THE PINCHAS SAPIR CENTER

FOR DEVELOPMENT

Tel Aviv University

המרכז לפיתוח ע"ש פנחס ספיר

ליד אוניברסיטת תל אביב

"עמותה רשומה"

580010221

THE PINHAS SAPIR CENTER FOR DEVELOPMENT

TEL AVIV UNIVERSITY

Monotone Activities in Fluctuating Contracting Environments

Daniel Bird[1] and Alexander Frug[2]

Discussion Paper No. 6-2018

July 2018

[1] Bird: Eitan Berglas School of Economics, Tel Aviv University (e-mail: dbird@post.tau.ac.il)

[2] Frug: Department of Economics and Business, Universitat Pompeu Fabra and Barcelona GSE (e-mail: alexander.frug@upf.edu)

אוניברסיטת תל-אביב, תל אביב יפו 6997801 | טל': 03-6409712 | פקס: 03-6405815

Tel Aviv University, Tel Aviv Yafo 6997801 Israel | Tel. 972-3-6409712 | Fax 972-3-6405815 | email: sapircen@tauex.tau.ac.il

Monotone Activities in Fluctuating Contracting Environments*

Daniel Bird Alexander Frug

July 25, 2018

Abstract

A common feature of dynamic contracting environments is that the terms of the periodic interaction (e.g., business opportunities) vary, perhaps stochastically, over time. We consider general fluctuating contracting environments with symmetric information and identify a systematic effect of this variability on qualitative properties of optimal contracts. First, we highlight this effect in a stationary model where the agent is incentivized to exert effort on multiple types of stochastically arriving tasks. We characterize the unique optimal contract and show that the agent's (task-specific) effort decreases and his wage increases over time. Next, we develop the notion of "separable activity," which reveals that the above properties are manifestations of the same, more general, monotonicity result. We then identify a condition on a separable activity that guarantees that the activity will evolve monotonically over time, in a direction that favors the agent, in any optimal contract and any fluctuating environment. The condition is tight in that, whenever it is violated, the monotonicity result is reversed in some contracting environments.

Keywords: Dynamic contracting, Stochastic opportunities.
JEL Classification: D86.

*Bird: Eitan Berglas School of Economics, Tel Aviv University (e-mail: dbird@post.tau.ac.il); Frug: Department of Economics and Business, Universitat Pompeu Fabra and Barcelona GSE (e-mail: alexander.frug@upf.edu). We are grateful to Jose Apesteguia, Eddie Dekel, Kfir Eliaz, Juan Jose Ganuza, Benny Moldovanu, Zvika Neeman, Antonio Penta, Sandro Shelegia, and Rani Spiegler for valuable suggestions and comments. We also appreciate the helpful comments of seminar participants at the 2017 meeting of the Israeli Economic Society and The Hebrew University Economic Theory Seminar 2017. Bird gratefully acknowledges support from the Foerder Foundation and the Sapir center.

1 Introduction

In this paper, we study a prevalent feature of ongoing economic interactions that has received scant attention to date, namely, that the terms of an interaction may change as various opportunities come and go. In a long-term principal-agent relationship, typically, the agent performs certain tasks on behalf of the principal, and the principal compensates the agent for his effort. In some cases, the agent is required to perform identical tasks every day. However, perhaps more often, the agent's assignments are selected in accordance with specific business opportunities or other stochastic events that require particular care. While the majority of the vast literature on contract theory studies the effects of information-related economic phenomena, e.g., adverse selection and moral hazard,[1] it has given much less attention to the inherent variability of the environment where the interaction takes place. We identify a systematic effect of this variability on the qualitative features of optimal contracts.

We start by considering a canonical dynamic contracting environment where the standard information asymmetries are muted but there is some uncertainty over the terms of the interaction in future periods. In particular, we assume that short-lived "business opportunities" (or tasks) of different types arrive stochastically over time according to a stationary distribution. When an opportunity arises, the agent can exert costly (task-specific) effort. In return, the principal offers compensation in the form of a periodic wage. We make the standard assumptions that the marginal cost of compensation is increasing while the marginal productivity of effort is decreasing. We assume that both the agent's effort and the task is currently available are perfectly observed by the principal.

A contract in this environment specifies the required task-specific effort and the amount of compensation in each period as functions of previous events. Even though the environment is stationary, the terms of the optimal contract exhibit distinct dynamic patterns. We show that under the unique optimal contract, the agent's wage increases and his (task-specific) effort decreases over time. In particular, the agent is assigned a "rank" that determines his wage and the composition of tasks he is required to perform, and his rank does not decrease over time.

[1] For recent reviews see, e.g., Bolton and Dewatripont (2005) and Edmans and Gabaix (2016).

A closer look at the above increasing-wage and decreasing-effort properties reveals that they are related in two ways. First, the changes are perfectly correlated in that the same stochastic event triggers changes in the periodic wage and the required effort. Second, whenever the required effort or wage is updated, they both shift in the direction that favors the agent. While the first point is specific to the stationary environment, the second one is more general. In fact, the aforementioned properties are different manifestations of the same deeper monotonicity property that holds in general dynamic contracting environments.

In the second part of the paper, we study this monotonicity property in more general contracting environments that change over time. Characterizing optimal mechanisms, even for relatively simple specifications in a non-stationary environment, is challenging. Rather than attempt to slightly generalize the previous stationary contracting environment and solve for optimal mechanisms, we will consider a very rich class of contracting environments while focusing on the dynamic properties of specific components of the contract. We impose no restrictions on the environment, except to assume that there is no asymmetric information.

We define the notion of "separable activity," which, broadly speaking, is a component of the interaction that is available in some (possibly random) periods and that satisfies two separation requirements. First, the payoffs from the activity are additively separable from the payoffs related to other parts of the interaction and, second, the players' activity-related actions do not affect the environment in future periods.

The periodic wage, the agent's effort exerted on a randomly arriving task, routine assignments whose availability does not change over time, joint production that requires a combination of the principal's resources and the agent's effort (say, capital and labor), non-monetary rewarding activities, working hours, etc., are prominent examples of components that are common in many contracting environments and often satisfy the defining conditions of separable activities.

We then present an intuitive (and oftentimes easy to verify) condition on the agent's payoffs from the activity, and show that if this condition is satisfied, then the level of the activity shifts in the agent's favor in any optimal contract in *any* fluctuating contracting environment. The central part of this condition is a function that specifies the agent's maximal activity-related payoff at each possible level of the activity. This function captures the agent's activity-specific incentives to deviate at different levels of the activity. The condition requires that the slope of this function be between zero and one. If this condition is violated, we show that there exist dynamic interactions where the monotonicity result is reversed. Finally, we show that if the slope of this function is negative, but close to zero, any shift in the principal's favor is small; however, if the slope is greater than one, the use of the activity can exhibit large swings in the principal's favor.

The rest of the paper proceeds as follows. Section 2 analyzes the optimal contract in a model where an agent is incentivized to exert effort on tasks that arrive according to an exogenous stationary distribution. Section 3 considers a general fluctuating contracting environment where the terms of periodic interactions change (perhaps endogenously) over time. In this section we define separable activities and derive the conditions under which the levels of such activities shift monotonically, in the agent's favor, as time goes by. Section 4 offers a review of the related literature and Section 5 concludes.

