On-call jobs: Contracts with lumpy effort

Daniel Bird[*] Alexander Frug[†]

May 14, 2021

Abstract

An agent performs randomly arriving tasks that are rare and "difficult," and a principal may provide frequent but low-value compensation. We characterize the optimal contract for both the case where the principal observes the arrival of opportunities and the case where the agent may conceal them from the principal. The optimal contracts reveal appealing qualitative features that vary with the assumed information structure and the players' discount rates. Unless the players use the same discount rate and opportunities are observable, the players' ability to realize gains from interaction decreases when the degree of "lumpiness" in opportunities increases.

Keywords: Lumpy effort, Random opportunities, Timing of compensation.
JEL Classifications: D82, D86.

[*]Eitan Berglas School of Economics, Tel Aviv University (e-mail: dbird@tauex.tau.ac.il).
[†]Department of Economics and Business, Universitat Pompeu Fabra and Barcelona GSE (e-mail: alexander.frug@upf.edu).
Acknowledgments to be added.

1 Introduction

Randomly arriving peaks that require concentrated effort are a feature of many professions and are common in many workplaces. Various on-call jobs and positions with irregular shifts (e.g., technicians, medical staff, firefighters, salespeople, manufacturers in industries with on-demand production) are prominent examples.[1] In some cases, workers receive immediate compensation whenever they exert high effort. Yet, due to institutional constraints (e.g., general salary regulations or limited periodic budget of a team manager) or preferences (efficient planning, dynamic self-control, etc.), fully compensating the worker *on the spot* after every instance of concentrated effort may be infeasible or inefficient. In this paper, we study the implications of the disparity between the lumpy nature of work requirements and the more evenly spread form of compensation provision.

We consider a dynamic interaction between an agent who performs *rare but difficult* tasks and a principal who is able to provide *frequent but low-value* compensation. A special feature in our model is that the provision of compensation is a time-consuming process. If both players could commit, it would hardly play any role: lumps of work could easily be counterbalanced by continual periods of compensation and, so long as the agent receives enough compensation ex ante, the lumpy nature of the work would be inconsequential. However, this is no longer the case when only the principal has commitment power: paying the agent in advance of his effort is not feasible if he is free to walk away at any moment. This constraint creates a dynamic spillover between different lumps of work and affects the structure of optimal contracts.

We set up a continuous-time infinite-horizon model with the following key components. Non-storable opportunities arrive according to a Poisson process. When an opportunity is available, the agent chooses an effort in $[0, 1]$, which we assume throughout is perfectly observable. Whenever the agent exerts positive effort, it generates an immediate benefit to the principal and a cost to the agent—both of which are linear in the agent's effort. The principal, on the other hand, chooses at every point in time (irrespective of opportunity arrivals) a *flow* wage in

[1] See Mas and Pallais (2020) for a recent review on "alternative work arrangements."

$[0, 1]$. Our choice to model effort differently from wage—i.e., as discrete Poisson events vs. continuous ones—is essential for capturing the randomness as well as the relative size of *lumpy* opportunities. The resulting asynchronicity of compensation and effort makes the players' discount rates key determinants. We do not restrict attention to the case where the players discount future payoffs identically but characterize optimal contracts for any pair of players' discount rates. While assuming identical discount rates is common, it can be misleading. As we will show, the lumpiness of opportunities is detrimental unless the players use precisely the same discount rate. In that case, if the principal observes the arrival of opportunities, then the degree of lumpiness is inconsequential.

First, we consider the case where the arrival of opportunities is observed by the principal.[2] For instance, consider a technical support center that receives all customer calls and allocates jobs to technicians. Also, it seems plausible to assume that the arrival of patients to hospitals or emergency calls to fire stations can be observed by the relevant supervisors. The structure of the optimal contract depends on the players' relative patience.

When the principal is *patient* relative to the agent, compensation begins immediately after the agent exerts effort. Following the agent's work on a given opportunity, the principal promises a fixed periodic payment for a certain amount of time; however, this promise is nullified upon the arrival of the next opportunity in order to make room for a new similar "conditional promise." Thus, in the optimal contract, the agent receives a fixed periodic wage so long as he has performed some work in a recent time span, after which his wage drops to zero until the next opportunity arrives. While compensation is responsive to work, the fact that the principal's compensation promises do not accumulate implies that there is only a partial correlation between the total amounts of effort and compensation.

