate contract. To see this, assume, first, that (PC) is slack and make, for all t, all payments contingent on the realization of $l = \bar{L}_t$ holding $\mathbb{E}^H\left[e^{-r_A t} db_t\right]$ and, thus, total compensation costs constant.[37] However, under this alternative contract (IC) is slack which allows to reduce $db_t > 0$ at some t for which $\mathbb{E}^H[db_t] > 0$, thus, reducing compensation costs. Second, assume that (PC) binds, which, together with $\kappa_{IC} > 0$ implies that not all payments can occur at $t = 0$. Then, using the same variation of the original contract constructed above we again arrive at a solution with slack (IC), which now allows to reduce $db_t > 0$ at some $t > 0$ for which $\mathbb{E}^H[db_t] > 0$ to $(1-x)db_t$ with $x \in (0,1)$

37. The associated compensation process b exists from Condition 1.

and add a lump-sum payment at $t=0$ of $\mathbb{E}^H\left[e^{-r_A t} x db_t\right]$ to still satisfy (PC). Again W is lower now due to differential discounting. For the case where $\bar{L}_0=\frac{\Delta k}{R+k_H}$ the result follows directly from the construction in Lemma A.3. **Q.E.D.**

It remains to characterize the optimal timing of pay. To do so, it is convenient to consider the following change of variables: For any admissible process b_t define

$$B:=\mathbb{E}^H\left[\int_0^{\bar{T}} e^{-r_A t} db_t\right],$$

$$w_s:=\mathbb{E}^H\left[\int_0^{s} e^{-r_A t} db_t\right]/B.$$

Then using Lemma A.4 we can rewrite Problem 1 in terms of (B,w) as follows:

$$W:=\min_{B,w_t} B\int_0^{\bar{T}} e^{\Delta r t} dw_t \quad \text{s.t.} \tag{W}$$

$$B\geq R+k_H, \tag{PC}$$

$$B\int_0^{\bar{T}} \bar{L}_t dw_t \geq \Delta k, \tag{IC}$$

$$dw_t\geq 0 \quad \forall t, \tag{LL}$$

where $w_{\bar{T}}=\int_0^{\bar{T}} dw_t=1$. We will now characterize the optimal timing and size of pay in terms of (B,w). We can then recover b via the transformation $\mathbb{E}^H[db_t]=e^{r_A t}Bdw_t$ for each $t\in\left[0,\bar{T}\right]$ together with Lemma A.4. Such a process b exists whenever Condition 1 holds and solves Problem 1.

Consider, first, the relaxed problem with slack (PC). Then, substituting out B from the objective function using (IC), the compensation design problem reduces to

$$W=\Delta k \min_{w_t} \frac{\int_0^{\bar{T}} e^{\Delta r t} dw_t}{\int_0^{\bar{T}} \bar{L}_t dw_t}. \tag{A.2}$$

Now, consider the curve $\left(\bar{L}_t, e^{\Delta r t}\right)$ parameterized by $t\in\left[0,\bar{T}\right]$, and its convex hull tracing out the set of $\left(\int_0^{\bar{T}} \bar{L}_t dw_t, \int_0^{\bar{T}} e^{\Delta r t} dw_t\right)$ achievable with any admissible weighting $(w_t)_{t=0}^{\bar{T}}$. Using the definition of the lower hull C (see main text and, for a graphical illustration, Figures 1 and 2) and defining $I:=\int_0^{\bar{T}} \bar{L}_t dw_t$, (A.2) can be rewritten as

$$W=\Delta k \min_{I\in\left[\bar{L}_0,\bar{L}_{\bar{T}}\right]} \frac{C(I)}{I}. \tag{A.3}$$

