

DISCUSSION PAPER SERIES

DP13643

ONLY TIME WILL TELL: A THEORY OF DEFERRED COMPENSATION

Roman Inderst and Marcus Opp

INDUSTRIAL ORGANIZATION

ONLY TIME WILL TELL: A THEORY OF DEFERRED COMPENSATION

Roman Inderst and Marcus Opp

Discussion Paper DP13643

Published 02 April 2019

Submitted 27 March 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Roman Inderst and Marcus Opp

ONLY TIME WILL TELL: A THEORY OF DEFERRED COMPENSATION

Abstract

This paper provides a complete characterization of optimal contracts in principal-agent settings where the agent's action has persistent effects. We model general information environments via the stochastic process of the likelihood-ratio. The martingale property of this performance metric captures the information benefit of deferral. Costs of deferral may result from both the agent's relative impatience as well as her consumption smoothing needs. If the relatively impatient agent is risk neutral, optimal contracts take a simple form in that they only reward maximal performance for at most two payout dates. If the agent is additionally risk-averse, optimal contracts stipulate rewards for a larger selection of dates and performance states: The performance hurdle to obtain the same level of compensation is increasing over time whereas the pay-performance sensitivity is declining. We derive testable implications for the optimal duration of (executive) compensation and the maturity structure of claims in financial contracting settings.

JEL Classification: D86

Keywords: Compensation design, duration of pay, moral hazard, Persistence, Principal-Agent Models, Informativeness principle

Roman Inderst - inderst@finance.uni-frankfurt.de
Goethe-Universität Frankfurt and CEPR

Marcus Opp - marcus.opp@hhs.se
Stockholm School of Economics

Only time will tell: A theory of deferred compensation*

Florian Hoffmann[†] Roman Inderst[‡] Marcus Opp[§]

February 4, 2019

Abstract

This paper provides a complete characterization of optimal contracts in principal-agent settings where the agent’s action has persistent effects. We model general information environments via the stochastic process of the likelihood-ratio. The martingale property of this performance metric captures the information benefit of deferral. Costs of deferral may result from both the agent’s relative impatience as well as her consumption smoothing needs. If the relatively impatient agent is risk neutral, optimal contracts take a simple form in that they only reward maximal performance for at most two payout dates. If the agent is additionally risk-averse, optimal contracts stipulate rewards for a larger selection of dates and performance states: The performance hurdle to obtain the same level of compensation is increasing over time whereas the pay-performance sensitivity is declining. We derive testable implications for the optimal duration of (executive) compensation and the maturity structure of claims in financial contracting settings.

Keywords: Compensation design, duration of pay, moral hazard, persistence, principal-agent models, informativeness principle.

JEL Classification: D86 (Economics of Contract: Theory).

*We thank Patrick Bolton, Peter DeMarzo, Willie Fuchs, Xavier Gabaix, Bengt Holmstrom, Gustavo Manso, Per Stromberg, Jean Tirole, and Jeffrey Zwiebel 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).

[†]Erasmus University Rotterdam. E-mail: hoffmann@ese.eur.nl.

[‡]Johann Wolfgang Goethe University Frankfurt. E-mail: nderst@finance.uni-frankfurt.de.

[§]Stockholm School of Economics. E-mail: marcus.opp@hhs.se.

1 Introduction

In many real-life principal-agent relationships, actions by agents have long-lasting – *not immediately observable* – effects on outcomes. 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 may not be assessed until well into the future. The list is certainly not exclusive and, yet, it suggests that the delay of observability is an important, if not defining feature of many moral hazard environments.

The goal of our paper is to address a basic research question. How does one 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 maximally informative outcomes, 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.

To focus on the optimal *intertemporal* provision of incentives in general information environments, we consider an otherwise parsimonious principal-agent setting with a one-time action (such as 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 (bonus) 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 compatibility and participation constraint. 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 formal 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, that while it may be natural that additional signals, such as, e.g., stock prices, become noisier over time, the informativeness contained in the history of signals, e.g., the entire *path of stock prices*, must weakly increase.¹

For ease of exposition, we initially assume bilateral risk-neutrality with agent limited liability (cf., Innes (1990) and Kim (1997)). This allows us to obtain clear-cut timing implications before analyzing the confounding factors of intertemporal smoothing motives. The key simplification of the compensation design problem in this setting results from the fact that the information relevant for the optimal timing of pay is fully summarized by the *maximal likelihood ratio* across signal histories, which is a deterministic increasing function of time due to the martingale property of the likelihood ratio process. 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. 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 of the agent. Then, by tracing out this maximal date- t likelihood ratio over time, we obtain a uni-dimensional *informativeness* measure quantifying by how much the principal is able to reduce incentive pay by deferring longer.

The timing of payouts trades off this information benefit of deferral with the dead-weight costs resulting from the agent’s relative impatience. When the agent’s outside option is low, more informative signals allow 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 agent’s outside option fixes her valuation of the compensation package, the principal’s rent extraction motive is absent. Then, optimal contracts minimize weighted average impatience costs subject to incentive compatibility. Using simple convexification arguments we show that optimal contracts *may* now require two payout dates in order to exploit significant changes in the growth rate of

¹The martingale property of the likelihood ratio process captures this information benefit of deferral in the sense of Kim (1995), since the likelihood ratio distribution at any subsequent date is a mean-preserving spread of the likelihood ratio distribution at an earlier date.

informativeness over time. In settings with two payout dates, the early payout date primarily targets the agent’s ex-ante participation constraint whereas the long-term payout date is used to tap late informative signals for incentive purposes.

While our abstract, broadly applicable model is not designed to match particular institutional details, its comparative statics implications are consistent with various stylized facts. For instance, we find that the duration of pay, the weighted average payout time, is shorter if the agent’s outside option is higher or the agency problem is weaker in the sense of lower effort costs. To the extent that one interprets better corporate governance as a way to reduce the benefit from shirking (e.g., via improved monitoring as in Holmstrom and Tirole (1997)), we, thus, obtain the prediction 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 the 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 long as firms in such industries receive relatively little information in early stages and exhibit high informativeness growth in later development stages. If informativeness grows faster at all points in time, our analysis implies that pay duration is shorter when the agent’s participation constraint binds, but longer when the principal has a rent-extraction motive. Extending our baseline binary-action model to allow for a continuum of actions we further show that the duration of pay may be non-monotonic in the induced effort level. As informativeness itself is a function of the chosen action the comparative statics depend on whether the principal “learns” faster under high or low effort.

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 maximum likelihood ratio history. Intuitively, starting from the “risk-neutral” payout time and contingency, it is now optimal to gradually spread out consumption across time and states in such a way that higher rewards are stipulated for higher performance. Due to the costs associated with relative impatience, the model then predicts an increasing performance hurdle over time: To obtain the same pay in subsequent periods performance needs to increase. Such a pattern is commonly observed, e.g., in compensation contracts of general partners

(carried interest) in private equity settings, which usually feature a growing absolute performance hurdle, i.e., a high hurdle rate (see Metrick and Yasuda (2010) and Robinson and Sensoy (2013)). Moreover, 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.

Finally, we apply our modeling framework to financial contracting and show how delayed observability of performance measures in moral hazard provides novel predictions for optimal security design (building on Innes (1990) and Hébert (2015)), in particular the optimal *maturity structure* of an entrepreneur’s financing decisions. We show that insiders optimally receive positive payouts if and only if performance is above a cutoff that is increasing in time. This implies a rich dynamic payoff structure for insiders’ and outsiders’ claims, resulting from the trade-off between the entrepreneur’s liquidity needs and the increased informativeness associated with new performance signals available to investors.²

Literature. The premise of our paper 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 information systems via the arrival of additional signals, so that these can be ranked in the sense of Holmstrom (1979) or Kim (1995). The key difference of our paper 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 liquidity needs 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) and Bergemann and Morris (2016)), and, in particular, its applications to moral hazard settings, Georgiadis and Szentes (2018) and Li and Yang (2017)).

More concretely, our paper belongs to a small, but growing literature that analyzes moral hazard setups in which the agent’s action has persistent effects. Hopenhayn and Jarque (2010) analyze 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*,

²In contrast, the optimal maturity structure in DeMarzo and Fishman (2007), Biais et al. (2007), DeMarzo and Sannikov (2006), and Oehmke et al. (2017) is derived from a repeated moral hazard model without persistence.

which is the focus of our paper. 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 system available to the principal. As is well known, in repeated-action settings the timing of pay may play an important role even when actions are immediately and perfectly observed (cf., Ray (2002)): In this literature, backloading of rewards to the agent has the benefit that it incentivizes both current as well as future actions. Work by Jarque (2010), Edmans et al. (2012), Sannikov (2014), or Zhu (2017) combines the effects of repeated 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

2.1 Setup

We consider a principal-agent problem in environments where the principal observes informative signals about the agent's action over time. Time is continuous $t \in [0, \bar{T}]$.³ At time 0, the agent A takes an unobservable action $a \in \mathcal{A} = \{a_L, a_H\}$. 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 $k_L = k_H - \Delta k$, where $\Delta k > 0$. As is standard, we suppose that the principal P wants to implement the high action, and we subsequently fully extend our analysis to the case where a is continuous (see Section 3.2.1).

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.⁴ 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. 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 $\mathbb{P}^L \ll \mathbb{P}^H$.

³The assumption of a finite horizon \bar{T} is not crucial for our results. 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.

⁴Formally, the index set of the stochastic process X_t can be any subset of $[0, \bar{T}]$.

The following three illustrative examples represent information environments covered by our framework.⁵

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

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 $a \in \mathcal{A}$ 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 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.

A compensation contract \mathcal{C} stipulates transfers from the principal to the agent as a function of the information available at the time of payout. The principal can commit to any such contract. Formally, a contract is represented by a cumulative compensation process b_t that is càdlàg and adapted to \mathcal{F}_t^X . We restrict attention to payout processes that satisfy limited liability of the agent, i.e., $db_t \geq 0$.

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.⁶ That is, the discount rates of the agent, r_A ,

⁵ Example 1 captures in reduced-form central features of Manso (2011) and Zhu (2017). Example 2 is a generalization of Hartman-Glaser et al. (2012) and Malamud et al. (2013) by allowing for arbitrary correlation structures between failure events. Example 3 is a one-time action version of Sannikov (2014).

⁶ See, e.g., DeMarzo and Duffie (1999), DeMarzo and Sannikov (2006), or Opp and Zhu (2015).