2 A Stationary Model of Random Opportunities

Consider an infinite-horizon (discrete-time) principal-agent interaction in which tasks arrive stochastically over time. When a particular task arrives, it is available for only one period, and during that period the agent can exert (a task-specific) effort $e \geq 0$. There are $I \in \mathbb{N}$ types of tasks and we assume that at most one task is available in each period. Denote the probability of task $i \in I$ being available in a given period by q_i. The immediate payoffs from exerting effort e_i on task i are $-e_i$ for the agent and $\pi_i(e_i) > 0$ for the principal, where $\pi_i(\cdot)$ are increasing, strictly concave, and differentiable.[2] Moreover, we assume that tasks are ordered in the sense that $\pi_i'(e) < \pi_{i+1}'(e)$ for all $i < I, e \geq 0$, and say that task i is better than task i' if $i > i'$. The principal incentivizes the agent to perform certain tasks by offering a periodic wage. The cost of providing u utils to the agent (in a given period) is $c(u)$, where $c(\cdot)$ is increasing, strictly convex, and differentiable. We assume that $\pi_i'(0) > c'(0)$ and $\lim_{e_i \to \infty} \pi_i'(e_i) < \lim_{u \to \infty} c'(u)$

[2]The assumption that at most one task is available is made solely for ease of exposition. If multiple tasks can arrive together, for each $j \in 2^I$ we can denote by q_j the probability that exactly the tasks in j are available. Moreover, if effort is allocated efficiently between the available tasks, the principal's payoff from the aggregate effort exerted on tasks in j is also increasing, concave, and a.e. differentiable.

for all $i \in I$; that is, there are no clearly redundant tasks and it is not profitable to incentivize infinite effort. Finally, we assume that both players are expected-utility maximizers who share the same discount factor δ.

At the beginning of each period, both players observe which (if any) task is available. The agent's choice of effort is observed at the end of the period. The principal has full commitment power, and thus can propose a binding contract at the beginning of the interaction. A history h_t consists of the information from $t = 1$ until $t - 1$, inclusive. A contract, then, is a pair of functions $(r(\cdot), u(\cdot))$. The function $r(\cdot)$ is referred to as a *job description*; $r(h_t) \in [0, \infty)^I$ specifies the required effort in period t, for every type of task, should that task become available (immediately) after h_t. The function $u(\cdot)$ is referred to as a *compensation plan*; $u(h_t)$ specifies the compensation at period t after[3] h_t. The agent does not have commitment power and can walk away at any time. If the agent exercises this option, the continuation payoff of each player is zero.

2.1 The Phase Mechanism

We begin by constructing a particular contract, referred to as the "phase mechanism" (henceforth PM), that we will later prove to be the unique optimal contract. The PM consists of multiple hierarchical phases. Within each phase, the agent's job description, $r(\cdot)$, and compensation plan, $u(\cdot)$, remain constant. The transition between phases is triggered by the arrival of a task that is better than all the tasks that arrived previously. Thus, the mechanism moves monotonically (perhaps with jumps) through the phases until the final absorbing phase is reached. When the mechanism changes phases, the periodic wage strictly increases and the job description weakly decreases in each dimension.

2.1.1 Auxiliary Problems

We now define I auxiliary problems that constitute the building blocks of the PM. For $i \in I$, let P_i denote an optimization problem where a task of type i is currently available, and the players interact until the arrival of a task that is better than i (i.e., a task of type $i' > i$). In this auxiliary problem, the principal is restricted to selecting a stationary effort vector $(e_j)_{j \leq i}$ and a fixed periodic compensation u. Compensation is provided in the final period of the interaction

[3] The implicit assumption that compensation cannot depend on current task availability clearly entails no loss of generality.

but not in the initial period.

As the agent can quit the contract at any time, the principal's choice must satisfy a number of IC constraints. These constraints guarantee that, for every $j \leq i$, the present value of future compensation net the expected cost of *future* effort is weakly greater than the immediate cost of effort, provided that a task of type j is currently available. Clearly, the agent's incentive to walk away when no task is available is weaker than his incentive to do so when (costly) effort is required. Thus, incentive compatibility in these periods follows from the incentive constraints in periods when effort is required.

Formally, the constraint for task j in problem i is given by

$$e_j \leq \frac{\delta}{1-\delta\lambda_i}(u - \sum_{k \leq i} q_k e_k), \qquad (IC_j(i))$$

where $\lambda_i = 1 - \sum_{k>i} q_k$ is the probability that a task better than task i does not arrive in a given period.

Problem P_i is defined as follows:

$$\max_{u,(e_j)_{j \leq i}} \pi_i(e_i) + \frac{\delta}{1-\delta\lambda_i}(\sum_{j \leq i} q_j \pi_j(e_j) - c(u))$$
$$\text{such that } IC_j(i) \text{ holds for all } j \leq i$$

P_i is a convex optimization problem; thus, it has a unique solution. We denote this solution by $u^{(i)}, (e_j^{(i)})_{j \leq i}$.

Next, we derive some important qualitative properties of the solutions of the auxiliary problems, properties that we will later use to establish the optimality of the PM.

Lemma 1. *In the solution to P_i, the only binding constraint is $IC_i(i)$.*

Proof. We start with three simple observations. First, at least one constraint must be binding, as otherwise slightly reducing u increases the value of the problem without violating any constraint. Second, all tasks are assumed potentially profitable ($\pi_j'(0) > c'(0)$); thus, for at least one task $e_j > 0$. Third, $e_i > 0$, as otherwise a higher value can be obtained by increasing e_i by ϵ and decreasing the task for which $e_j > 0$ by $\frac{q_j}{q_i}\epsilon$.

Assume by way of contradiction that $IC_j(i)$ is binding, for some $j < i$. This implies that $e_j \geq e_i$. Consider the following modification for $\epsilon > 0$. Decrease e_j by $\epsilon\frac{1-\delta\lambda_i}{\delta q_j}$ and increase e_i by $\epsilon\frac{1-\delta\lambda_i}{1+\delta(q_i-\lambda_i)}$. It is straightforward to verify that this modification does not violate any of the constraints of the problem P_i.

The first-order effect of this modification on the value of the problem is

$$\begin{aligned}
&\epsilon(\pi_i'(e_i)\frac{1-\delta\lambda_i}{1+\delta(q_i-\lambda_i)}(1+\frac{\delta}{1-\delta\lambda_i q_i}) - \pi_j'(e_j)\frac{1-\delta\lambda_i}{\delta q_j}\frac{\delta}{1-\delta\lambda_i}) \\
&= \epsilon(\pi_i'(e_i) - \pi_j'(e_j))
\end{aligned}$$

As $e_j \geq e_i$ the ranking of the tasks implies that this effect is positive; thus, for a small enough ϵ this modification increases the value of the problem. $\square$

By Lemma 1, the solution of problem P_i can be obtained by maximizing the Lagrangian:

$$\max_{u,(e_j)_{j\leq i}} \pi_i(e_i) + \frac{\delta}{1-\delta\lambda_i}(\sum_{j\leq i} q_j\pi_j(e_j) - c(u)) - \mu_i(e_i - \frac{\delta}{1-\delta\lambda_i}(u - \sum_{j\leq i} q_j e_j))$$

The FOCs of this problem stipulate that the marginal cost of periodic compensation is equal to the marginal benefit from every implemented task. Essentially, the solution ignores the randomness of task arrival and the dynamics of the environment.

Lemma 2. *$\pi_j'(e_j^{(i)}) = c'(u^{(i)})$, for all $j \leq i$ with $e_j^{(i)} > 0$, and $\pi_j'(e_j^{(i)}) \leq c'(u^{(i)})$ when $e_j^{(i)} = 0$.*

Proof. This follows immediately from the FOCs, which are necessary and sufficient for optimality as the objective is concave. $\square$

The simple characterization of the solution to P_i enables us to compare the the agent's compensation in different auxiliary problems.