For any given I, $\frac{C(I)}{I}$, graphically corresponds to the slope of a ray through the origin and the respective point on C. The objective is, thus, to minimize the slope (see slope of green line in right panels of Figure 2). If there exists a unique solution, then the minimizer defines an extreme point on C which, by definition of an extreme point (cf., Krein and Milman, 1940), can be achieved with a single payout at date $\hat{T}$, $\left(\bar{L}_{\hat{T}}, e^{\Delta r\hat{T}}\right)$ (see *e.g.* the left panels of Figure 2). In the case of multiple global minimizers of $\frac{C(I)}{I}$, both the smallest and the largest minimizer must define an extreme point. Disregarding (PC), any minimizer (or a combination thereof) solves the principal's relaxed problem. Now, since $\hat{T}$ features the lowest informativeness among the set of minimizers, the associated compensation value to the agent $\Delta k/\bar{L}_{\hat{T}}$ is highest. Hence, a necessary and sufficient condition for PC to be slack is that $R\leq\bar{R}$.

The proof of statement 2 in the Theorem is in the main text. It just remains to show that $\hat{T}$ is an upper bound on payment dates with binding (PC). To see this, note first that (PC) binds, $R>\bar{R}$, if and only if $\frac{\Delta k}{R+k_H}<\bar{L}_{\hat{T}}$, where $\frac{\Delta k}{R+k_H}$ is the weighted average informativeness required by (IC) for binding (PC). If the optimal contract with binding (PC) features a single payment at T_1, the result follows directly from $\bar{L}_{T_1}=\frac{\Delta k}{R+k_H}<\bar{L}_{\hat{T}}$ together with the fact that $\bar{L}_t$ is increasing (see *e.g.* the upper panels of Figure 2). In the remaining case in which two payout dates $T_S<T_L$ are optimal, it follows from the fact that T_S, T_L and $\hat{T}$ must all define extreme points (see *e.g.* Figure 1 and the lower panels of Figure 2). If T_L were greater than $\hat{T}$, *i.e.*, $\bar{L}_{T_L}>\bar{L}_{\hat{T}}>\frac{\Delta k}{R+k_H}=w_S\bar{L}_{T_S}+(1-w_S)\bar{L}_{T_L}$, then $\left(\bar{L}_{\hat{T}}, e^{\Delta r\hat{T}}\right)$ could not be an extreme point. Contradiction. **Q.E.D.**

Proof of Lemma 1. It remains to derive expression (6). To see this, suppose to the contrary that for any given T_S there exists a $T'>T_S$ such that $\left(\bar{L}_{T'}-\bar{L}_{T_S}\right)/\left(e^{\Delta r T'}-e^{\Delta r T_S}\right)>\left(\bar{L}_{T_L}-\bar{L}_{T_S}\right)/\left(e^{\Delta r T_L}-e^{\Delta r T_S}\right)$. Then, if $\bar{L}_{T'}>\frac{\Delta k}{R+k_H}$, there exists a

$w_S' \in (0,1)$ such that $w_S'\bar{L}_{T_S} + (1-w_S')\bar{L}_{T'} = \frac{\Delta k}{R+k_H}$ and $w_S' e^{\Delta r T_S} + (1-w_S')e^{\Delta r T'} < w_S e^{\Delta r T_S} + (1-w_S)e^{\Delta r T_L}$ contradicting optimality of payouts at T_S and T_L. Similarly, if $\bar{L}_{T'} \leq \frac{\Delta k}{R+k_H}$ there exists a $w_S' \in [0,1)$, such that $w_S'\bar{L}_{T'} + (1-w_S')\bar{L}_{T_L} = \frac{\Delta k}{R+k_H}$ and $w_S' e^{\Delta r T'} + (1-w_S')e^{\Delta r T_L} < w_S e^{\Delta r T_S} + (1-w_S)e^{\Delta r T_L}$ again contradicting optimality.

For completeness, we further state the explicit characterization of db_{T_S} and db_{T_L}:

$$\begin{aligned} db_{T_S} &= (R+k_H)w_S \frac{e^{r_A T_S}}{\mathbb{P}^H\left(L_{T_S} = \bar{L}_{T_S}\right)} && \text{if} \quad L_{T_S} = \bar{L}_{T_S}, \\ db_{T_L} &= (R+k_H)(1-w_S)\frac{e^{r_A T_L}}{\mathbb{P}^H\left(L_{T_L} = \bar{L}_{T_L}\right)} && \text{if} \quad L_{T_L} = \bar{L}_{T_L}, \\ db_t &= 0 && \text{else.} \end{aligned}$$

Q.E.D.