When the principal is *impatient* relative to the agent, initially the agent works (whenever an opportunity arrives) but does not receive compensation. At some point, the contract moves to a new phase in which the agent works and receives a periodic wage. Finally, the contract moves to a third, absorbing phase, in which the agent does not work but enjoys a periodic wage indefinitely. The transition times

[2] Whether or not the opportunities are directly observed by the agent is immaterial.

between the phases are fully determined by the arrival time of the first opportunity. Hence, there is a very low correlation between the total wage paid to the agent and the total amount of effort he exerts.

Next we analyze the case where the availability of opportunities is privately observed by the agent. For instance, when a prospective buyer arrives, the manager of a car dealership can forgo the sales opportunity (and conceal it from the company headquarters) simply by not being obliging. Also, the specification may fit an adapted version of the visiting-technician story where, instead of receiving jobs from the general support center, technicians are directly contacted by potential customers.

When opportunities can be concealed by the agent, the contract must incentivize the agent not only to exert the desired effort on opportunities that are known to be available, but also to reveal their availability. We show that, regardless of the players' relative patience, the optimal contract is fully characterized by two thresholds over the agent's continuation utility: (1) compensation provision is postponed so long as the current debt to the agent is below u^W, and (2) effort is exerted on available opportunities so long as the current debt to the agent is below u^O. In contrast to the case where the principal observes the arrival of opportunities, now there is perfect bookkeeping of the agent's realized work. That is, the overall compensation provided to the agent exactly equals his total effort (both of which are discounted according to the agent's discount factor). Perhaps surprisingly, even when the principal is impatient relative to the agent, unless he is "extremely" impatient, he front-loads at least some (and sometimes all) of the agent's compensation.

A main comparative statics result in our context relates to the degree of *lumpiness* of opportunities. Although making opportunities lumpier (i.e., larger but rarer) has different effects on the optimal contract in the various specifications that we consider, we find that in all but the case where $r_a = r_p$ and opportunities are publicly observable, lumpier opportunities reduce the value from the interaction. This finding is consistent with recent empirical evidence on the detrimental effect of lumpy work requirements.[3] For example, Mas and Pallais (2017) find that the

[3]This traditional view is also prevalent in sociology; see, e.g., Kalleberg (2011).

average worker in their experiment is willing to give up 20 percent of his wages to avoid a schedule set by an employer on short notice.

Related Literature

Our paper complements works that study the optimal timing of compensation (e.g., Salop and Salop 1976; Lazear 1981; Carmichael 1983) and, in particular, those that analyze the mixture between short- and long-term incentives in settings where information changes over time (e.g., Sannikov 2008; Garrett and Pavan 2012, 2015); see Edmans and Gabaix (2009) and Pavan (2017) for a review of the literature. We contribute to this literature by analyzing the effect of the stochastic arrival of opportunities—a friction that has been overlooked by previous work.

Our work also contributes to the recent literature that studies optimal contracting under different discount factors. Opp and Zhu (2015) study relational contracting in a repeated moral hazard setting, Frankel (2016) studies dynamic delegation, Hoffmann, Inderst and Opp (2020) study a one-shot moral hazard problem in which there is delay in the arrival of information, and Krasikov, Lamba and Schacherer (2020) analyze a canonical adverse selection problem.

Finally, our work is related to the growing literature that studies interactions with randomly arriving favors/projects. One strand of this literature, e.g., Möbius (2001), Hauser and Hopenhayn (2008), and Samuelson and Stacchetti (2017), analyzes the optimal exchange of favors. Specifically, these papers characterize efficient equilibria in repeated games where, occasionally, the players have opportunities to grant favors to their counterparts.

The other strand of this literature studies principal-agent problems. Forand and Zápal (2020) and Bird and Frug (2020) derive the optimal contract under symmetric information: Forand and Zápal (2020) study a model with no transfers in which projects of different types arrive randomly over time, whereas Bird and Frug (2020) study a canonical employment model in which the agent's productivity of effort varies over time. Li, Matouschek and Powell (2017), Bird and Frug (2019), and Lipnowski and Ramos (2020) consider transfer-free environments with asymmetric information. More specifically, Li, Matouschek and Powell (2017) derive the optimal relational contract when the agent has private information on project availability. Bird and Frug (2019) derive the optimal contract under full commitment in a setting where projects of different types arrive randomly over time, and the project's type is observed privately by the agent. Lipnowski and Ramos (2020) characterize efficient equilibria when the agent has private information on project payoffs. We contribute to this strand of the literature by studying the consequences of lumpy work requirements and highlighting the impact of the players' relative patience and the informational structure on the qualitative properties of the optimal contract. In particular, we establish that lumpiness generally reduces the potential gains from the interaction, a result that would be obscured under the standard assumption in the literature that players share the same discount factor.