Lemma 3. *$u^{(i)} > u^{(i')}$ for all $i > i'$.*

Proof. First, suppose that there exists i such that $u^{(i+1)} < u^{(i)}$. We will show that this implies that there exists $j \leq i$ for which $e_j^{(i+1)} < e_j^{(i)}$, leading to the following violation of Lemma 2: since $c'(u^{(i+1)}) < c'(u^{(i)})$ and $\pi'_j(e_j^{(i+1)}) > \pi'_j(e_j^{(i)})$ it cannot be the case that $\pi'_j(e_j^{(i+1)}) = c'(u^{(i+1)})$ and $\pi'_j(e_j^{(i)}) = c'(u^{(i)})$.

We now show that $e_j^{(i+1)} < e_j^{(i)}$ for some $j \leq i$. Recall that in problem P_{i+1} the binding constraint is $IC_{i+1}(i+1)$. I.e.,

$$e_{i+1}^{(i+1)} = \frac{\delta}{1-\delta\lambda_{i+1}}(u^{(i+1)} - \sum_{k \leq i+1} q_k e_k^{(i+1)})$$

This means that $IC_i(i+1)$ is equivalent to

$$e_i^{(i+1)} \leq \frac{\delta}{1-\delta\lambda_i}(u^{(i+1)} - \sum_{k \leq i} q_k e_k^{(i+1)}) \tag{1}$$

Note that the RHS is the agent's expected payoff until the arrival of a task of type $i+1$ or better.

If $u^{(i+1)} < u^{(i)}$ and $e_j^{(i+1)} \geq e_j^{(i)}$ for all $j \leq i$, then from the fact that $IC_i(i)$ is binding in problem P_i, it follows that the inequality in equation (1) is violated, which, in turn, implies that $IC_i(i+1)$ is violated in P_{i+1}.

Now, assume to the contrary that $u^{(i+1)} = u^{(i)}$. The only vector $(e_j)_{j \leq i}$ for which $\pi_j(e_j) = c'(u^{(i+1)})$ is the solution of P_i. By Lemma 2 this implies that $e_j^{(i+1)} = e_j^{(i)}$ for all $j \leq i$. From the binding constraint of P_{i+1} it follows that $e_{i+1}^{(i+1)} = 0$, which contradicts Lemma 2 as $\pi'_{i+1}(0) > \pi'_i(e_i)$ for all e_i. □

The combination of Lemmas 2 and 3 has immediate implications for the effort exerted on different tasks within a problem and the effort exerted on the same task in different problems. In particular, higher effort is exerted on better tasks within each problem, and the effort exerted on task $j \leq i'$ in $P_{i'}$ is lower than the effort exerted on the same task in problem P_i for $i > i'$. Moreover, in both cases the comparison is strict when the choice is interior.

Corollary 1. *Let $j \leq i$.*
1. For $j > 1$, $e_j^{(i)} \geq e_{j-1}^{(i)}$, with a strict inequality whenever $e_j^{(i)} > 0$.
2. For $i < I$, $e_j^{(i)} \geq e_j^{(i+1)}$, with a strict inequality whenever $e_j^{(i)} > 0$.

2.1.2 Definition and Optimality of the PM

Denote by $\mathscr{I}(h_t)$ the best task that has arrived at least once in the history h_t. If no task has arrived set $\mathscr{I}(h_t) = 0$ and, with a slight abuse of notation, let $u^{(0)} = 0$. We refer to $\mathscr{I}(h_t)$ as the current phase of the contract. The PM is then defined as follows: if the agent followed his job description in all previous periods,

$$r_i(h_t) = \begin{cases} e_i^{(\mathscr{I}(h_t))} & \text{for } i \leq \mathscr{I}(h_t) \\ e_i^{(i)} & \text{for } i > \mathscr{I}(h_t) \end{cases}$$
$$u(h_t) = u^{(\mathscr{I}(h_t))},$$

whereas, in the off-path histories following a deviation, the agent receives no wage and is required to exert no effort.

It is easy to see that the PM is incentive compatible and that the IC constraint is binding (only) in periods when the PM changes phases or a task of type $\mathscr{I}(h_t)$ is available.

Proposition 1. *The PM is the unique optimal contract.*

Proof. We start by showing that restricting attention to stationary solutions does not reduce the value of P_i. Consider the general problem for this environment where the principal can choose any incentive-compatible pair consisting of a of history-dependent job definition and a compensation plan. Since this is a convex maximization problem where the objective function is separable in all arguments, if the stationary candidate $u^{(i)}, (e_j^{(i)})_{j \leq i}$ obtained from the solution to P_i is suboptimal, there exists an improvement such that the required effort on one specific task is modified at one particular history, and only the compensation offered immediately after that history is modified and set at the lowest level under which all IC constraints are satisfied. By Lemma 2, the marginal benefit from every implemented task equals the marginal cost of compensation. Since the cost of compensation is convex and the productivity of effort is concave, every such modification will reduce the total expected value for the principal. Therefore, the stationary solution to P_i specifies the unique optimal contract in the auxiliary environment.

We now return to the general environment and denote by C_0 the class of all incentive-compatible contracts for which, whenever a task that is better than all previously available tasks arrives, the agent's continuation utility is zero. It is immediate that the PM is the unique optimal contract in the class C_0. To see this, notice that the restriction to contracts in C_0 implies that it is sufficient to show that the PM attains the highest expected value between any two (subsequent) earliest arrivals of tasks that are superior to all previously available ones. But this follows directly from the observation given in the previous paragraph, and the construction of the PM.

Finally, we show that relaxing the restriction that the solution has to be in C_0 is not profitable for the principal. Suppose that the PM is suboptimal in the class of all contracts. Since, as before, the principal solves a convex optimization problem that is separable in all arguments, there must exist a profitable modification of the following form: i) at a given history in phase $k < I$, the PM is marginally altered in the direction that reduces the agent's expected payoff in the phase (i.e., either the required effort is increased or compensation is decreased), and ii) at a later history that is part of phase $k' > k$, the PM is marginally changed such that the resulting contract is incentive compatible. However, since the marginal cost of compensation and the marginal benefit from effort during phase k are below those of $k' > k$ under the PM, any such modification reduces the principal's expected payoff, a contradiction. The optimal contract is unique due to the concavity of the objective function. □

Even though the principal and agent interact in a stationary environment, their relationship exhibits a dynamic that can be described using the metaphor of a ratchet that only allows advancement in one direction and never returns to previous levels. Under the unique optimal contract, in any realization, the periodic wage and effort exerted on every type of task are given by monotonic step functions. When the wage or the required (task-specific) effort is updated, they jump to a new level, in a direction that favors the agent, and stay at that level until the next stochastic event causes another jump in the same direction. To an outsider observer, this dynamic may resemble promotions, even though the environment is stationary and has no frictions, apart from the random arrival of tasks.

We conclude this section with two observations. First, the derivation of the PM assumes that the principal has full commitment power. However, from the qualitative properties of the PM (decreasing job description and increasing wage) it follows that the principal's continuation value is lowest in the absorbing phase I. Therefore, if the principal's expected continuation payoff in phase I is greater than the cost of providing the periodic compensation

$$\frac{\delta}{1-\delta}(\sum_{i\in I} q_i\pi_i(e_i^{(I)}) - c(u^{(I)})) \geq c(u^{(I)}),$$

the principal has no incentive to renege on his commitment, and the PM constitutes an equilibrium in a dynamic game.