Proof of Corollary 1. To show statement 1, we restrict attention to twice differentiable functions $\bar{L}_t$. Strict convexity of C (over the domain for which (PC) binds) follows directly from standard arguments, see *e.g.* Gollier (2004) (Chapter 2, Proposition 2). Statement 1 then follows directly from Lemma 1 together with Theorem 1.

To show statement 2, it is sufficient to note that by definition $\bar{L}_t$ takes on only countably many values such that $\bar{L}_t = \frac{\Delta k}{R+k_H}$ has no solution t for almost all $\frac{\Delta k}{R+k_H} \in \left(0, \bar{L}_{\hat{T}}\right)$. Since, $R > \bar{R}$ implies that (PC) binds, the result then follows from Lemma 1. **Q.E.D.**

Proof of Corollary 2. From Theorem 1, these comparative statics hold trivially if $\bar{L}_{\hat{T}} < \frac{\Delta k}{R+k_H}$ so that (PC) is slack. In this regime 1, the duration $\hat{T}$ does not depend on R and Δk. When $\bar{L}_{\hat{T}} \geq \frac{\Delta k}{R+k_H} > \bar{L}_0$, (PC) and (IC) bind (regime 2), and the result follows directly from $\int_0^{\bar{T}} \bar{L}_t dw_t = \frac{\Delta k}{R+k_H}$, which is decreasing in R and increasing in Δk, together with Lemma 1. Finally, when $\bar{L}_{\hat{T}} \geq \bar{L}_0 \geq \frac{\Delta k}{R+k_H}$, Lemma A.3 shows that (IC) is slack (regime 3) and the duration is equal to zero independently of R and Δk. Now note that, as R increases or Δk decreases, we either stay within a given regime or move from regime 1 to regime 2 to regime 3 and the result follows. **Q.E.D.**

Proof of Corollary 3. Given differentiability of $\bar{L}_t$ and the assumption of an interior payout date, the optimal payment date with slack (PC) is characterized by (3), *i.e.*, $g'\left(\hat{T}\right) = \frac{\Delta r}{\phi}$, and the first statement follows. Next, note that concavity of $\bar{L}_t$ implies strict convexity of C such that with binding (PC) Lemma 1 implies a single payout at date T_1 satisfying $\bar{L}_{T_1} = \frac{\Delta k}{R+k_H}$. The second statement then follows directly as for each t, informativeness $\bar{L}_t$ is strictly increasing in ι and ϕ. **Q.E.D.**

Proof of Proposition 1. First, note that Problem 1 with a risk-neutral agent and payment bounds (PB) as given in Condition 2 is a linear programming problem with finite value. Existence of an optimal contract then follows directly from the assumption that the set of contracts satisfying the constraint set is non-empty. Let $m_t(l)$ denote the flow payment at time t following realization of $l \in \mathbf{L}_t$ where $\mathbf{L}_t$ denotes the support of the date-t likelihood ratio distribution. Then, using Lemma A.1, the Lagrangian can be written as follows

$$\begin{aligned} \mathcal{L} = \int_0^{\bar{T}} \int_{\mathbf{L}_t} e^{-r_A t}\left[e^{\Delta r t} - \kappa_{IC} l - \kappa_{PC} - \left(\kappa_{LL}^t(l) - \kappa_{PB}^t(l)\right)\right] m_t(l) dF_t(l) dt \\ + \kappa_{IC}\Delta k + \kappa_{PC}(k_H + R), \end{aligned}$$