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. When the agent is risk-neutral (see Section 3), the instantaneous utility transfer to the agent, denoted as 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

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

$$V := \mathbb{E}^H \left[\int_0^{\bar{T}} e^{-r_A t} dv_t \right] - k_H \geq R, \quad (PC)$$

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

$$db_t \geq 0 \quad \forall t. \quad (LL)$$

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

The principal's objective is to minimize the present value of wage cost W (discounted at the principal's rate r_P). 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 .⁷ Second, incentive compatibility (IC) requires that it is optimal for the agent to choose action a_H given \mathcal{C} .

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), limits on transfers (Condition 2) or the utility function (Condition 3) that guarantee existence of a solution to Problem 1.

⁷Since the agent in our model only 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 $t = 0$.

2.2 Preliminary analysis

Before proceeding with the formal analysis it is useful to apply insights from optimal static compensation design to our dynamic setup. In particular, in standard static principal-agent settings, the principal evaluates agent performance in likelihood ratio units and cares about additional signals if and only if they affect likelihood ratios (see, e.g., Holmstrom (1979) and Kim (1995)). As the agent's action has persistent effects in our setting, additional informative signals arrive over time, augmenting observed histories. Thus, the analogous date- t performance measure is the likelihood ratio of the entire *history* of realized signals $h^t = \{x_j\}_{0 \leq j \leq t}$. For instance, in Example 1 the date-2 history $\check{h}^2 = (f, s)$ has probability mass $p^H = \rho(1 - \rho^2)$ under the high action and $p^L = \frac{1}{4}$ under the low action so that the (transformed) likelihood ratio satisfies $l_2(\check{h}^2) = 1 - \frac{p^L}{p^H}$,⁸ and similarly for each possible realization of h^2 . For general information environments, the likelihood ratio process satisfies

$$L_t := 1 - \mathbb{E}^H \left[\frac{d\mathbb{P}^L}{d\mathbb{P}^H} \middle| \mathcal{F}_t^X \right], \quad (1)$$

where existence of the Radon-Nikodym derivative $\frac{d\mathbb{P}^L}{d\mathbb{P}^H}(\omega)$ follows from the Radon-Nikodym Theorem given $\mathbb{P}^L \ll \mathbb{P}^H$. The stochastic process L_t then directly implies a date- t distribution over likelihood ratios which we denote by $F_t(l) = \Pr(L_t \leq l)$ with support $\mathbf{L}_t \subseteq (-\infty, 1]$. The supremum and infimum of \mathbf{L}_t are denoted by \bar{L}_t and \underline{L}_t respectively. Importantly, the only restriction that consistent “updating” about the agent's action puts on the likelihood ratio process is evident from (1): For any signal process X_t the corresponding likelihood ratio process L_t is a *martingale* with respect to \mathcal{F}_t^X under \mathbb{P}^H so that the unconditional expectation satisfies $\mathbb{E}^H[L_t] = 0$.

It is now convenient to treat the (univariate) likelihood ratio process as the primitive of the information environment (rather than the potentially multi-dimensional signal process X_t). In particular, to ensure statistical identifiability of the actions a_L and a_H , we assume that $F_{\bar{T}}(0) < 1$. The following Lemma justifies this reduced-form approach.

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

Intuitively, if two different date- t signal histories share the same performance measure (in terms of likelihood ratios), it is optimal to pay out the same amount (strictly so, if the agent is risk-averse). Formally, this result relies on both time-separability of the

⁸Following Tirole (2006), this transformation of the standard likelihood ratio p^H/p^L is made for expositional reasons in expressing incentive constraints.

agent’s preferences as well as the absence of private savings.⁹ While these assumptions are trivially satisfied in our baseline setting with a risk-neutral agent, absence of private savings is key for preserving tractability in the risk-averse setting (as in Sannikov (2008)).

Figure 1 plots the support and a sample path of the likelihood ratio processes generated by our three example information environments. In general, as the likelihood ratio




Figure 1. Likelihood ratio process for examples. 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. 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 $a_L = -3/2$.

is a martingale the upper (lower) support must be weakly increasing (decreasing) over time. In Example 1 (see left panel), where informative signals only arrive at discrete points in time, the support corresponds to step functions, i.e., is constant in-between two information dates. The sample path corresponds to failures for $t = 1, 2, 3$ and a success in $t = 4$. The notion that an initial failure as well as later successes are indicative of the desired long-run action is then captured by the increase in the likelihood ratio for $t \in \{1, 4\}$ and the decrease for $t \in \{2, 3\}$. Within Example 2, when the failure rate is strictly lower under the high action, as depicted in the middle panel for a case with $j = 1$, the absence of failure up to time t implies that the likelihood ratio moves along its upper support \bar{L}_t , which is a smooth and strictly increasing function of time. Then, as

⁹More precisely, these assumptions ensure that in our setting any variable that forecasts higher moments of the *future* likelihood ratio distribution is irrelevant for *today's* compensation as long as *today's* likelihood ratio is unaffected. However, when the agent’s preferences are not time-separable or a risk-averse agent can save privately, exposure to future uncertainty will be relevant, and it is not necessarily sufficient to only condition on the path of L_t .

soon as failure occurs, the likelihood ratio jumps down and stays constant thereafter as no more information arrives. The lowest possible value \underline{L}_t corresponds to failure at the earliest possible time (date 0). Finally, for the case of Example 3 (with $\rho = 1$) where the relevant performance signal x_t is normally distributed for any $t > 0$, the likelihood ratio distribution has full support $(-\infty, 1)$ for all $t > 0$ (see Shiryaev (1978)). The sample path is continuous and increasing at any t with “good news,” i.e., when dx_t is sufficiently high, while it is decreasing following “bad news” corresponding to dx_t sufficiently low.

We structure our subsequent analysis as follows. First, to build intuition, we consider the case of a risk-neutral agent, which allows for a closed-form expression of payout dates. We then consider the additional implications resulting from agent risk-aversion.

3 Contracting with risk-neutral agent

3.1 Baseline model

Our preliminary goal is to rewrite Problem 1 so as to make the trade-offs for optimal contract design fully transparent. To this end, we introduce three auxiliary variables that capture the three distinct levers that the principal possesses in providing incentives, namely i) the total expected value of payments to the agent, ii) how to spread out expected payments across *time*, and iii) how to spread out expected payments across *states* for a given t . Formally, we capture the first lever by the agent’s time-0 valuation of the contract \mathcal{C} :

$$B := \mathbb{E}^H \left[\int_0^{\bar{T}} e^{-r_A t} dv_t \right], \quad (2)$$

Second, we define w_s as the fraction of the compensation value B that the agent derives from stipulated payouts up to time s , i.e.,

$$w_s := \mathbb{E}^H \left[\int_0^s e^{-r_A t} dv_t \right] / B, \quad (3)$$

so that $w_{\bar{T}} = \int_0^{\bar{T}} dw_t = 1$. Hence, $\int_0^{\bar{T}} t dw_t$ measures the weighted average payout time, i.e., the *duration of the compensation contract*, consistent with the empirical analysis by Gopalan et al. (2014).¹⁰ Third, to measure how the principal spreads out expected pay

¹⁰We, thus, employ a duration measure analogous to the *Macauley* 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.

at date t across likelihood ratio states, we define for all t with $dw_t > 0$

$$\gamma_t(l) := F_t(l) \frac{\mathbb{E}^H [dv_t | L_t \leq l]}{\mathbb{E}^H [dv_t]}, \quad (4)$$

so that $\gamma_t(\bar{L}_t) = \int_{\mathbf{L}_t} d\gamma_t(l) = 1$. Hence, if the principal decided to pay out an equal amount across all states then $\gamma_t(l) = \Pr(L_t \leq l) = F_t(l)$.

To avoid Mirrleesian existence problems in our baseline setting with bilateral risk-neutrality and without payment bounds, we impose the following sufficient condition:¹¹

Condition 1 For each date t , $\Pr^H(\bar{L}_t) > 0$.

Then, since with bilateral risk-neutrality $dv_t = db_t$, we can rewrite the compensation design Problem 1 with B, w_t and $\gamma_t(l)$ as control variables:

Problem 1*

$$W := \min_{B, w_t, \gamma_t(l)} B \int_0^{\bar{T}} e^{\Delta r t} dw_t \quad s.t. \quad (W)$$

$$B \geq R + k_H, \quad (PC)$$

$$B \int_0^{\bar{T}} \int_{\mathbf{L}_t} l d\gamma_t(l) dw_t \geq \Delta k, \quad (IC)$$

$$dw_t \geq 0, d\gamma_t(l) \geq 0 \quad \forall t. \quad (LL)$$

Problem 1* reveals that total compensation costs to the principal are given by the multiplicative interaction of the agent's valuation of the compensation package B and weighted average impatience costs $\int_0^{\bar{T}} e^{\Delta r t} dw_t$. Note that, due to bilateral risk-neutrality, the objective function of the principal does not depend on γ directly since the marginal cost of providing utility to the agent is constant across states. Instead, the choice of γ affects the principal's compensation costs only indirectly via the (IC) constraint. Here, the term $\int_0^{\bar{T}} \int_{\mathbf{L}_t} l d\gamma_t(l) dw_t$ represents the weighted average likelihood ratio of performance signals used in the contract, which we refer to as *contract informativeness* $I_{\mathcal{C}}$.

3.1.1 Maximal-incentives contracts

We derive optimal contracts in two steps. First, we show that it is without loss of generality to focus on the class of “maximal-incentives” contracts, yielding the choice of

¹¹ Existence of an optimal contract can be ensured even when Condition 1 is violated, such as in the case of Example 3, if one imposes bounds on transfers (see Condition 2 in Section 3.2.2) or sufficient risk aversion (see Condition 3 in Section 4).

γ . Second, we derive the optimal timing of pay as governed by w . Then, given γ and w , the agent's valuation of the compensation package, B , follows from (IC) or (PC).

Definition 1 *For any given $t \geq 0$, maximal-incentives contracts (\mathcal{C}_{MI} -contracts) stipulate rewards only for the maximum likelihood ratio \bar{L}_t . That is, for all t ,*

$$d\gamma_t(l) = 0 \quad \forall l < \bar{L}_t.$$

The construction of this contract class draws on insights from static moral hazard models (see e.g., Innes (1990) or, for a textbook treatment, Tirole (2006)). In static principal-agent models with bilateral risk-neutrality and limited liability of the agent maximal-incentives contracts are optimal as long as there is a relevant incentive constraint.¹² Due to the absence of risk-sharing considerations, the agent is only rewarded for the outcome which is most indicative of the recommended action, i.e., the outcome with the highest likelihood ratio given this action, \bar{L}_t (and obtains zero for all other outcomes due to limited liability). This basic logic can be readily extended to our dynamic setting.