Second, were we to assume that $\pi_i(\cdot)$ is only weakly concave for some i or that $c(\cdot)$ is only weakly convex, the PM would remain an optimal contract, albeit not the unique optimal contract. In the first case, any optimal contract would retain the structure of the PM, with the exception that the task with a linear segment in its payoff could be implemented at a non-stationary effort within a phase (the average discounted effort exerted on the task within the phase would not change). In the second case, optimal contracts could also differ from the PM by the postponement of compensation within a phase or between two phases that share the same marginal cost of compensation. Note that this implies that compensation need not increase with the phase or over time.

Clearly, in the extreme case of a linear cost of compensation, in addition to our mechanism, a trivial optimal contract exists where, upon observing a desired effort, the principal fully compensates the agent in the following period. While this mechanism seems natural in the case of a linear cost of compensation, it cannot be approached as a limit of optimal contracts where the cost of compensation is strictly convex. It is the strict convexity of the cost of compensation that constitutes the link between different periods in our baseline model.

3 General Dynamic Contracting Environments

In this section, we establish that the main qualitative properties of the optimal mechanism characterized for the stationary environment of the previous section are, in fact, different manifestations of a more general feature that holds in dynamic contracting environments.

We consider dynamic interactions between a principal and an agent where the terms of the periodic interaction are stochastic, information is symmetric, and actions are perfectly observed. Specifically, at the beginning of each period t, the terms of period t's interaction are drawn from a commonly known distribution $f(h_t)$, after which the players take actions and receive payoffs. The argument of the distribution function h_t denotes the complete history of all one-period interactions and the players' moves in periods[4] $1, ..., t-1$. As in the previous section, we assume that the principal has commitment power, that both players share the same discount factor, and that even though action availability is publicly observed, the agent cannot be compelled to take any specific action.

As the calendar time, previous "opportunities," and players' past moves may affect the terms of future interactions, this specification is fairly general. The stationary interaction of the previous section clearly belongs to this class of environments. Another example is the focus of Harris and Holmström (1982) and Holmström (1983) who studied environments where there is uncertainty about the worker's productivity in future periods. However, in addition to these relatively simple and well-behaved environments, our modeling approach can also accommodate a wide variety of scenarios including, but not limited to, seasonality, long-term projects, storable investment opportunities, R&D-type investments, etc.

Attempting to derive a complete characterization of optimal contracts in such a general environment is a fool's errand. However, general insights into certain components of the interaction can be derived. To do so, we now develop the notion of separable activities.

3.1 Separable Activities

One possible way to model a separable activity is by specifying a separate game that the players play in selected periods, in addition to other parts of the interaction, and imposing certain separability conditions. As our sole objective is to emphasize the dynamic monotonicity of separable activities, we follow a "reduced-form" approach.[5] Henceforth, we use the terms "separable activity" and "activity" interchangeably.

Our definition will specify a set of possible activity levels (e.g., the set of

[4]The "one-period interaction" is, of course, a simultaneous two-player game that includes a specification of the action space and the players' payoffs, and the entire interaction is a dynamic stochastic game with perfect monitoring.

[5]Our reduced-form approach is akin to the modeling strategies used in Rotemberg and Saloner (1986) and Ray (2002).

the agent's possible effort levels), the players' payoffs from each possible activity level (e.g., the agent's cost and the principal's benefit from every level of the agent's effort), and the agent's maximal activity-related payoff at every intended level (e.g., exert no effort). Intuitively, when this "maximal payoff" is greater than the agent's payoff from a given activity level, it means that the agent can increase his activity-related payoff by deviating within the activity related component of the interaction. Of course, as that specific activity is only one component of a more general interaction, other components (or future appearances of that activity) can be used as incentives not to deviate. Finally, we will impose a (unidirectional) independence restriction between the activity and the contracting environment by which the selected level of the activity does not impact the terms of the interaction in future periods.

3.1.1 Formal Definition of Separable Activities

A separable activity consists of: an interval of possible levels $L \subset \mathbb{R}$ that we normalize to be equal to the agent's payoff from each level (the agent's payoff from level l is l); an increasing and strictly convex function $\kappa : L \to \mathbb{R}$ specifying the principal's cost of engaging in the activity at each level; and a continuous function $D : L \to \mathbb{R}$ specifying the agent's maximal activity-specific payoff for every intended activity level; such that all payoffs enter the player's utilities in an additively separable manner and, for every two histories h_t and $\hat{h}_t$ that differ *only* in the selected activity levels, $f(h_t) = f(\hat{h}_t)$.

Notice that even though our definition rules out the possibility of affecting the future through the selection of specific levels for the separable activity, we do allow the availability of the activity to depend on past actions as these actions change the argument in $f(\cdot)$. This allows us to capture scenarios where, unlike in the environment considered in the previous section, the availability of certain activities is endogenous. In Appendix B we discuss the role of the assumptions that the activity levels belong to a connected support and that the function $D(\cdot)$ is continuous.

3.1.2 Examples of Separable Activities

We now offer some specific examples and illustrate how they can be modeled as separable activities. First, consider the stationary environment of Section

2. For wage, the relevant domain is $L = [0, \infty)$; thus, $\kappa(l) = c(l)$ and, assuming that the agent cannot alter his wage, we have $D(l) = l$. For task-specific effort, note that the agent's payoff is non-positive and so the relevant domain is $L = (-\infty, 0]$, and the principal's "cost" is $\kappa(l) = -\pi(-l)$. As the agent's optimal "effort-specific" deviation is to exert no effort, $D(l) = 0$ for every l.

Notice that, more generally, the activity that represents the agent's effort may induce a less trivial form of $D(\cdot)$. For example, assume that the agent's level of effort is contractible; however, the agent may decide whether to provide high-quality effort (at a cost of l) or low-quality effort (at a cost of αl). If the penalty for breaking a contract is severe, the optimal deviation is to exert low-quality effort and, thus, $D(l) = \alpha l$. In this case, considering the severe penalty for breaking the contract (by exerting zero effort) as an activity-related payoff is the only consistent way to model the effort as a separable activity (otherwise, deviation to level zero changes the set of the principal's moves in the future and violates the separability condition).

The above examples focus on a special class of activities that are controlled by a single player. The discussion in the previous paragraph illustrates that even in this class, non-trivial forms of the function $D(\cdot)$ may arise. Another class of activities where non-trivial functions $D(\cdot)$ arise even more naturally consists of activities where in order to produce a given level of the activity both players need to take certain actions. We refer to this class as "joint production." For example, suppose a good is produced from "labor" provided by the agent and "capital" provided by the principal. In this case, the agent's best deviation can depend on the specific details of the production function (e.g., are labor and capital complements or substitutes?).

First, assume the principal and agent must jointly provide one unit of resources. In some cases, the agent may be able to allocate some, say $\frac{1}{2}$, of the resources provided by the principal for alternative uses. Thus, if we denote the agent's contribution by $x \in [0, 1]$ and the principal's contribution by $1 - x$, the agent's alternative value is $\frac{1-x}{2}$. The possible levels of this activity are $L = [-1, 0]$, where $-l$ is the agent's contribution to joint production, and $D(l) = \frac{1+l}{2}$. Alternatively, suppose that one unit of output is produced from one unit of labor and k units of capital, and that the agent's alternative value from reassigning capital elsewhere when he has x units of capital under his control is $g(x)$. Then, the levels of the activity are $L = (-\infty, 0]$, where $-l$ is the level of production, and $D(l) = g(-lk)$.