where $F_t(l) = \mathbb{P}^H(L_t \leq l)$ denotes the ex ante distribution over date-t likelihood ratios if the agent chooses action a_H and $\kappa_{IC}, \kappa_{PC}, \kappa_{LL}^t(l), \kappa_{PB}^t(l) \geq 0$ are the Lagrange multipliers on the respective constraints. Since $\mathcal{L}$ is linear in $m_t(l)$ for all (t,l), we almost everywhere either have $m_t(l) = 0$ (with $\kappa_{LL}^t(l) \geq 0$ and $\kappa_{PB}^t(l) = 0$) or $m_t(l) = \bar{b}$ (with $\kappa_{LL}^t(l) = 0$ and $\kappa_{PB}^t(l) \geq 0$). Now from $\partial\mathcal{L}/\partial m_t(l) = 0$ we get

$$e^{\Delta r t} - \kappa_{IC} l - \kappa_{PC} = \kappa_{LL}^t(l) - \kappa_{PB}^t(l),$$

such that $m_t(l) = \bar{b}$ if and only if $e^{\Delta r t} - \kappa_{IC} l - \kappa_{PC} \leq 0$. The case with slack (PC) then corresponds to $\kappa_{PC} = 0$ and applies whenever $R \geq \bar{R}(\bar{b})$, where the cut-off value is given by $\bar{R}(\bar{b}) := \int_0^{\bar{T}} \int_{\widetilde{\mathbf{L}}_t(\bar{b})} e^{-r_A t}\bar{b} dF_t(l) dt - k_H$ with $\widetilde{\mathbf{L}}_t(\bar{b}) := \left\{ l \in \mathbf{L}_t : e^{-\Delta r t} l \geq \frac{1}{\widetilde{\kappa}_{IC}(\bar{b})} \right\}$ for each t and $\widetilde{\kappa}_{IC}(\bar{b})$ denoting the Lagrange multiplier on (IC) in the solution to the compensation design problem with (PB) but absent (PC). **Q.E.D.**

Proof of Proposition 2. Consider the compensation design problem with the same change of variables as in Lemma A.1, *i.e.*, $v_t = e^{-r_A t} u(c_t)$. This ensures that we have a convex programming problem. The corresponding Lagrangian can then be written as

$$\mathcal{L}=\int_0^{\bar{T}}\int_{\mathbf{L}_t}\Big[e^{-r_P t}u^{-1}(e^{r_A t}\nu_t(l))-\kappa_{IC}l\nu_t(l)-\kappa_{PC}\nu_t(l)-\kappa_{LL}^t(l)\nu_t(l)\Big]dF_t(l)dt$$
$$+\kappa_{IC}\Delta k+\kappa_{PC}(k_H+R),$$

where $\kappa_{IC},\kappa_{PC},\kappa_{LL}^t(l)\geq 0$ denote the Lagrange multipliers on the respective constraints. Optimizing point-wise with respect to $\nu_t(l)$, we obtain that

$$\frac{e^{\Delta rt}}{u'(c_t(l))}-\kappa_{IC}l-\kappa_{PC}=\kappa_{LL}^t(l), \tag{A.4}$$

for almost every (t,l). First note that for given (t,l) the left-hand side of (A.4) is bounded above by $\frac{e^{\Delta rt}}{u'(0)}-\kappa_{IC}l-\kappa_{PC}$ due to strict concavity of u and $u'(0)>0$ (see Condition 3). Hence, the limited liability constraint binds, $\kappa_{LL}^t(l)>0$, if and only if $\frac{e^{\Delta rt}}{u'(0)}-\kappa_{IC}l-\kappa_{PC}\geq 0$. Otherwise, (A.4) has a solution $c_t(l)>0$, where $c_t(l)$ is bounded above since $u'(c)\to 0$ as $c\to\infty$ (see Condition 3). Existence of an optimal contract follows from the assumption that the set of contracts strictly satisfying the constraints is non-empty. **Q.E.D.**