Lemma 2 *There always exists an optimal contract from the class of \mathcal{C}_{MI} -contracts. If the shadow price on (IC), κ_{IC} , is strictly positive, any optimal contract is a \mathcal{C}_{MI} -contract.*

By inspection of (IC), it is easy to see that contracts outside the class of \mathcal{C}_{MI} -contracts are strictly suboptimal. By shifting rewards towards the maximum likelihood-ratio, the principal would either be able to reduce the agent's valuation of the compensation package B or reduce deadweight impatience costs by moving payments to an earlier date (or both) while preserving incentive compatibility and satisfying (PC).

Lemma 2 greatly increases tractability as, by restricting attention to the class of \mathcal{C}_{MI} -contracts, one *only* needs to keep track of the upper support \bar{L}_t (see Figure 1) rather than the entire likelihood ratio distribution over time. In particular, we can now write (IC) succinctly as

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

where $I_{\mathcal{C}} = \int_0^{\bar{T}} \bar{L}_t dw_t$ is the contract informativeness of a \mathcal{C}_{MI} -contract.

Before turning to the characterization of the optimal duration of compensation contracts, it is instructive to map the abstract maximum likelihood-ratio states back to

¹²If the incentive constraint is not relevant for compensation costs, the compensation design problem generically has a multiplicity of solutions.

concrete signal histories within the context of our leading examples. Given the optimality of \mathcal{C}_{MI} -contracts, the signal histories that are relevant for incentive compensation at each t are those that induce the maximal likelihood ratio \bar{L}_t and will be denoted by h_{MI}^t . In Example 1 (with $\rho = \frac{2}{3}$ and $\bar{T} = 2$), the histories that give rise to \bar{L}_t are failure in period 1, $h_{MI}^1 = (f)$, and subsequent success in period 2, i.e., $h_{MI}^2 = (f, s)$. We note that, in general, when we do not assume independence of performance signals over time, h_{MI}^{t+1} need not be a continuation history of h_{MI}^t (see Online-Appendix B.1.1). For our Example 2, whenever the high action implies a lower default hazard rate (as in the middle panel of Figure 1) the most informative history at each date t , h_{MI}^t , is survival up to t , which is summarized by $x_t = 0$. In both of these examples Condition 1 holds. In contrast, for the Brownian motion Example 3, where x_t is normally distributed with unbounded support and monotonically increasing likelihood ratio, the maximum likelihood ratio history does not exist, so that Condition 1 is violated.

3.1.2 Optimal payout times

The principal's choice of payout times t (via w_t) can be interpreted as choosing the costly quality of the information system. Intuitively, optimal payout times are pinned down by the trade-off between *impatience costs*, measured by $e^{\Delta rt}$, and gains in *informativeness*, as measured by the maximal likelihood ratio at date t .

Proposition 1 *Informativeness $I(t) := \bar{L}_t \in [0, 1]$ is an increasing function of time.*

Formally, the result follows immediately from the fact that the likelihood ratio L_t is a martingale. Intuitively, since the principal can observe the entire history of signals he could always choose to ignore additional signals if he wanted to do so. Thus, $I(t)$ must be an increasing function of time, formalizing the notion of “time will tell.” As an illustration, consider the specification of Example 2 plotted in the middle panel of Figure 1. In this case, survival up to time t (captured by the survival time distribution $S(t|a)$) is the most informative history, so that $I(t) = 1 - \frac{S(t|a_L)}{S(t|a_H)}$. Hence, informativeness grows at a faster rate, as measured by $I'(t) = \frac{S(t|a_L)}{S(t|a_H)} [\lambda(t|a_L) - \lambda(t|a_H)]$, the greater the difference in the hazard rate under the low action $\lambda(t|a_L)$ and the high action $\lambda(t|a_H)$. In contrast, if the hazard rate under both actions is identical at time t , the principal learns nothing from the absence (or occurrence) of failure so that $I'(t) = 0$.

Our subsequent analysis of the timing of pay considers optimal payout dates for the case when (PC) is slack and when (PC) binds separately. Since it is only interesting to analyze optimal payout times when (IC) is relevant, we initially suppress the trivial

case of slack (IC).¹³ Theorem 1 then synthesizes these results and provides conditions for when each case applies.

PC slack. First, consider the case when (PC) is slack. Then, using $B = \frac{\Delta k}{\int_0^{\bar{T}} I(t) dw_t}$ from (IC*) the objective function in Problem 1* simplifies to

$$W = \Delta k \min_{w_t} \frac{\int_0^{\bar{T}} e^{\Delta r t} dw_t}{\int_0^{\bar{T}} I(t) dw_t}. \quad (5)$$

Thus, the optimal timing reflects the principal's *rent extraction* (RE) motive: The principal can reduce the size of pay as measured by the agent's valuation of the compensation package, B , 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*,”

$$T_{RE} = \arg \max_t e^{-\Delta r t} I(t), \quad (6)$$

where the discount rate Δr reflects the effective cost of deferral. It is now easy to see that a cost-minimizing contract with slack (PC) requires only a single payout date.¹⁴ If $I(t)$ is differentiable at the optimal T_{RE} solving (6), we obtain the intuitive characterization

$$\left. \frac{d \log I(t)}{dt} \right|_{t=T_{RE}} = \Delta r. \quad (7)$$

That is, the principal defers until the (log) growth rate of informativeness, $\frac{d \log I}{dt}$, equals the (log) growth rate of impatience costs, Δr .

PC binds. In contrast, when (PC) binds the agent's valuation of the compensation package is fixed at $R + k_H$, so that the principal's rent extraction motive is absent. Using $B = R + k_H$ Problem 1* becomes

$$W = (R + k_H) \min_{w_t} \int_0^{\bar{T}} e^{\Delta r t} dw_t \quad (8)$$

¹³When the (IC) constraint is irrelevant for compensation costs, the principal can achieve the minimum wage cost imposed by (PC), $W = B = R + k_H$, by making all payments at time 0.

¹⁴In the knife-edge case of multiple global maximizers, one can use Pareto optimality as a criterion to select the earliest payout date. The agent strictly prefers the one with the earliest payout date since $V = \frac{\Delta k}{I(T_{RE})} - k_H$, while the principal is indifferent.

subject to (IC*)

$$I_{\mathcal{E}} = \int_0^{\bar{T}} I(t)dw_t = \frac{\Delta k}{R + k_H} \quad (9)$$

Hence, the principal's objective now is to choose the payout dates that achieve a given *weighted average* informativeness of $I_{\mathcal{E}} = \frac{\Delta k}{R+k_H}$ at lowest *weighted average* impatience costs, $\int_0^{\bar{T}} e^{\Delta r t} dw_t$. It is, hence, useful to define

Definition 2 *The cost of informativeness, $\check{C}(I_{\mathcal{E}})$, is the minimum weighted average impatience cost for a given value of contract informativeness $I_{\mathcal{E}}$.*

Lemma 3 *Let $C(I_{\mathcal{E}}) := e^{\Delta r \inf\{t: I(t) \geq I_{\mathcal{E}}\}}$, then the cost of informativeness, $\check{C}(I_{\mathcal{E}})$, is given by the lower convex envelope of $C(I_{\mathcal{E}})$.*

The function $C(I_{\mathcal{E}})$ may be interpreted as the minimum cost of generating an informativeness level of at least $I_{\mathcal{E}}$ by using contracts that only stipulate payments at a *single payout date* t . Since optimal contracts do not necessarily restrict payouts to a single payout date, but instead allow for any possible weighted average via w_t , the minimum cost of achieving any contract informativeness is appropriately measured by the lower convex envelope $\check{C}(I_{\mathcal{E}})$. By construction, $\check{C}(I_{\mathcal{E}})$ is an increasing and convex function mapping $I_{\mathcal{E}} \in [0, I(\bar{T})]$ into $[1, e^{\Delta r \bar{T}}]$.

Figure 2 illustrates how C and \check{C} can be constructed graphically for two continuous, strictly increasing informativeness processes. In the top left panel of Figure 2, informativeness is strictly concave in time t (Example 2 with $j = 1$ and a Poisson process) while impatience costs grow exponentially. Since $C = e^{\Delta r I^{-1}(I_{\mathcal{E}})}$ is thus already strictly convex (see right top panel of Figure 2) there is no benefit from convexification and C and $\check{C}(I_{\mathcal{E}})$ coincide.¹⁵ In contrast, in the bottom left panel of Figure 2, the underlying informativeness process features two phases of high growth. As a result, the impatience costs associated with single-date contracts, $C = e^{\Delta r I^{-1}(I_{\mathcal{E}})}$, exhibit non-convexities, so that, for $I_{\mathcal{E}} \in (I(T_S), I(T_L))$, minimal impatience cost are strictly lower, $\check{C} < C$ (see bottom right panel of Figure 2). The subsequent Lemma 4 links the benefits of convexification to the required number of payouts and their respective timing under an optimal contract.

Lemma 4 *Timing of \mathcal{C}_{MI} -contracts with binding (PC)*

- 1) **Single-date:** *The optimal contract can be implemented with a single payout date T_1 if and only if $\check{C}\left(\frac{\Delta k}{R+k_H}\right) = C\left(\frac{\Delta k}{R+k_H}\right)$. Then, T_1 solves $I(T_1) = \frac{\Delta k}{R+k_H}$.*
- 2) **Two-dates:** *Otherwise, the contract requires a short-term payout date T_S and a long-term date $T_L > T_S \geq 0$. The respective payout dates define the boundary points of the*

¹⁵More generally, as long as I is weakly concave, C will be strictly convex.