3.2 Monotone Use of Separable Activities

It turns out that the key behind the monotonicity result illustrated in Section 2 for wage and task-specific effort is the behavior of the function $D(l)$. Consider the following condition.

Condition $\mathcal{D}$.

$$0 \leq \frac{D(l_2) - D(l_1)}{l_2 - l_1} \leq 1 \quad \forall l_1 < l_2 \in L$$

That is, $D(l)$ is a weakly increasing function with a slope no greater than one. To develop some intuition of the condition, first note that for the separable activities in the stationary model of Section 2 condition $\mathcal{D}$ is satisfied on the boundaries: for wage, the slope equals 1, and for effort it is zero. For an example where the function $D(\cdot)$ is decreasing consider the last example of joint production in the previous subsection where both labor and capital need to increase in order for output to increase. Increasing the level of the activity (decreasing production) decreases the amount of capital provided by the principal, and hence decreases the agent's benefit from misusing the capital he was assigned.

To see how the other bound can be violated, consider an activity that represents the amount of time the principal allows the agent to acquire human capital, and suppose that low levels of l represent investment in general skills while higher levels correspond to acquisition of interaction-specific skills. It may be the case that at low levels of the activity, the agent's optimal use of time is to accumulate general skills ($D(l) = l$), while at higher levels of l the agent will prefer to acquire only the general skills and then spend the rest of his allotted time doing other things ($D(l) > l$). Thus, the slope of $D(\cdot)$ is greater than one for some levels of l.

Our main result is to establish that if condition $\mathcal{D}$ holds, the use of the activity (weakly) increases over time under any optimal contract in *any* possible interaction. To better understand why this is so, first consider how the principal would like to use a specific activity while holding the rest of the interaction fixed if $D(\cdot)$ is constant. The key tension is between the desire to smooth the use of the activity (due to the convexity of $\kappa(\cdot)$) and the need to provide the agent with enough utility from the activity to satisfy his forward-looking IC constraints at every point in time.

Suppose for a moment that the agent takes costly actions only at the very beginning of the interaction and the activity is used to provide sufficient incentives. In this case, it is optimal to use the activity at a constant level. However, when the terms of the interaction change over time, the agent's expected discounted utility (from the rest of the interaction) may decrease over time. In this case a constant level of the activity either violates IC constraints (if the level is set according to the agent's initial utility) or over-compensates the agent. This, in turn implies that, generically, the principal would like the use of the activity to increase over time.

Now, consider the impact of "smoothing" a decrease in the level of the activity on the agent's incentive to deviate. In the early period smoothing decreases the use of the activity. Thus, if $D(\cdot)$ is non-decreasing, smoothing decreases the agent's incentive to deviate in the early period. In the later period, where smoothing increases the level of the activity, the agent's payoff from the activity has increased, but his benefit from deviating may also have increased. If the slope of $D(\cdot)$ is less than one, the first increase is greater than second one, and thus smoothing does not incentivize the agent to deviate in the later period.

Let $l(h_t)$ denote the level of the activity at history h_t and let τ_i denote the time at which the activity is available for the $i-th$ time in a realized infinite history.

Proposition 2. *Suppose that condition $\mathcal{D}$ holds. Then, for any $i<j$, $l_{\tau_i} \leq l_{\tau_j}$ almost surely under any optimal contact.*

Proof. Consider an incentive-compatible contract C in which the realized sequence $(l_{\tau_1}, l_{\tau_2}, \ldots)$ decreases with positive probability. Let h_t be a history after which l declines between periods $t=\tau_s$ and $t'=\tau_{s+1}$ with positive probability. There exist $\Delta>0, p>0$, such that the set Ω of all histories of length t' that are consistent with h_t and for which $l(h_t)-\Delta \geq l(h_{t'})$ satisfies $Pr(\Omega|h_t)=p$. Fix an $\epsilon>0$ for which $\epsilon+\frac{\epsilon}{p\delta^{t'-t}}<\Delta$.

Consider the contract $\hat{C}$ that is obtained from C by modifying the level of the activity as follows: $\hat{l}(h_t)=l(h_t)-\epsilon$, and, at every history $h_{t'} \in \Omega$,

$\hat{l}(h_{t'}) = l(h_{t'}) + \frac{\epsilon}{\delta^{t'-t}p}$. Notice that the original contract is modified only at histories during which the agent followed the recommendation. Moreover, by assumption, this change does not impact the unmodeled part of the interaction or the future availability of the activity. First, we show that $\hat{C}$ is incentive compatible. Then, we show that it increases the principal's expected value from the interaction.

For all histories h_s such that $s \geq t'$ and $h_s \notin \Omega$, the modified contract is identical to the original one. At $h_{t'} \in \Omega$, the agent's continuation utility from following the contract is increased by $\frac{\epsilon}{\delta^{t'-t}p}$ while his best alternative value increases by $D(l(h_{t'}) + \frac{\epsilon}{\delta^{t'-t}p}) - D(l(h_{t'})$. By condition $\mathcal{D}$ this increase is less than $\frac{\epsilon}{\delta^{t'-t}p}$. This, in turn, implies that the agent's incentive to follow the recommendation is weakly greater at all histories with a length of between t and t'. For all histories of length t other than the designated h_t, the contracts C and $\hat{C}$ are identical. Consider h_t where $\hat{l}(h_t) < l(h_t)$. By construction, if the agent follows the modified contract at period t, the expected increase in l at t' balances the decrease in l at h_t. Moreover, by condition $\mathcal{D}$, decreasing l at h_t weakly reduces the agent's value for violating the contract at h_t. Thus, the modified contract is IC at h_t. Finally, it follows that for all histories h_s such that $s < t$ (regardless of whether these histories are consistent with h_t), the agent's continuation utility from any action is unchanged.

To show that this modification is profitable for the principal it is sufficient to show that his expected cost from the activity conditional on reaching h_t under $\hat{C}$ is lower than that under C. Let μ denote the distribution of $l_{t'}$ induced by the distribution of histories in Ω conditional on h_t (a well defined distribution as a truncation of the distribution of $l_{t'}$, conditional on h_t). It is sufficient to show that

$$\kappa(l(h_t) - \epsilon)) + \delta^{t'-t}p \int \kappa(l(h_t') + \frac{\epsilon}{\delta^{t'-t}p})d\mu < \kappa(l(h_t)) + \delta^{t'-t}p \int \kappa(l(h_t'))d\mu$$

Since $\kappa(\cdot)$ is convex, it has a right-hand derivative. With a slight abuse of notation, we denote this derivative by $\kappa'(\cdot)$:

$$\kappa(l(h_t)-\epsilon)+\delta^{t'-t}p\int\kappa(l(h_t')+\frac{\epsilon}{\delta^{t'-t}p})d\mu<$$

$$\kappa(l(h_t))-\epsilon\kappa'(l(h_t)-\epsilon)+\delta^{t'-t}p\Bigg(\int\kappa(l(h_{t'}))d\mu+\frac{\epsilon}{\delta^{t'-t}p}\kappa'(l(h_t)-\Delta+\frac{\epsilon}{\delta^{t'-t}p})\Bigg)=$$

$$\kappa(l(h_t))+\delta^{t'-t}p\int\kappa(l(h_{t'}))d\mu-\epsilon\Bigg(\kappa'(l(h_t)-\epsilon)-\kappa'(l(h_t)-\Delta+\frac{\epsilon}{\delta^{t'-t}p})\Bigg)<$$

$$\kappa(l(h_t))+\delta^{t'-t}p\int\kappa(l(h_{t'}))d\mu,$$

where the last inequality follows from the convexity of $\kappa(\cdot)$ and the choice of ϵ.