Proof of Proposition 3. The expressions in (9) and (10) follow from implicit differentiation of (8) holding fixed the Lagrange multipliers at their optimal level. Then the cross derivative is given by

$$\frac{\partial^2 c}{\partial t\partial l}=-\frac{\partial c}{\partial t}\frac{\partial c}{\partial l}\frac{ARA'(c)}{ARA(c)},$$

which has same sign as $ARA'(c)$. **Q.E.D.**

Acknowledgments. We thank five anonymous referees, the editor Veronica Guerrieri, as well as Michalis Anthropelos, Alvin Chen, Peter DeMarzo, Darrell Duffie, Willie Fuchs, Xavier Gabaix, Nicolae Gârleanu, Sebastian Gryglewicz, Bengt Hölmstrom, Gustavo Manso, Vincent Maurin, John Morgan, Sebastian Pfeil, Sven Rady, Jan Starmans, Per Stromberg, Steven Tadelis, Jean Tirole, and Johan Walden for valuable comments on earlier drafts. In addition, we thank seminar participants at Berkeley (Micro Theory), MIT (Organizational Econ), Harvard (Organizational Econ), Stanford GSB, University of Rochester (Theory), Harvard Business School, University of Bonn (Micro Theory), University of Mannheim (Applied Micro), University of British Columbia, Stockholm School of Economics, University of Calgary, INSEAD, Econometric Society Summer meeting Lisbon, AFA Philadelphia, FIRS Barcelona, EFA Oslo, CEPR Gerzensee. Inderst gratefully acknowledges financial support from the European Research Council (Advanced Grant 22992 "Regulating Retail Finance") and from the German Research Foundation (DFG, Leibniz Preis). Opp gratefully acknowledges financial support from the Marianne & Marcus Wallenberg Foundation.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

BARANCHUK, N., KIESCHNICK, R. and MOUSSAWI, R. (2014), "Motivating Innovation in Newly Public Firms", *Journal of Financial Economics*, **111**, 578–588.

BERGEMANN, D. and MORRIS, S. (2016), "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium", *American Economic Review*, **106**, 586–591.

BIAIS, B., MARIOTTI, T., PLANTIN, G. and ROCHET, J. (2007), "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications", *Review of Economic Studies*, **74**, 345–390.

DEMARZO, P. and DUFFIE, D. (1999), "A Liquidity-based Model of Security Design", *Econometrica*, **67**, 65–99.

DEMARZO, P. M. and FISHMAN, M. J. (2007), "Optimal Long-term Financial Contracting", *Review of Financial Studies*, **20**, 2079–2128.

DEMARZO, P. M. and SANNIKOV, Y. (2006), "Optimal Security Design and Dynamic Capital Structure in a Continuous-time Agency Model", *Journal of Finance*, **61**, 2681–2724.

EDMANS, A., GABAIX, X., SADZIK, T. and SANNIKOV, Y. (2012), "Dynamic CEO Compensation", *Journal of Finance*, **67**, 1603–1647.

FUDENBERG, D. and TIROLE, J. (1990), "Moral Hazard and Renegotiation in Agency Contracts", *Econometrica*, **58**, 1279–1319.

GEORGIADIS, G. and SZENTES, B. (2020), "Optimal Monitoring Design", *Econometrica*, forthcoming.

GJESDAL, F. (1982), "Information and Incentives: The Agency Information Problem", *Review of Economic Studies*, **49**, 373–390.