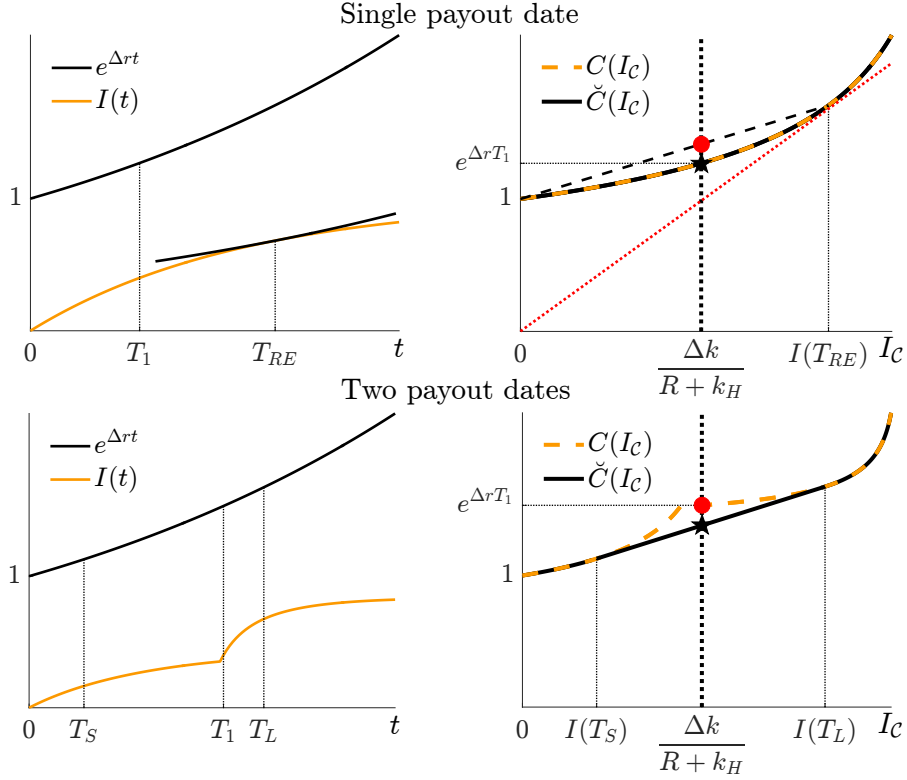


Figure 2. Convexification benefits and the number of payout dates: The upper panels plot a case with strictly convex C (so that a given level of informativeness is optimally achieved with one payout date). In the lower panels, contracts with single payout dates are strictly suboptimal for any $I_\ell \in (I(T_S), I(T_L))$.

linear segment of \check{C} that contains $I_\ell = \frac{\Delta k}{R+k_H}$. The fraction of B derived from payouts at date T_S is given by $w_S = \frac{I(T_L) - \frac{\Delta k}{R+k_H}}{I(T_L) - I(T_S)}$.

Economically, non-convexities in C , and hence the optimality of multiple payment dates, arise if I features sufficient changes in the growth rate of informativeness relative to the growth rate of impatience costs, such that convexification is generically required in discrete-information settings (as in Example 1, see also Online Appendix Section B.1.2). To build further intuition, it is now instructive to revisit the continuous-information examples plotted in Figure 2.

In the upper panel, where the growth rate of informativeness is constant, the use of two payout dates is suboptimal for any level of the outside option R , which governs variation in the (exogenously) required contract informativeness $I_\ell = \frac{\Delta k}{R+k_H}$. Compared to the optimal timing of pay with slack (PC), where contract informativeness, $I_\ell = I(T_{RE})$, is chosen optimally to reduce agency rents,¹⁶ the duration of pay is *shorter* when (PC)

¹⁶It is easy to verify that if $I(T_{RE}) < \frac{\Delta k}{R+k_H}$, then (PC) is slack (see Theorem 1).

binds. To illustrate this point graphically, notice that the optimal timing choice in (6) implies that $I(T_{RE})$ can be identified by the point on $(I_{\mathcal{C}}, C(I_{\mathcal{C}}))$ that minimizes the slope of a ray through the origin (see red dotted line in right top panel of Figure 2). This minimum slope can be interpreted as the shadow price on (IC), κ_{IC} . One may have conjectured that the optimal contract with binding (PC) can then generally be decomposed into the optimal contract with slack (PC), which pays out at date T_{RE} , and an additional sufficiently high (unconditional) date-0 payment to satisfy (PC) at lowest possible impatience costs. This particular conjecture is always wrong when C is strictly convex. The candidate contract (indicated by the red circle in the right top panel of Figure 2) turns out to produce strictly higher wage costs to the principal than the optimal contract that concentrates all payouts at a single payout date T_1 (indicated by the black star).

In the lower panel, where the underlying informativeness process features two phases of high growth, the optimal contract requires two payout dates for all R satisfying $\frac{\Delta k}{R+k_H} \in (I(T_S), I(T_L))$. In particular, to tap “late” increases in informativeness, the optimal contract now makes a payment at a long-term date T_L to target (IC) and an additional short-term payment at date T_S mainly to satisfy (PC) at lower impatience costs. The optimal choice of T_S and T_L generates a strict improvement over the single-date contract that pays out exclusively at date T_1 (see red circle in top right panel). We note that while $T_S > 0$ in this example, it is possible that an optimal contract features an up-front payment of $T_S = 0$ (see Online Appendix Section B.1.2 for a discussion).

3.1.3 Optimal contracts and comparative statics

Synthesizing the cases with (PC) binding and (PC) slack, we can now fully characterize optimal contracts. Together with the conditions for the optimality of \mathcal{C}_{MI} -contracts derived in Lemma 2 we thereby obtain a characterization of optimal contracts based on the solution to Problem 1*.¹⁷ For completeness, the characterization also includes the less interesting case when (IC) is slack ($\kappa_{IC} = 0$).

Theorem 1 *In an optimal contract, a strictly positive bonus is paid out if and only if $L_{T^*} = \bar{L}_{T^*}$, where the optimal payout dates T^* are characterized as follows.*

1. If $R \leq \bar{R} = \frac{\Delta k}{I(T_{RE})} - k_H$, (PC) is slack. The optimal payout date is $T^* = T_{RE}$ as defined in (6). The agent values the compensation package at $B^* = \frac{\Delta k}{I(T_{RE})}$.

¹⁷ Given B , w and γ (or h_{MI}^t), one then obtains $db_t = dv_t$ from $\mathbb{E}^H[dv_t]e^{-r_A t} = Bdw_t$.

2. If $R > \bar{R}$, (PC) binds, so that $B^* = R + k_H$ and $I_\ell = \frac{\Delta k}{R+k_H}$. If $I(0) \leq \frac{\Delta k}{R+k_H}$, (IC) binds and the optimal contract requires maximally two payout dates T^* as characterized in Lemma 4. If $I(0) > \frac{\Delta k}{R+k_H}$, (IC) is slack and all payments are made at date 0.

Theorem 1 summarizes the intuitive characterization of the timing of optimal contracts in general information environments. From this characterization we also obtain the associated wage cost to the principal:

$$W = \begin{cases} \frac{\Delta k}{I(T_{RE})} e^{\Delta r T_{RE}} & R \leq \bar{R} \\ (R + k_H) \check{C} \left(\frac{\Delta k}{R+k_H} \right) & R > \bar{R} \end{cases}. \quad (10)$$

Depending on the particular application at hand W can then be substituted into the principal's objective function to determine whether implementing a_H is indeed optimal, the second step in the structure of Grossman and Hart (1983).

Using the characterization of the optimal timing of pay in Theorem 1, it is now also possible to analyze its comparative statics

Corollary 1 *The duration of the 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 the agent's valuation of the compensation package, B , and more informative performance signals, I_ℓ , are substitutes for providing incentives to the agent, i.e., $BI_\ell = \Delta k$. When an increase in the agent's binding outside option R exogenously raises the value of pay to the agent, this substitutability implies that the principal optimally shortens the duration of the compensation package, such as to reduce contract informativeness (strictly so if (IC) and (PC) bind). In contrast, if the agency problem gets more severe, i.e., Δk increases, then the principal relies on both more informative performance signals (longer 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 (see empirical evidence cited in Introduction).

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 $I(t)$ should primarily capture variation in the nature of information arrival across tasks or industries. To this end, it is useful to 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 2 *Suppose $I(t) = \iota e^{\phi g(t)}$ 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 2 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 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 required contract informativeness *level* of $\frac{\Delta k}{R+k_H}$ at an earlier date. Hence, the principal responds by shortening the duration of the contract in response to increases in ι and ϕ . Firms in industries with high R&D expenditures may feature low informativeness levels in the beginning (low ι), but higher informativeness growth at subsequent dates, e.g., upon patent acceptance or rejection. Hence, we would expect firms in such industries to exhibit longer pay duration (see evidence in Baranchuk et al. (2014) and Gopalan et al. (2014)).

3.2 Extensions

3.2.1 Continuous actions

We now extend our setup to a continuous action set, $a \in \mathcal{A} = [0, \bar{a}]$. The main benefit of employing a continuous-action set is that it allows us to conduct a comparative statics analysis of the duration of pay, $\int t dw_t^*$, in the incentivized action a . Formally, we now consider a family of filtered probability spaces $\{(\Omega, \mathcal{F}^X, (\mathcal{F}_t^X)_{0 \leq t \leq \bar{T}}, \mathbb{P}^a); a \in [0, \bar{a}]\}$, and, to avoid degeneracies assume that \mathbb{P}^{a_1} is equivalent to \mathbb{P}^{a_2} for all $a_1, a_2 \in [0, \bar{a}]$. The associated cost function $k(a)$ satisfies the usual conditions, i.e., it is strictly increasing and strictly convex with $k(0) = k'(0) = 0$ as well as $k'(\bar{a}) = \infty$. To mirror the structure of our analysis so far, we will focus on optimal compensation design, i.e., characterize cost-minimizing contracts to implement a given action a (the first problem in Grossman and Hart (1983)) and relegate the optimal action choice by the principal to Online-Appendix B.2. Also, as is common in static moral hazard problems with continuous actions (see e.g., Holmstrom (1979) and Shavell (1979)) we assume that the first-order approach is valid and provide a sufficient condition in Lemma 5 below. Hence, for each

a , we replace (IC) by the following first-order condition

$$\frac{\partial}{\partial a} \mathbb{E}^a \left[\int_0^{\bar{T}} e^{-r_A t} db_t \right] = k'(a). \quad (\text{IC-FOC})$$

As now local incentives matter according to (IC-FOC), the appropriate measure of agent performance, analogous to the likelihood ratio in the binary action case, is the *score* function which measures the (local) sensitivity of the likelihood function with respect to the action. Formally, denoting by \mathbb{P}_t^a the restriction of \mathbb{P}^a to \mathcal{F}_t^X , we define for each $a > 0$ the likelihood function $\mathcal{L}_t(a|\omega) := \frac{d\mathbb{P}_t^a}{d\mathbb{P}_t^0}(\omega)$ which exists from the Radon-Nikodym Theorem. To illustrate the close connection to the binary-action case we then denote the score by

$$L_t(a) := \frac{\partial}{\partial a} \log \mathcal{L}_t(a|\omega).$$

Here, we impose standard Cramér-Rao regularity conditions used in statistical inference (cf. e.g., Casella and Berger (2002)) by stipulating, in particular, that the score $L_t(a)$ exists and is bounded for any (t, h^t) . Then, letting $[\underline{L}_t(a), \bar{L}_t(a)] \subseteq (-\infty, \infty)$ denote the support of the date- t score distribution,¹⁸ we can define the appropriate informativeness function for action a analogous to the binary action case by

$$I(t|a) := \bar{L}_t(a), \quad (11)$$

where we again assume that $\bar{L}_t(a)$ exists and has positive probability mass (cf., Condition 1). Now, since the score is a martingale, $I(t|a)$ is an increasing function of time (cf., Proposition 1).