If such h_t is reached with positive probability, the modified contract is better than the original one. Otherwise, our assumption that the realized level of the activity decreases with positive probability under C implies that there is a t for which there is a positive measure of histories h_t such that 1) the activity is available for the $j-th$ time at h_t and 2) conditional on h_t, there is a positive measure of histories $h_{t'}$ that are consistent with h_t in which the activity is available for the $j+1-th$ time and $l(h_t)>l(h_{t'})$. For each such h_t, perform the modification as specified above and note that these modifications do not interact with one another as they modify distinct histories. It follows that the modified contract outperforms the original one. □

The general framework developed in this paper enables us to draw connections between seemingly unrelated existing results. For example, although Harris and Holmström (1982) and Holmström (1983) study a competitive market where there is uncertainty about the worker's skill, their setting can easily be embedded in our framework. Adding alternative employment opportunities implies that at any point in time the agent can quit his contract and receive the expected payoff associated with his best alternate offer. Therefore, competition can be incorporated into the model by assuming the general interaction includes "quit" actions, which provide the agent with a payoff equal to that of the best outside offer and in all subsequent periods both players get a payoff of zero. Proposition 2 directly establishes the downward rigidity of wages that is established in both papers.

In more recent work Forand and Zápal (2018) study the key properties of optimal contracts in an environment where different projects arrive stochastically over time according to an exogenous distribution. Projects are separable activities and, as the agent's choice is whether or not to implement a project at the suggested probability, all projects satisfy condition $\mathcal{D}$. Therefore, a small adaptation of Proposition 2 to weakly convex $\kappa(\cdot)$ would imply that there exists an optimal contract in which each project is used monotonically over time and all shifts are in the direction that favors the agent.[6] This, in combination with the immediate observation that, in any given period (after any history), using expensive rather than cheap rewards to provide incentives (projects that benefit the agent but are costly to the principal) as well as incentivizing inefficient rather than efficient investments (projects that benefit the principal but are costly to the agent) is strictly dominated, provides a simple proof for the main result of Forand and Zápal (2018).

Next, we provide a converse to Proposition 2. Namely, we show that if condition $\mathcal{D}$ is violated, there exist interactions where the use of the activity decreases over time. The intuition for this converse result is also straightforward. Consider a contract in which the agent's IC constraint is binding in two periods between which the use of the activity decreases. Moreover, assume that condition $\mathcal{D}$ is violated on the interval connecting these two levels. Then, any attempt to smooth the use of the activity between the two periods will lead to a violation of the IC constraint in one of the two periods. Thus, to prove this result, it is sufficient to construct a counterexample in which under the optimal contract the level of the activity decreases and the agent's IC constraints are binding in all periods. The construction of the counterexample is given in Appendix A.

Proposition 3. *If condition $\mathcal{D}$ is violated, then there exist interactions in which the use of the activity decreases over time.*

Our final result shows that due to this difference small violations of condition $\mathcal{D}$ on the two bounds have an asymmetric impact on the monotonic use of an activity. In particular, if the slope of $D(\cdot)$ is negative but bounded from below, then the maximal drop in the level of the activity is also bounded from below in any interaction. By contrast, if the slope of $D(\cdot)$ exceeds one, then there exist interactions in which the decrease in the use of the activity is large (relative to the size of the interval on which this occurs).

[6]The strict convexity of $\kappa(\cdot)$ is only required in order to establish that this occurs in *every* optimal contract.

To establish proposition 2 we showed that condition $\mathcal{D}$ guarantees that if the principal smooths out a decrease in the level of the activity between periods t and $t' > t$, the new contract is incentive-compatible at both t and t'. However, that analysis may have obscured an intrinsic difference between the incentive compatibility constraints in both periods. If smoothing violates the forward-looking IC constraint at t, incentive compatibility can always be restored by increasing the agent's continuation utility via an increase in $l_{t'}$. By contrast, if smoothing violates the IC constraint at t', the principal may not be able to increase the agent's continuation utility from the contract.

If the slope of $D(\cdot)$ is negative, smoothing out a decrease in the level of the activity might violate the IC constraint at t. However, should this occur the principal can then restore incentive compatibility at t by increasing $l_{t'}$. Thus, a decrease in the level of the activity is consistent with optimality only if the gain from (partially) smoothing its use is outweighed by the cost of increasing its average (discounted) level. Due to the convexity of κ, the principal's gain from slightly reducing the size of the reduction in the level of the activity is increasing in the size of the reduction. By contrast, if the slope of $D(\cdot)$ is bounded from below by $-c$, the cost of restoring incentive compatibility is proportional to c. Thus, when $-c \leq \frac{D(l_2)-D(l_1)}{l_2-l_1} \leq 1 \quad \forall l_2 > l_1$ there exists an upper bound on the maximal decrease in the level of the activity and, moreover, this bound converges to zero with c.

On the other hand, if the slope of $D(\cdot)$ is greater than one, smoothing out a decrease in the level of the activity might violate the IC constraint at t'. In particular, if the IC constraint at t' is binding, any attempt to smooth out the reduction in the level of the activity will violate the IC constraint at t'. This, in turn, implies that when $1+c \leq \frac{D(l_2)-D(l_1)}{l_2-l_1} \quad \forall l_2 > l_1$, large reductions in the level of the activity are consistent with optimality if the principal cannot alter the contract in the agent's favor after time t' via other means.

Formally, for any $c > 0$ define $X_c = \{x : \inf_{l \in L} \frac{\kappa'(l+x)}{\kappa'(l)} < 1 + c\}$ and let $\bar{x}_c = \sup\{X_c\}$ (recall that $\kappa'(\cdot)$ is the right-hand derivative of $\kappa(\cdot)$). Note that $\lim_{c \to 0} \bar{x}_c = 0$.

Proposition 4. *For any $c > 0$,*

1. *If $-c \leq \frac{D(l_2)-D(l_1)}{l_2-l_1} \leq 1$, then almost surely $l_{\tau_i} \geq l_{\tau_j} - \bar{x}_c \quad \forall i < j$.*
2. *If there exists an interval $\tilde{L} \subset L$ such that $1+c \leq \frac{D(l_2)-D(l_1)}{l_2-l_1} \quad \forall l_1 < l_2 \in$*

$\tilde{L}$, then there exist environments in which the use of the activity decreases by $|\tilde{L}|$.

Proof. Assume to the contrary that in an optimal contract with strictly positive probability the use of the activity decreases by more than $\Delta > \bar{x}_c$ between the $i-th$ and $j-th$ use for $j > i$. Let h_t be a history after which l declines between period $t = \tau_i$ and $t' = \tau_j$ by more than Δ with strictly positive probability. There exist $p > 0$, such that the set Ω of all histories of length t' that are consistent with h_t and for which $l(h_t) > l(h_{t'}) + \Delta$ satisfies $Pr(\Omega|h_t) = p$. For each such history we show that there exists a profitable modification of the contract that does not violate any IC constraint.