GOLLIER, C. (2004), *The Economics of Risk and Time*, Vol. 1 (Cambridge, MA: The MIT Press).
GOPALAN, R., MILBOURN, T., SONG, F. and THAKOR, A. V. (2014), "Duration of Executive Compensation", *Journal of Finance*, **69**, 2777–2817.
GREEN, B. and TAYLOR, C. R. (2016), "Breakthroughs, Deadlines, and Self-reported Progress: Contracting for Multistage Projects", *American Economic Review*, **106**, 3660–3699.
GROSSMAN, S. J. and HART, O. D. (1983), "An Analysis of the Principal-agent Problem", *Econometrica*, **51**, 7–45.
HARTMAN-GLASER, B., PISKORSKI, T. and TCHISTYI, A. (2012), "Optimal Securitization with Moral Hazard", *Journal of Financial Economics*, **104**, 186 – 202.
HÉBERT, B. (2018), "Moral Hazard and the Optimality of Debt", *Review of Economic Studies*, **85**, 2214–2252.
HERMALIN, B. E. and KATZ, M. L. (1991), "Moral Hazard and Verifiability: The Effects of Renegotiation in Agency", *Econometrica*, **59**, 1735–1753.
HOFFMANN, F., INDERST, R. and OPP, M. M. (2020), "The Economics of Deferral and Clawback Requirements" (Unpublished Working Paper, ESE Rotterdam, University of Frankfurt, Stockholm School of Economics).
HOLMSTROM, B. (1979), "Moral Hazard and Observability", *Bell Journal of Economics*, **10**, 74–91.
HOLMSTROM, B. and TIROLE, J. (1997), "Financial Intermediation, Loanable Funds, and the Real Sector", *Quarterly Journal of Economics*, **112**, pp. 663–691.
HOPENHAYN, H. and JARQUE, A. (2010), "Unobservable Persistent Productivity and Long Term Contracts", *Review of Economic Dynamics*, **13**, 333–349.
INNES, R. D. (1990), "Limited Liability and Incentive Contracting with Ex-ante Action Choices", *Journal of Economic Theory*, **52**, 45–67.
JACOD, J. and SHIRYAEV, A. N. (2003), *Limit Theorems for Stochastic Processes*, Vol. 2 (Berlin, Heidelberg, New York: Springer).
JARQUE, A. (2010), "Repeated Moral Hazard with Effort Persistence", *Journal of Economic Theory*, **145**, 2412–2423.
JEWITT, I., KADAN, O. and SWINKELS, J. M. (2008), "Moral Hazard with Bounded Payments", *Journal of Economic Theory*, **143**, 59–82.
KAMENICA, E. and GENTZKOW, M. (2011), "Bayesian Persuasion", *American Economic Review*, **101**, 2590–2615.
KIM, S. K. (1995), "Efficiency of an Information System in an Agency Model", *Econometrica*, **63**, 89–102.
KREIN, M. and MILMAN, D. (1940), "On Extreme Points of Regular Convex Sets", *Studia Mathematica 9*, 133–138.
LEVITT, S. and SNYDER, C. (1997), "Is No News Bad News? Information Transmission and the Role of Early Warning in the Principal-agent Model", *RAND Journal of Economics*, **28**, 641–661.
LI, A. and YANG, M. (2020), "Optimal Incentive Contract with Endogenous Monitoring Technology", *Theoretical Economics*, **15**, 1135–1173.
MALAMUD, S., RUI, H. and WHINSTON, A. (2013), "Optimal Incentives and Securitization of Defaultable Assets", *Journal of Financial Economics*, **107**, 111 – 135.
MANSO, G. (2011), "Motivating Innovation", *Journal of Finance*, **66**, 1823–1860.
OPP, M. M. and ZHU, J. Y. (2015), "Impatience versus Incentives", *Econometrica*, **83**, 1601–1617.
PLANTIN, G. and TIROLE, J. (2018), "Marking to Market versus Taking to Market", *American Economic Review*, **108**, 2246–2276.
ROGERSON, W. P. (1985), "Repeated Moral Hazard", *Econometrica*, **53**, 69–76.
SANNIKOV, Y. (2008), "A Continuous-time Version of the Principal: Agent Problem", *Review of Economic Studies*, **75**, 957–984.
SANNIKOV, Y. (2014), "Moral Hazard and Long-run Incentives" (Unpublished Working Paper, Princeton University).
TIROLE, J. (2006), *Theory of Corporate Finance* (Princeton: Princeton University Press).
VARAS, F. (2018), "Managerial Short-Termism, Turnover Policy, and the Dynamics of Incentives", *Review of Financial Studies*, **31**, 3409–3451.
ZHU, J. Y. (2018), "Myopic Agency", *Review of Economic Studies*, **85**, 1352–1388.