It is then immediate that our preceding characterization readily extends to the continuous action case simply by replacing $I(t)$ with $I(t|a)$. For completeness, we restate Theorem 1 in Online-Appendix B.2. We can now also provide a sufficient condition for the validity of the first-order approach.

Lemma 5 *If $\Pr^{\tilde{a}}(\bar{L}_t(a))$ is strictly concave in \tilde{a} for the optimal payout dates $T^*(a)$, then the first-order approach is valid for action a .¹⁹*

For brevity's sake our comparative statics analysis focuses on the case when (PC) is slack. Then, the optimal payment date solves $T^*(a) = T_{RE}(a) = \arg \min_t e^{-\Delta r t} I(t|a)$. Since informativeness, $I(t|a)$, itself is now a function of the implemented action, the

¹⁸ The score is no longer bounded above by one, but this is irrelevant for the further analysis.

¹⁹ The condition is reminiscent of the convexity of the distribution function condition (CDFC) in static models (cf. e.g., Rogerson (1985a)).

comparative statics depend on the characteristics of the signal process. To provide further intuition, it is useful to assume that the first-order condition in (7) applies and has a unique solution. Then, by the implicit function theorem, the sign of the comparative statics of pay duration $T_{RE}(a)$ in a depends on whether the (log) growth rate of informativeness, $\frac{d \log I}{dt}$, increases or decreases in a , i.e.,

$$\text{sgn} \left(\frac{dT_{RE}(a)}{da} \right) = \text{sgn} \left(\frac{d}{da} \frac{d \log I(t|a)}{dt} \Big|_{t=T_{RE}(a)} \right). \quad (12)$$

The following Lemma illustrates that all comparative statics are generically possible even within a fixed parametric Example.²⁰

Lemma 6 *Consider Example 2 with $j = 1$ and generalized Gamma survival time distribution $S(t|a) := \frac{\Gamma(\beta, (\lambda(a)t)^p)}{\Gamma(\beta, 0)}$ where $\lambda(a) > 0$ is a strictly decreasing function of a , $p > 0$ is a constant, and $\Gamma(\beta, x) := \int_x^\infty s^{\beta-1} e^{-s} ds$ denotes the upper incomplete Gamma function. Then, the payout date $T^*(a)$ of the cost-minimizing compensation contract is*

- 1) *strictly increasing in the action if $\beta > 1$,*
- 2) *independent of the action if $\beta = 1$, and,*
- 3) *decreasing in the action if $\beta < 1$.*

First, consider the special case of an exponential distribution ($\beta = p = 1$). With a continuous action set, any i.i.d. process yields an informativeness function that is linear in time (here: $I(t|a) = \frac{t}{a^2}$). As a result, the log-growth rate is independent of the action, here, $\frac{1}{t}$, implying a payout date of $T_{RE}(a) = \frac{1}{\Delta r}$ for all actions a . In this case (and for all other i.i.d. processes), the timing of the bonus alone would not provide any information about the induced action. In contrast, when $\beta < 1$, information grows faster for lower effort so that a shorter duration is indicative of *higher*, rather than lower incentives. The opposite comparative static holds for $\beta > 1$. Knowledge of the sensitivity of the information process to the action (here captured by the parameter β) is thus crucial for understanding optimal deferral for different levels of the agent's action a .

3.2.2 Payment bounds and security design

So far, the focus of our paper was to provide a tractable characterization of the optimal timing of pay (see Theorem 1). However, in cases where the associated maximal-incentives contracts prescribe high rewards for low-probability events, resource constraints such as limited liability on the side of the principal or regulatory constraints,

²⁰The validity of the first-order approach has to be ensured via appropriate parameterization.

such as bonus caps, may become economically 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 consider the following additional constraint:

Condition 2 *The principal faces an upper constraint on the flow payment, $db_t \leq \bar{b}dt$, 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, and, hence, to extend our analysis to information settings where a solution to the original Problem 1 does not exist, including, in particular, the Brownian motion Example 3. We note that the bound \bar{b} may (realistically) be time and/or state dependent to reflect, 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), but we omit this for notational convenience. However, we do assume that \bar{b} is sufficiently high, so that a_H remains implementable.²¹

We now again consider a binary action set and, following the structure of our baseline analysis in the absence of payment bounds, initially study the case with (PC) slack.

Lemma 7 *Suppose (PC) is slack, then there exists a value $\kappa_{IC} \geq \frac{e^{\Delta r T_{RE}}}{I(T_{RE})}$ such that the principal pays the maximum rate, $db_t = \bar{b}dt$, if the discounted likelihood ratio satisfies $e^{-\Delta r t} l_t \geq \frac{1}{\kappa_{IC}}$ and $db_t = 0$, otherwise.*

By construction, adding payment bounds must increase the shadow value on (IC), κ_{IC} , relative to the benchmark without bounds $\frac{e^{\Delta r T_{RE}}}{I(T_{RE})}$. Since the principal can no longer satisfy (IC) by exclusively relying 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 cost-benefit trade-off results both in a wider selection of payout dates and payout states compared to the case without payment bounds, a feature that is also present when (PC) binds. In particular, (PC) will bind whenever the agent's valuation of the contract described in Lemma 7 does not match the agent's outside option R .

²¹ Trivially, if $\bar{b} \rightarrow 0$, 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 whenever $l_t > 0$.

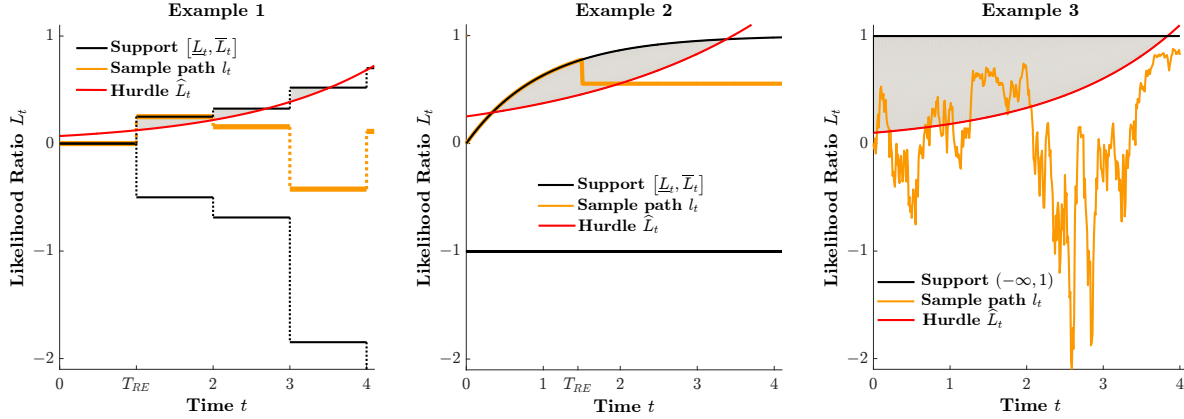


Figure 3. Payment bounds and payout dates for 3 examples: The graphs plot the respective information environments of Figure 1 and the performance hurdle \hat{L}_t when (PC) is slack.

Proposition 2 *There exists $\bar{R}(\bar{b})$ such that, if $R \leq \bar{R}(\bar{b})$, Lemma 7 applies. Otherwise, the shadow price of (PC), κ_{PC} , is strictly positive and the principal pays out the maximum rate if and only if*

$$\kappa_{IC}l_t + \kappa_{PC} \geq e^{\Delta r t}. \quad (13)$$

Regardless of whether (PC) binds or not, we obtain the following robust prediction.

Corollary 3 *The minimal likelihood ratio of the performance signal, $\hat{L}_t = \frac{e^{\Delta r t} - \kappa_{PC}}{\kappa_{IC}}$, required for a strictly positive payout at date t is strictly increasing in t .*

The increase in the performance hurdle over time results from the fact that 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 an appropriately detrended performance measure (i.e., in terms of likelihood ratios, not necessarily just raw profits).

Figure 3 plots the increasing performance hurdle \hat{L}_t (in red) as well as the payment region between \hat{L}_t and \bar{L}_t (shaded in grey) for our 3 examples (for $\kappa_{PC} = 0$). In Example 1, the optimal contract without payment bounds, would exclusively pay out at date $T_{RE} = 1$, whereas the optimal contract in the presence of payment bounds pays out for a wider range of dates. Here, it is particularly instructive to compare the “bad-news” Example 2 to the Brownian Example 3. For Example 2, one requires initially a sufficiently long period without failure (here, ~ 0.5 years) to make it into the payment region. However, eventually, the agent is paid for some time even after failure has occurred (~ 2 years for depicted sample path), which may be interpreted as severance pay. Still, once

the agent moves out of the payment region after a bad shock, he will never return since the performance hurdle is increasing by Corollary 3 and there is no more news after failure. Instead, within the Brownian motion Example 3, there will be immediate payments for sufficiently positive realizations of the likelihood ratio at time 0. Further, even after a sequence of negative realizations that move the agent out of the payment region, the agent may recover for sufficiently positive good news (see plotted sample path).²²

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 design setting where an entrepreneur subject to moral hazard seeks financing from a competitive financial market. In such an application signals usually correspond to output/profit realizations which are commonly assumed to satisfy the MLRP (as in as in Innes (1990)). The entrepreneur then chooses the dynamic structure of insiders’ and outsiders’ cash flows to maximize his payoffs, 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.²³

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 makes the following assumption on the agent’s utility function.

Condition 3 $u : [0, \infty) \rightarrow [0, \infty)$ 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 (following 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

²²This result is reminiscent of the “waiting for news” feature in the dynamic adverse selection environment of Daley and Green (2012).