Changing the level of the activity at t to $\tilde{l}_t = l_t - \epsilon$ and at t' to $\tilde{l}_{t'} = l_{t'} + \alpha\epsilon$ is profitable for sufficiently small ϵ if $\alpha < \frac{1}{p\delta^{t'-t}} \frac{\kappa'(l_t)}{\kappa'(l_{t'})}$. Moreover, if $\alpha > \frac{1}{p\delta^{t'-t}}$ such a change slackens IC constraints at $\{1, 2, \ldots, t-1, t+1, \ldots, t'-1\}$, has no impact on IC constraints after t', and, as the slope of $D(\cdot)$ is less than one, does not violate the IC constraints at t'. Thus, $l_t > l_{t'} + \Delta$ only if any small decrease in l_t, which is offset by a subsequent increase in $l_{t'}$ that maintains IC at t, is not profitable.

The supremum of the marginal increase in the agent's on-path payoff from a profitable modification is $(\frac{\kappa'(l_{t'}+\Delta)}{\kappa'(l_{t'})} - 1)$. Thus, a sufficiently small modification of the type suggested above is IC at t if $\frac{\kappa'(l_{t'}+\Delta)}{\kappa'(l_{t'})} - 1 \geq c$. Therefore, a decrease of size Δ in the level of the activity can be part of an optimal contract only if $\inf_{l\in L} \frac{\kappa'(l+\Delta)}{\kappa'(l)} < 1 + c$. However, as $\Delta > \bar{x}_c$, which is the supremum of the set $\{x : \inf_{l\in L} \frac{\kappa'(l+x)}{\kappa'(l)} < 1 + c\}$, there exist profitable modifications of the contract.

We prove the second part of the proposition by example. To construct a decrease of size $|\tilde{L}|$ in the use of the activity let $D(l) = (1 + c)l$ on $\tilde{L}$, and consider the counterexample used in the the second part of Proposition 3. By selecting $L_2 = \tilde{L}$ we have an environment in which $l_1 = \max_{l\in\tilde{L}} l, l_2 = \min_{l\in\tilde{L}} l$ under the optimal contract. □

4 Literature Review

This paper contributes to the literature on environments with stochastic availability of opportunities. Möbius (2001) and Hauser and Hopenhayn (2008) study a repeated game in which each player occasionally has an opportunity to grant a favor to his counterpart at a cost to himself. These papers differ from ours in that they analyze a symmetric game in which the availability of favors is privately observed. Moreover, these assumptions lead to efficient equilibria being symmetric and including phases in which favors are denied. Thus, these papers do not derive monotonicity results similar to ours. Samuelson and Stacchetti (2017) characterize the Pareto fortifier of a repeated game where the set of favor types is arbitrary, their availability is publicly observed, and both players must agree to extend a favor. However, their analysis is not informative about the strategies used to obtain these values.

A more closely related paper is our previous work, Bird and Frug (2018), which utilizes a mechanism design approach to analyze a stationary environment in which the agent privately observes the availability of rewards and investment opportunities when monetary transfers are absent. In that paper we characterize the optimal combination of rewards that are allowed at each state, and establish that the agent is compensated via "time allowances" (i.e., he is allowed to enjoy all rewards that arrive in a fixed time interval). Consequently, in contrast to the present paper, the agent's compensation may decrease over time. This qualitative difference is easily explained by the need to incentivize the agent to reveal available investments, a need that is not present when their availability is publicly observed.

A second closely related paper is Forand and Zápal (2018) who study a model, similar to that of Bird and Frug (2018), in which there is symmetric information about action availability but the environment is non-stationary. Their main result, which was derived independently of our work, is twofold. First, they show that there exists an optimal, albeit not unique, mechanism under which once the agent enjoys a certain reward (or forgoes an investment), he does so indefinitely. Second, as we demonstrated in Bird and Frug (2018), they show that the ratio of the player's utility from each action is the key characteristic for determining which actions are allowed in each state.

The main difference between the present paper and the aforementioned papers is of course the generality of our environment that allows us to identify a fundamental common force that appears in many contracting environments. Instead of simplifying a complex environment, we study the dynamic properties of certain components of the the dynamic contract. In addition, our analysis is also applicable to interactions that evolve endogenously.

Our results also resemble those of Ray (2002) who studies the dynamics of optimal contracts where the principal has limited commitment power. In his setting the set of available actions does not change over time; however, the principal's limited commitment induces periodic contracts to shift in the agent's favor as time goes by.

This work is also related to the vast literature that analyzes the dynamics of wage. Harris and Holmström (1982) and Holmström (1983) show that fluctuations in the the worker's outside option generates an increasing wage profile. A similar conclusion is drawn by Burdett and Coles (2003) and Postel-Vinay and Robin (2002) who study wage dynamics in the presence of search frictions. He (2012) shows that if the agent has access to a private savings account, his wage must also increase over time.

Rogerson (1985) and Holmström and Milgrom (1987) focus on the optimal inter-temporal provision of incentives when the agent's action is not observed. They show that in order to solve the moral hazard problem, the agent's wage must decrease if a bad outcome is observed and thus his wage is non-monotone. Hoffmann and Pfeil (2010) further show that with moral hazard random shocks that are beyond the agent's control can also lead to ambiguous changes in the agent's wage. Recent work (e.g., Sannikov, 2008; Garrett and Pavan, 2012) focuses on changes in the worker's productivity and studies the optimal promotion and termination of the agent in addition to the optimal timing of compensation.

5 Conclusion

The main contribution of this paper is to show that in environments with fluctuations in action availability and symmetric information on *both* action availability and payoffs, the agent's compensation (or the level at which any other separable activity is used) is non-decreasing; however, Bird and Frug (2018) show that this is not the case if the agent has private information on action availability. Intuitively, a potential decrease in the agent's compensation can be a useful screening device if the agent, for whose type the contract is intended, believes the wage-decrease occurs at a lower probability than agents of other types do.

When the agent has private information over payoffs, wage reductions may also be a valuable screening device. However, in the special case where the agent (privately) learns his type at the beginning of the interaction and his type does not evolve over time, providing a non-decreasing wage schedule not only minimizes the cost of incentivizing the agent to implement tasks, but also minimizes his information rents. In this special case, after the agent chooses his contract (from the menu of available contracts) he can no longer impact the evolution of his job description. Thus, the agent chooses the contract he is supposed to choose if and only if his payoff from doing so is (weakly) greater than his payoff from choosing another contract and following its terms until any possible stopping time. Therefore, smoothing out a wage decrease by postponing compensation makes choosing the wrong contract less profitable for all agent types.

Appendix A: Proofs

Proof of Proposition 3

Proof. The simplest counterexamples can be obtained if the principal's payoff from the agent misusing the activity is such that any misuse of the activity should be avoided. In particular, setting the principal's payoff from any deviation to $-\kappa(l)-C$ for a sufficiently large C will suffice for the following examples.

Consider the following interaction with no discounting that we use to construct both counterexamples (l', l'' are defined below for each type of violation of condition $\mathcal{D}$).

$t=0$ Agent decides whether to initiate interaction.

$t=1$ Activity is available.

$t=2$ Principal chooses G_1 or B_1, which gives the agent a respective payoff of $-l'$ or $-D(l')$. Both actions give the principal a payoff of zero, but the latter action ends the interaction.

$t=3$ Activity is available.

$t=4$ Principal chooses G_2 or B_2, which gives the agent a respective payoff of $-l''$ or $-D(l'')$. Both actions give the principal a payoff of zero, but the latter action ends the interaction.

$t = 5$ Principal gets a large positive payoff.

Denote the suggested level of activity in period t by l_t.