²³It 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.

to this connection, the characterization of the optimal contract under “standard” risk-aversion can be thought of as a generalization of Lemma 7 and Proposition 2. Assuming that the set of contracts strictly satisfying the constraints in Problem 1 is non-empty, we obtain

Proposition 3 *There exist non-negative shadow prices, κ_{IC} and κ_{PC} , such that $c_t > 0$ if $\kappa_{IC}l_t + \kappa_{PC} > \frac{e^{\Delta r t}}{u'(0)}$ and $c_t = 0$, otherwise. If (PC) is slack, optimal interior consumption is chosen such that the marginal cost of transferring utility to the agent $\frac{1}{u'(c_t)}$ is proportional to the discounted likelihood ratio $e^{-\Delta r t}l_t$, with constant κ_{IC} . If (PC) binds, optimal interior consumption solves*

$$\frac{e^{\Delta r t}}{u'(c_t)} = \kappa_{IC}l_t + \kappa_{PC}. \quad (14)$$

Compared to Proposition 2, the optimality condition in (14) now also reflects the inverse marginal utility term $\frac{1}{u'(c_t)}$ (in addition to the impatience costs $e^{\Delta r t}$).²⁴ 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 $T^* = \arg \max_t e^{-\Delta r t} (\kappa_{IC}\bar{L}_t + \kappa_{PC})$, 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.²⁵

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 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} . In fact, once κ_{PC} is sufficiently high (and 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.

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

²⁴ As long as limited liability does not interfere, (14) implies that the inverse of the *discounted* marginal utility, $1/(e^{-\Delta r t}u'(c_t))$, 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 a repeated-action model (see Rogerson (1985b)).

²⁵ 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 (6) regardless of the value of κ_{IC} .




Figure 4. Contour plot of optimal contract under agent risk-aversion: The graphs plots optimal compensation to the agent for the respective information environments of Figure 1. 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 $l_{0+} \rightarrow 1$ as the likelihood ratio distribution has full support for all $t > 0$.

relative impatience as well as their interplay. For this, it is useful to make explicit the dependence of the agent's compensation on time and the likelihood ratio, i.e., $c_t = c(t, l_t)$ mapping for each $t \in [0, \bar{T}]$ the set of likelihood ratios \mathbf{L}_t into the positive reals.

Proposition 4 *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, $\frac{\partial c_t}{\partial l_t}$, and time, $\frac{\partial c_t}{\partial t}$, satisfy within the payment region:*

$$\frac{\partial c_t}{\partial l_t} = \frac{1}{ARA(c_t)} \frac{1}{l_t + \frac{\kappa_{PC}}{\kappa_{IC}}} > 0, \quad (15)$$

$$\frac{\partial c_t}{\partial t} = -\frac{\Delta r}{ARA(c_t)} \leq 0. \quad (16)$$

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

Proposition 4 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, $\frac{\partial c_t}{\partial l_t} > 0$, as in Holmstrom (1979). *Relative impatience* implies that the performance hurdle is increasing over time, i.e., holding the likelihood ratio fixed, rewards are decreasing over time, $\frac{\partial c_t}{\partial t} < 0$.²⁶ These comparative statics can be directly

²⁶ Intuitively, one may think of this latter effect either as the change in consumption at t when no

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. Empirically, such an increasing performance hurdle is an important feature of compensation contracts in private equity settings (see evidence cited in Introduction). 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, $\frac{\partial c_t}{\partial t} < 0$, 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.

Given that optimal contracts under risk-aversion are now fully characterized, it is instructive to further investigate properties of the optimal contract for two commonly used utility functions.

Corollary 4 *If the agent has CARA utility, $u(c) = 1 - e^{-\rho c}$ with $\rho > 0$, then*

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

If the agent has CRRA utility, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$,²⁷ 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,²⁸ we obtain:

$$\kappa_{IC} = \frac{\Delta k}{2 \int_0^{\bar{T}} e^{-(\Delta r + r_A)t} \mathbb{E}^H [L_t^2] dt}, \text{ and, } \kappa_{PC} = (\Delta r + r_A) \frac{R + k_H}{2}. \quad (17)$$

Thus, 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} , which allows us to obtain a precise metric for quantifying the notion of “only time will tell” even for a setting with risk-aversion.²⁹ The *Fisher Information*, i.e., the variance of

further information arrives, or as comparing two signal histories of different length that imply the same likelihood ratio.

²⁷ Technically, if $\gamma > 1$, CRRA utility violates the boundedness from below in Condition 3. However, this will be immaterial for the existence of optimal contracts as long as the date- \bar{T} likelihood ratio is bounded from below.

²⁸ Limited liability does not bind for any (t, l_t) combination if the likelihood ratio distribution at date \bar{T} is bounded below and R is sufficiently high.

²⁹ Of course, the basic idea that “only time will tell” is still generically true under risk-aversion, since

the likelihood ratio distribution, $\mathbb{E}^H [L_t^2]$, is the sufficient statistic for quantifying the efficiency of the information system for any given t (cf., Jewitt et al. (2008)). Then, since L_t is a martingale, $\mathbb{E}^H [L_t^2]$, is an increasing function of time.³⁰

5 Conclusion

This paper 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 directly modeling the stochastic process of the likelihood-ratio and relying on the martingale property of this performance metric. 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., Biais et al. (2007), DeMarzo and Fishman (2007), and, DeMarzo and Sannikov (2006)). 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 that the performance hurdle is increasing over time, i.e., holding performance fixed, rewards are decreasing over time. Finally, the *interaction* of impatience and risk aversion implies that the pay-performance

the date- t information systems can be ranked in the sense of Kim (1995) or Holmstrom (1979). However, the desire to smooth consumption over time and states implies that the entire likelihood-ratio distribution becomes relevant for compensation costs and a concrete quantification of the benefit of deferral depends on the specification of $u(\cdot)$.

³⁰Technically, *Fisher Information* is defined as the variance of the score (see Section 3.2.1 with a continuum of actions) so that $\mathbb{E}^H [L_t^2]$ should be interpreted as the appropriate adaption of Fisher Information to the case of a binary action.

sensitivity is decreasing over time for any utility function exhibiting decreasing absolute risk aversion. Central predictions of our model are consistent with empirically observed contracts. For example, Metrick and Yasuda (2010) and Robinson and Sensoy (2013) document evidence on an increasing performance hurdle in private equity settings.

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 *multidimensional* 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 (2017))? 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.³¹

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 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

Proof of Lemma 1. The result follows directly from an application of the Halmos-Savage theorem showing that for each t , L_t is a sufficient statistic for $\{x_s\}_{0 \leq s \leq t}$ with respect to a . Then, as the expected cost of transferring utility is time separable, the

³¹To 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 (2017)).

result follows from a straightforward dynamic extension of the arguments in Holmstrom (1979) (Proposition 3). Consider the case of a risk averse agent (the risk neutral case is analogous) and rewrite the compensation design problem with contingent discounted agent utility $v_t := e^{-r_A t} u(c_t)$ as choice variable:

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

$$V := \mathbb{E}^H \left[\int_0^{\bar{T}} v_t dt \right] - k_H \geq R, \quad (\text{PC})$$

$$\mathbb{E}^H \left[\int_0^{\bar{T}} L_t v_t dt \right] \geq \Delta k, \quad (\text{IC})$$

$$v_t \geq 0 \quad \forall t. \quad (\text{LL})$$

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

Lemma A.1 *If the agent is risk-neutral, the shadow value on (IC), κ_{IC} , is zero if and only if $\bar{L}_0 \geq \frac{\Delta k}{R+k_H}$.*

Proof of Lemma A.1 From (PC) and (LL) together with differential discounting we have that $W \geq R + k_H$. Hence, we need to show that $W = R + k_H$ if and only if $\bar{L}_0 \geq \frac{\Delta k}{R+k_H}$. To show sufficiency, consider the \mathcal{C}_{MI} -contract delivering total expected pay $B = (R + k_H)$ with a single payment at $t = 0$ ($w_0 = 1$). This contract trivially satisfies (PC) and (LL). Since the incentive constraint associated with \mathcal{C}_{MI} -contracts is given by $B \int_0^{\bar{T}} \bar{L}_t dw_t \geq \Delta k$ (see IC*), the incentive constraint for $B = (R + k_H)$ and $w_0 = 1$ becomes $(R + k_H) \bar{L}_0 \geq \Delta k$, which is satisfied if $\bar{L}_0 \geq \frac{\Delta k}{R+k_H}$. To show the necessary part, observe that any contract with $W = R + k_H$ cannot feature any delay due to differential discounting, i.e., must satisfy $w_0 = 1$. 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 B contingent on \bar{L}_0 . Now suppose that $\bar{L}_0 < \frac{\Delta k}{R+k_H}$, then the contract requires $B > R + k_H$ in order to satisfy (IC), implying $W > R + k_H$. **Q.E.D.**

Proof of Lemma 2. Take first the case where $\kappa_{IC} > 0$, implying $W = B \int_0^{\bar{T}} e^{\Delta r t} dw_t > R + k_H$ (see Lemma A.1). The proof is by contradiction. So assume that under the optimal contract there exists some t with $dw_t > 0$, for which $\gamma_t(l) > 0$ for some $l < \bar{L}_t$.

Then, there exists another feasible contract with $\gamma_t(l) = 0$ whenever $l < \bar{L}_t$ that yields strictly lower compensation costs. To see this, observe that this new contract maximizes, for given w_t and B , the left-hand side in (IC). So, assume, first, that (PC) is slack. Then, holding w_t constant, the new contract allows for a strictly lower B , thus reducing W . Second, assume that (PC) binds, which, from $\kappa_{IC} > 0$ implies that $w_0 < 1$. Then, holding B constant, the new contract allows to reduce some w_t , $t > 0$, and increase w_0 resulting in lower W . Finally, consider the case when $\kappa_{IC} = 0$. Then, the proof of Lemma A.1 implies that there exists an optimal contract from the class of \mathcal{C}_{MI} contracts. **Q.E.D.**

Proof of Proposition 1. The result follows directly from the martingale property of L_t . **Q.E.D.**

Proof of Lemma 3. See main text. **Q.E.D.**

Proof of Lemma 4. As has been shown in the main text, the optimal \mathcal{C}_{MI} -contract with binding participation constraint requires contract informativeness of $I_\mathcal{C} = \frac{\Delta k}{R+k_H}$ with associated cost of informativeness of $\check{C}(\frac{\Delta k}{R+k_H})$. As, from Lemma 3, $\check{C}(I_\mathcal{C})$ is the lower convex envelope of $C(I_\mathcal{C})$, it is immediate that at most 2 payout dates are sufficient for achieving $\check{C}(\frac{\Delta k}{R+k_H})$. These are generally characterized by $I_S = I(T_S)$ and $I_L = I(T_L)$, where $I_S = \sup \left\{ I_\mathcal{C} \leq \frac{\Delta k}{R+k_H} : \check{C}(I_\mathcal{C}) = C(I_\mathcal{C}) \right\}$ and $I_L = \inf \left\{ I_\mathcal{C} \geq \frac{\Delta k}{R+k_H} : \check{C}(I_\mathcal{C}) = C(I_\mathcal{C}) \right\}$, which we refer to as the boundary points of $\check{C}(I_\mathcal{C})$ around $I_\mathcal{C} = \frac{\Delta k}{R+k_H}$. If $\check{C}(\frac{\Delta k}{R+k_H}) = C(\frac{\Delta k}{R+k_H})$ we have $I_S = I_L$ and, hence, B is paid out at a single date $T_S = T_L =: T_1$. Else, there are two payment dates, $T_S < T_L$ and the fraction of B paid out at T_S is obtained from (9). **Q.E.D.**

Proof of Theorem 1. It follows from Lemma 2 that, given Condition 1, Problem 1 has a solution within the class of \mathcal{C}_{MI} -contracts, i.e., the solution to Problem 1* solves Problem 1 with bilateral risk neutrality. Consider now, first, the relaxed problem ignoring (PC). Then, as shown in the main text, the optimal payout time is given by T_{RE} as characterized in (6) which implies from (IC) that $B = \Delta k / I(T_{RE})$. Then (PC) is indeed satisfied, if and only if $B \geq R + k_H$ which is equivalent to $R \leq \bar{R} := \Delta k / I(T_{RE}) - k_H$. Else, (PC) must bind under the optimal contract, i.e., $B = R + k_H$. The optimal timing of pay then depends on whether (IC) is relevant for compensation costs, which from Lemma A.1 is the case if and only if $I(0) < \frac{\Delta k}{R+k_H}$. Hence, if $I(0) < \frac{\Delta k}{R+k_H}$, the optimal payout times are as characterized in Lemma 4, while for $I(0) \geq \frac{\Delta k}{R+k_H}$ all payouts are made at date 0. **Q.E.D.**

Proof of Corollary 1. From Theorem 1, these comparative statics hold trivially if $I(T_{RE}) < \frac{\Delta k}{R+k_H}$ so that (PC) is slack. In this regime 1, the duration T_{RE} does not

depend on R and Δk . When $I(T_{RE}) \geq \frac{\Delta k}{R+k_H} > I(0)$, (PC) and (IC) bind (regime 2), and the result follows directly from $I_\ell = \frac{\Delta k}{R+k_H}$, which is decreasing in R and increasing in Δk , together with Lemma 4. Finally, when $I(T_{RE}) \geq I(0) \geq \frac{\Delta k}{R+k_H}$, 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 2. Given differentiability of $I(t)$ and the assumption of an interior payout date, the optimal payment date with slack (PC) is characterized by (7), i.e., $g'(T_{RE}) = \frac{\Delta r}{\phi}$, and the first statement follows. Next, note that with binding (PC) concavity of $I(t)$ implies strict convexity of $C(I_\ell)$ such that Lemma 4 implies a single payout at date T_1 satisfying $I(T_1) = \frac{\Delta k}{R+k_H}$. The second statement then follows directly as for each t , informativeness $I(t)$ is strictly increasing in ι and ϕ . **Q.E.D.**

Proof of Lemma 5. The assumption in the Lemma is sufficient to ensure that, given a contract as characterized in Theorem B.1, the agent's problem $\max_{\tilde{a}} \left\{ \mathbb{E}^{\tilde{a}} \left[\int_0^T e^{-rAt} dv_t \right] - k_H \right\}$ is strictly concave. Hence, the first-order condition in (IC-FOC) is both necessary and sufficient for incentive compatibility. **Q.E.D.**

Proof of Lemma 6. Direct computation gives

$$\frac{d \log I(t|a)}{dt} = \frac{d \log \left(\frac{S_a(t|a)}{S(t|a)} \right)}{dt} = \frac{p}{t} \left(\beta - (\lambda t)^p + \frac{(\lambda t)^{p\beta} e^{-(\lambda t)^p}}{\int_{(\lambda t)^p}^{\infty} s^{\beta-1} e^{-s} ds} \right),$$

such that

$$\frac{d^2 \log I(t|a)}{dadt} = \left[\frac{(\lambda t)^p - (\lambda t)^{p\beta} e^{-(\lambda t)^p} (\beta - (\lambda t)^p) \int_{(\lambda t)^p}^{\infty} s^{\beta-1} e^{-s} ds + (\lambda t)^{p\beta} e^{-(\lambda t)^p}}{\left(\int_{(\lambda t)^p}^{\infty} s^{\beta-1} e^{-s} ds \right)^2} \right] \frac{(-p^2 \lambda'(a))}{\lambda t}.$$

Now, define $x := (\lambda T_{RE}(a))^p$, where $T_{RE}(a)$ solves (7) for given a , then

$$\begin{aligned} \operatorname{sgn} \left(\frac{d^2 \log I(t|a)}{dadt} \Big|_{t=T_{RE}(a)} \right) &= \operatorname{sgn} \left(x \Gamma^2(\beta, x) - (\beta - x) x^\beta e^{-x} \Gamma(\beta, x) - (x^\beta e^{-x})^2 \right) \\ &= \operatorname{sgn} \left((\beta - 1) [x (e^{-x} x^\beta + \Gamma(\beta, x))] \Gamma(\beta - 1, x) - \Gamma(\beta, x) e^{-x} x^\beta \right), \end{aligned}$$

where we have used that $\Gamma(\beta, x) = (\beta - 1) \Gamma(\beta - 1, x) + e^{-x} x^{\beta-1}$ for all $\beta > 0$. The result then follows as the term in square brackets is strictly positive by properties of the Gamma function. **Q.E.D.**

Proof of Lemma 7. 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 $l \in \mathbf{L}_t$. Then, using Lemma 1, the Lagrangian can be written as follows

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

where $\kappa_{IC}, \kappa_{PC}, \kappa_{LL}^t(l), \kappa_{PB}^t(l) \geq 0$ denote 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

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

such that $m_t(l) = \bar{b}$ if and only if $e^{\Delta rt} - \kappa_{IC}l - \kappa_{PC} \leq 0$. The case with slack (PC) then corresponds to $\kappa_{PC} = 0$. **Q.E.D.**

Proof of Proposition 2. To characterize the respective cutoff, fix \bar{b} and consider the implied Lagrange multiplier on (IC), $\tilde{\kappa}_{IC}(\bar{b})$, that results from the solution to the compensation design problem with payment bounds absent a participation constraint (see characterization in Lemma 7). Now define for each t the set $\tilde{\mathbf{L}}_t(\bar{b}) := \left\{ l \in \mathbf{L}_t : e^{-\Delta rt}l \geq \frac{1}{\tilde{\kappa}_{IC}(\bar{b})} \right\}$. Then we have $\bar{R}(\bar{b}) := \int_0^{\bar{T}} \int_{\tilde{\mathbf{L}}_t(\bar{b})} e^{-rAt}\bar{b}dF_t(l)dt - k_H$. The characterization of the optimal contract with binding (PC) then follows from the arguments in the proof of Lemma 7. **Q.E.D.**

Proof of Proposition 3. Consider the same change of variables as in Lemma 1, i.e., $v_t = e^{-rAt}u(c_t)$. This ensures that we have a convex programming problem. The corresponding Lagrangian can then be written as

$$\begin{aligned} \mathcal{L} = & \int_0^{\bar{T}} \int_{\mathbf{L}_t} [e^{-rPt}u^{-1}(e^{rAt}v_t(l)) - \kappa_{IC}lv_t(l) - \kappa_{PC}v_t(l) - \kappa_{LL}^t(l)v_t(l)] dF_t(l) dt \\ & + \kappa_{IC}\Delta k + \kappa_{PC}(k_H + R), \end{aligned}$$

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 $v_t(l)$, we obtain that

$$\frac{e^{\Delta rt}}{u'(c_t(l))} - \kappa_{IC}l - \kappa_{PC} = \kappa_{LL}^t(l), \quad (18)$$

for almost every (t, l) . First note that for given (t, l) the left-hand side of (18) 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 Assumption 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, (18) has a solution $c_t(l) > 0$, where $c_t(l)$ is bounded above since $u'(c) \rightarrow 0$ as $c \rightarrow \infty$ (see Assumption 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 4. The expressions in (15) and (16) follow from implicit differentiation of (14). Then the cross derivative is given by

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

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

References

- Baranchuk, N., R. Kieschnick, and R. Moussawi, 2014, Motivating innovation in newly public firms, *Journal of Financial Economics* 111, 578 – 588.
- Bergemann, D., and S. Morris, 2016, Information design, bayesian persuasion, and bayes correlated equilibrium, *American Economic Review* 106, 586–91.
- Biais, B., T. Mariotti, G. Plantin, and J. Rochet, 2007, Dynamic security design: Convergence to continuous time and asset pricing implications, *Review of Economic Studies* 74, 345–390.
- Casella, G., and R. L. Berger, 2002, *Statistical inference*, volume 2 (Duxbury Pacific Grove, CA).
- Daley, B., and B. Green, 2012, Waiting for news in the market for lemons, *Econometrica* 80, 1433–1504.
- DeMarzo, P., and D. Duffie, 1999, A liquidity-based model of security design, *Econometrica* 67, 65–99.

- DeMarzo, P. M., and M. J. Fishman, 2007, Optimal long-term financial contracting, *The Review of Financial Studies* 20, 2079–2128.
- DeMarzo, P. M., and Y. Sannikov, 2006, Optimal security design and dynamic capital structure in a continuous-time agency model, *Journal of Finance* 61, 2681–2724.
- Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov, 2012, Dynamic ceo compensation, *The Journal of Finance* 67, 1603–1647.
- Georgiadis, G., and B. Szentes, 2018, Optimal monitoring design, Technical report, Northwestern University and London School of Economics.
- Gjesdal, F., 1982, Information and incentives: The agency information problem, *The Review of Economic Studies* 49, 373–390.
- Gopalan, R., T. Milbourn, F. Song, and A. V. Thakor, 2014, Duration of executive compensation, *The Journal of Finance* 69, 2777–2817.
- Green, B., and C. R. Taylor, 2016, Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects, *American Economic Review* 106, 3660–99.
- Grossman, S. J., and O. D. Hart, 1983, An analysis of the principal-agent problem, *Econometrica* 51, 7–45.
- Hartman-Glaser, B., T. Piskorski, and A. Tchisty, 2012, Optimal securitization with moral hazard, *Journal of Financial Economics* 104, 186 – 202.
- Hébert, B., 2015, Moral hazard and the optimality of debt, Technical report, Stanford University.
- Holmstrom, B., 1979, Moral hazard and observability, *Bell Journal of Economics* 10, 74–91.
- Holmstrom, B., and J. Tirole, 1997, Financial intermediation, loanable funds, and the real sector, *The Quarterly Journal of Economics* 112, pp. 663–691.
- Hopenhayn, H., and A. Jarque, 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.