In an optimal mechanism the principal must incentivize the agent to participate and then incentivize him to select the correct level of l_t while he chooses actions G_1, G_2. As $l \leq D(l)$ it is w.l.o.g. to assume that after the agent misuses the activity the principal chooses B_t. Thus, the IC constraints are

$$\begin{aligned} & l_1 + l_3 - l' - l'' \geq 0 && IC_0 \\ & l_1 + l_3 - l' - l'' \geq D(l_1) - D(l') && IC_1 \\ & l_3 - l'' \geq D(l_3) - D(l'') && IC_3 \end{aligned}$$

We now show that in the optimal contract the activity is used at level $l_1 = l'$ and then $l_2 = l''$. Note that for this contract all ICs are binding.

First we consider the case where the slope of $D(\cdot)$ is negative for some $l \in L$. Due to the continuity of $D(\cdot)$ there exists an interval $L_1 \subset L$ and $c > 0$ such that for all $x'' < x' \in L_1, \quad D(x') - D(x'') < -c(x' - x'')$. Choose $l'' < l' \in L_1$ such that for any $\epsilon \in [0, l' - l'']$,

$$\kappa(l') + \kappa(l'') < \kappa(l' - \epsilon) + \kappa(l'' + \epsilon(1 + c))$$

Such values exist as $\kappa(\cdot)$ is convex, increasing, and continuous and this inequality holds with equality when $l' = l'', c = 0$.

Since $\kappa(\cdot)$ is convex, if there exists a better contract it must have $l_1, l_3 \in (l'', l')$. Note that in this range IC_3 is non-binding. As $l_1 < l'$ and D is decreasing in this range, IC_1 implies IC_0. Thus, if $l_1 = l' - \epsilon$, a necessary condition for the contract to be IC is

$$l_3 \geq l'' + \epsilon(1 + c) \tag{2}$$

Thus, relative to the initial contract, l_3 must be increased by at least $(1 + c)$ times the decrease in l_1. However, by the choice of l'', l', any l_3 that satisfies this will increase the principal's total cost from the activity.

Next, we consider the case where the slope of $D(\cdot)$ is greater than one for some $l \in L$. In this case, there exists an interval $L_2 \subset L$ where $\forall x'' < x' \in L_2 \quad D(x') - D(x'') > x' - x''$. With a slight abuse of notation, denote $[l'', l'] \equiv L_2$.

In this interval $D(\cdot)$ is strictly increasing; thus IC_0 implies IC_1. Since $\kappa(\cdot)$ is increasing, IC_0 holds with equality, and since $\kappa(\cdot)$ is convex, a contract can be more profitable (for the principal) than the one suggested above only if $l_1, l_3 \in L_2$. By construction, $l_3 = l''$ is the only such value that satisfies IC_3. □

Appendix B

A Role of Selected Assumptions

Connected Support for Activity Levels

Our definition of an activity requires that its level have a connected support. This assumption, which at first glance may seem to be a mere simplification, is in fact necessary for showing that the use of an activity increases over time. Consider an infinite interaction with a discount factor of $\delta = \frac{1}{2}$. In the first two periods the agent can exert an effort in the set $\{0, 1, 2\}$ on a task (the principal prefers higher effort) and in each period the principal can provide compensation worth 2 utils to the agent.[7]

Requiring high effort and then low effort while providing the maximal compensation is not IC as the agent's discounted utility in period 1 is $\frac{2}{1-\delta} - 2 - \delta = \frac{3}{2} < 2$. Clearly, this implies that requiring high effort in both periods is not IC either. However, requiring low effort and then high effort is IC, as in both periods the agent's discounted continuation payoff from following the contract equals his payoff from exerting no effort.[8] Therefore, the optimal contract requires the agent's effort to increase over time even though condition $\mathcal{D}$ holds.

Discontinuous $D(l)$

Our definition of an activity assumes that $D(l)$ is a continuous function. Unlike the connected support assumption, this assumption has no impact on our results. If a function is discontinuous it must violate condition $\mathcal{D}$; thus the only question is if a discontinuity in $D(\cdot)$ implies that there exist environments where the use of the activity decreases. We claim that a discontinuity in $D(\cdot)$ leads to either a potential decrease in the use of the activity or the non-existence of the

[7]The agent's optimal deviation is always to take the money and exert no effort.

[8]In the second period the agent's payoff from the contract is $-2 + \frac{2}{1-\delta} = 2$, while in the first period it is $-1 - 2\delta + \frac{2}{1-\delta} = 2$.

optimal contract. To see this, consider the case where $D(\cdot)$ is non-decreasing and has a discontinuity point at l^* (the arguments for a non-increasing function are analogous). If D is not right-continuous at l^* then by a counterexample identical to the one used in Proposition 3 we can find environments in which the use of activity decreases from $l > l^*$ to l^*. If D is not left-continuous at l^* then in the same counterexample the principal will wish to increase l_2 as long as it is strictly less than l^*, leading to the non-existence of the optimal contract.

References

Bird, Daniel, and Alexander Frug. 2018. "Dynamic nonmonetary incentives." Tel Aviv University, Mimeo.

Bolton, Patrick, and Mathias Dewatripont. 2005. *Contract Theory.* MIT Press.

Burdett, Ken, and Melvyn Coles. 2003. "Equilibrium wage-tenure contracts." *Econometrica*, 71(5): 1377–1404.

Edmans, Alex, and Xavier Gabaix. 2016. "Executive compensation: A modern primer." *Journal of Economic Literature*, 54(4): 1232–1287.

Forand, Jean Guillaume, and Jan Zápal. 2018. "Production priorities in dynamic relationships." Mimeo.

Garrett, Daniel F., and Alessandro Pavan. 2012. "Managerial turnover in a changing world." *Journal of Political Economy*, 120(5): 879–925.

Harris, Milton, and Bengt Holmström. 1982. "A theory of wage dynamics." *Review of Economic Studies*, 49(3): 315–333.

Hauser, Christine, and Hugo Hopenhayn. 2008. "Trading favors: Optimal exchange and forgiveness." *Collegio Carlo Alberto*, Working paper 88.

He, Zhiguo. 2012. "Dynamic compensation contracts with private savings." *Review of Financial Studies*, 25(5): 1494–1549.

Hoffmann, Florian, and Sebastian Pfeil. 2010. "Reward for luck in a dynamic agency model." *Review of Financial Studies*, 23(9): 3329–3345.

Holmström, Bengt. 1983. "Equilibrium long-term labor contracts." *Quarterly Journal of Economics*, 98: 23–54.

Holmström, Bengt, and Paul Milgrom. 1987. "Aggregation and linearity in the provision of intertemporal incentives." *Econometrica*, 55(2): 303–328.

Möbius, Markus. 2001. "Trading favors." Harvard University Mimeo.

Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Equilibrium wage dispersion with worker and employer heterogeneity." *Econometrica*, 70(6): 2295–2350.

Ray, Debraj. 2002. "The time structure of self-enforcing agreements." *Econometrica*, 70(2): 547–582.

Rogerson, William P. 1985. "Repeated moral hazard." *Econometrica*, 53(1): 69–76.

Rotemberg, Julio J., and Garth Saloner. 1986. "A supergame-theoretic model of price wars during booms." *American Economic Review*, 76(3): 390–407.

Samuelson, Larry, and Ennio Stacchetti. 2017. "Even up: Maintaining relationships." *Journal of Economic Theory*, 169: 170–217.

Sannikov, Yuliy. 2008. "A continuous-time version of the principal-agent problem." *The Review of Economic Studies*, 75(3): 957–984.