- Jarque, A., 2010, Repeated moral hazard with effort persistence, *Journal of Economic Theory* 145, 2412–2423.
- Jewitt, I., O. Kadan, and J. M. Swinkels, 2008, Moral hazard with bounded payments, *Journal of Economic Theory* 143, 59 – 82.
- Kamenica, E., and M. Gentzkow, 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.
- Kim, S. K., 1997, Limited liability and bonus contracts, *Journal of Economics and Management Strategy* 6, 899–913.
- Levitt, S., and C. Snyder, 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 M. Yang, 2017, Optimal incentive contract with endogenous monitoring technology, Technical report, Duke University.
- Malamud, S., H. Rui, and A. Whinston, 2013, Optimal incentives and securitization of defaultable assets, *Journal of Financial Economics* 107, 111 – 135.
- Manso, G., 2011, Motivating innovation, *The Journal of Finance* 66, 1823–1860.
- Metrick, A., and A. Yasuda, 2010, *Venture capital and the finance of innovation* (Wiley).
- Oehmke, M., H. Zhong, and C. Huang, 2017, A theory of multi-period debt structure, 2017 Meeting Papers 1645, Society for Economic Dynamics.
- Opp, M. M., and J. Y. Zhu, 2015, Impatience versus incentives, *Econometrica* 83, 1601–1617.
- Plantin, G., and J. Tirole, 2018, Marking to market versus taking to market, *American Economic Review* 108, 2246–76.
- Ray, D., 2002, The time structure of self-enforcing agreements, *Econometrica* 70, 547–582.

- Robinson, D. T., and B. A. Sensoy, 2013, Do private equity fund managers earn their fees? compensation, ownership, and cash flow performance, *The Review of Financial Studies* 26, 2760–2797.
- Rogerson, W. P., 1985a, The first-order approach to principal-agent problems, *Econometrica* 53, 1357–67.
- Rogerson, W. P., 1985b, Repeated moral hazard, *Econometrica* 53, 69–76.
- Sannikov, Y., 2008, A continuous-time version of the principal: Agent problem, *The Review of Economic Studies* 75, 957–984.
- Sannikov, Y., 2014, Moral hazard and long-run incentives, Unpublished working paper, Princeton University.
- Shavell, S., 1979, Risk sharing and incentives in the principal and agent relationship, *The Bell Journal of Economics* 10, 55–73.
- Shiryaev, A. N., 1978, *Optimal stopping rules*, volume 1 (Springer Science & Business Media).
- Tirole, J., 2006, *Theory of Corporate Finance* (Princeton University Press, Princeton).
- Varas, F., 2017, Managerial short-termism, turnover policy, and the dynamics of incentives, *The Review of Financial Studies* hxx088.
- Zhu, J., 2017, Myopic agency, *Review of Economic Studies* forthcoming.

Appendix B Online-Appendix

B.1 Further results for binary action set

B.1.1 Non-i.i.d. example

To illustrate that h_{MI}^{t+1} need not be a continuation history of h_{MI}^t we consider the non-i.i.d. information environment depicted in Figure 5 with $a \in \mathcal{A} \subset [0, 1]$. For this concrete




Figure 5. Example information process. This graph plots an example information environment with discrete information arrival as in Example 1 but where performance signals are not independent over time.

specification, a success is the most informative signal in $t = 1$, i.e., $h_{MI}^1(a) = (s)$, while $h_{MI}^2(a)$ changes with the concrete values of $a_H > a_L$: When $a_H > 1 - a_L$, the maximally informative history at $t = 2$ is a continuation history of $h_{MI}^1(a)$, in particular, $h_{MI}^2(a) = (s, s)$. When $a_H < 1 - a_L$, however, $h_{MI}^2(a) = (f, s)$, i.e., in this case early failure followed by a success is the best indicator of the agent taking the intended action.

B.1.2 Optimality of up-front payment - Example

In this Appendix we provide some further discussion of the optimality of an up-front payment when (PC) is binding using the illustrative example in Figure 6. In this example information environment, information arrives at discrete points in time, such that the function $C(I_{\mathcal{C}})$ is generically non-convex and two payout dates are optimal according to Lemma 4. In particular, for any $I_{\mathcal{C}} = \frac{\Delta k}{R+k_H} \in (I(0), I(t^*))$, as illustrated in the graph, the




Figure 6. Example with optimal upfront payment. This graph plots an example information environment with discrete information arrival in which an upfront payment may be optimal with a binding participation constraint.

optimal payment dates are $T_S = 0$ and $T_L = t^*$. More generally, an up-front payment is optimal for $I_{\mathcal{I}}$ sufficiently small, if $\lim_{I_{\mathcal{I}} \rightarrow 0^+} [C(I_{\mathcal{I}}) - \check{C}(I_{\mathcal{I}})] > 0$, which arises whenever informativeness is sufficiently convex in some interval $[0, t]$.

B.2 Continuous action set

B.2.1 Optimal compensation design

In this Appendix, we formally characterize the optimal contract for the model with continuous action choice described in Section 3.2.1. In particular, given the conditions in Lemma 5 hold, a simple extension of the arguments leading to Theorem 1 with binary action set gives the following result for a continuous action set:

Theorem B.1 *Suppose $I(0|a) \leq \frac{k'(a)}{R+k(a)}$, then (IC) is relevant for compensation costs and action a is optimally implemented with a \mathcal{C}_{MI} -contract.*

1) *If $R \leq \bar{R}(a) = \frac{k'(a)}{I(T_{RE}(a)|a)} - k(a)$, (PC) is slack, the unique optimal payout date is $T^*(a) = T_{RE}(a)$ which solves $T_{RE}(a) = \arg \max_t e^{-\Delta r t} I(t|a)$, and the agent values the compensation package at $B^* = \frac{k'(a)}{I(T_{RE}(a)|a)}$.*

2) *Otherwise, (PC) binds, so that $B^* = R + k(a)$ and $I_{\mathcal{I}} = \frac{k'(a)}{R+k(a)}$. Payments are optimally made at maximally two payout dates $T^*(a)$ which are characterized as follows:*

If $C(I_{\mathcal{E}}|a) = e^{\Delta r \inf\{t: I(t|a) \geq I_{\mathcal{E}}\}}$ and its lower convex envelope $\check{C}(I_{\mathcal{E}}|a)$ coincide at $I_{\mathcal{E}} = \frac{k'(a)}{R+k(a)}$, there is a single payout at $T_1(a)$ which solves $I(T_1|a) = \frac{k'(a)}{R+k(a)}$. Else there are two payout dates $T_S(a) < T_L(a)$ corresponding to the boundary points of the linear segment of \check{C} that contains $I_{\mathcal{E}} = \frac{k'(a)}{R+k(a)}$.

Theorem B.1 summarizes the characterization of the optimal contract for the case with a relevant (IC) constraint. It remains to characterize the (less interesting) case when (IC) is irrelevant for compensation costs such that \mathcal{C}_{MI} -contracts do not apply. Intuitively, this is the case if the principal receives sufficiently precise signals at time 0 (and $R > 0$). In particular, if $I(0|a) > \frac{k'(a)}{R+k(a)}$, \mathcal{C}_{MI} -contracts would provide excessive incentives, violating (IC-FOC).³² Hence, deferral is not needed to provide incentives:

Lemma A.2 *If $I(0|a) > \frac{k'(a)}{R+k(a)}$, \mathcal{C}_{MI} -contracts do not apply. (PC) binds and all payments are made at time 0, $w^*(0) = 1$, and $B^* = R + k(a)$.*

We have now completely characterized optimal compensation contracts to implement any given action a . The associated wage cost to the principal follows immediately:

$$W(a) = \begin{cases} \frac{k'(a)}{I(T_{RE}(a)|a)} e^{\Delta r T_{RE}(a)} & R \leq \bar{R}(a) \\ (R + k(a)) \check{C}\left(\frac{k'(a)}{R+k(a)} \middle| a\right) & R > \bar{R}(a) \end{cases} \quad (19)$$

B.2.2 Optimal action choice

So far, the analysis has focused on the principal's costs to induce a given action, $W(a)$. In this Appendix we discuss the principal's preferences over actions and the resulting equilibrium action choice, the second problem in Grossman and Hart (1983). We capture the benefits of an action a to the principal by a strictly increasing and concave bounded function $\pi(a)$. Here, $\pi(a)$ could simply be interpreted as the principal's utility derived from action a , or may, more concretely, correspond to the *present value* of the (gross) profit streams under action a . For instance, take Example 2 with $j = 1$ and an exponential arrival time distribution $S(t|a) = e^{-\frac{t}{a}}$, $a > 0$, and suppose that the agent is a bank employee generating a consumer loan or mortgage of size 1, designed as a perpetuity with flow payment f . Through exerting (diligence) effort a , the agent can decrease the likelihood with which a loan subsequently defaults, in which case the asset becomes worthless. For this specification, we can write the bank's (the principal's) expected

³²To see this note that when the principal makes the minimum size of the compensation package required by (PC), $B = R + k(a)$, contingent on $I(0|a) = \bar{L}_0(a)$ (generated by $h_{MI}^0(a)$), the least informative signal within the class of \mathcal{C}_{MI} -contracts, then the marginal benefit of increasing the action to the agent is $I(0|a)(R + k(a))$ which exceeds the marginal cost, $k'(a)$.

discounted revenue for given a as $\pi(a) = \frac{f}{r_P + \frac{1}{a}} - 1$. Generally, given any (gross) profits $\pi(a)$ and compensation costs $W(a)$ the equilibrium action then solves

$$a^* = \arg \max_{a \in \mathcal{A}} \pi(a) - W(a), \quad (20)$$

and, given a solution a^* , the chosen payout times are characterized by Theorem B.1 (and Lemma A.2).