The paper proceeds as follows. In Section 2 we present the model. In Sections 3 and 4 we analyze, respectively, the case where opportunities are observable and concealable. We offer concluding remarks in Section 5. All proofs are relegated to the Appendix.

2 Model

We consider an infinite-horizon continuous-time contracting problem in which *opportunities* arrive stochastically over time according to a Poisson process with arrival rate $\mu > 0$. The no-effort action, $\alpha = 0$, is always available to the agent. When an opportunity arrives, and only then, in addition to the no-effort action, the agent can exert effort $\alpha \in (0, 1]$. The agent's effort $\alpha \in [0, 1]$ induces an immediate benefit of $\alpha \cdot B$ to the principal and a cost of $\alpha \cdot C$ to the agent, where $B > C > 0$. At each instant, the principal chooses a flow wage $w \in [0, 1]$. We assume that both the agent's marginal utility from wage and the principal's marginal cost of wage are constant and equal to 1.

The players maximize expected discounted payoffs. We denote the agent's discount factor by r_a and focus on the case where there is no fundamental shortage of incentives. That is, we assume that the agent's discounted payoff from setting $w = 1$ indefinitely exceeds his expected discounted cost of full-intensity work on all opportunities that arrive, even if one is currently available. Formally, we assume that[4]

$$C + \frac{\mu C}{r_a} < \frac{1}{r_a}.$$

We denote the principal's discount factor by r_p and refer to the principal as *patient* if $r_p \leq r_a$ and as *impatient* if $r_p > r_a$.

Throughout the paper we assume that the agent's effort is perfectly observed by the principal, but we vary our assumptions about whether or not the principal observes the arrival of opportunities. If the principal does observe the arrival of opportunities ("*observable opportunities*"), then a public history h_t specifies for every $s < t$ whether or not an opportunity was available and the agent's choice of effort.[5] If the principal does not observe the arrival of opportunities ("*concealable opportunities*"), then a public history h_t contains only the agent's effort choices. Given the Poisson arrival of opportunities, any private information that the agent has about the availability of opportunities in the past is irrelevant, and so there is no need to keep track of his private information. Hence, to reduce notation and terminology we refer to a public history as a history under both information structures we consider. We denote the set of all histories of length t by H_t and the set of all finite histories by $H = \cup_{t \in \mathbb{R}_+} H_t$.

At the beginning of the interaction, the principal specifies a *work schedule*

$$\alpha : H \to [0, 1],$$

which assigns a required effort to every history should an opportunity arrive at that history, and he commits to a *compensation policy*

$$w : H \to [0, 1],$$

which maps histories into a flow wage. A pair $\langle \alpha(\cdot), w(\cdot) \rangle$ is referred to as a *contract*. Without loss of generality, we assume that after the principal *detects* a deviation from the specified work schedule, he pays a wage of zero indefinitely.[6] We

[4]Allowing for the opposite inequality would add trivial cases with corner solutions that would not add much of substance but would needlessly impede the exposition.

[5]In this case, whether or not the agent observes the arrival of opportunities is immaterial.

[6]The set of detectable deviations depends on the information structure. When opportunities are

say that the contract $\langle \alpha(\cdot), w(\cdot) \rangle$ is *measurable* if, at every history, the agent's continuation utility and the principal's continuation value are well defined. That is, the expectations

$$\mathbb{E}\left(\int_{s=t}^{\infty} e^{-r_a(s-t)}\left(w(h_s) - \mu C\alpha(h_s)\right) ds | h_t\right),$$
$$\mathbb{E}\left(\int_{s=t}^{\infty} e^{-r_p(s-t)}\left(\mu B\alpha(h_s) - w(h_s)\right) ds | h_t\right)$$

exist for every $h_t \in H$.

We say that the contract $\langle \alpha(\cdot), w(\cdot) \rangle$ is *incentive compatible* if it is measurable and, for every $h_t \in H$, it is optimal for the agent to choose $\alpha = \alpha(h_t)$ (conditional on the availability of an opportunity), given the continuation of the contract. Note that if a deviation to a positive effort level is profitable at a given history, then so is a deviation to no effort. Hence, a contract is incentive compatible if and only if the agent (weakly) prefers to follow $\alpha(\cdot)$ than to deviate to $\alpha = 0$ from some point onward.[7] Since the agent can guarantee himself a payoff of zero by never exerting effort, there is no need to impose an explicit individual rationality constraint. We restrict attention to incentive-compatible contracts.

In our analysis and discussion of the results, we often compare contracts in terms of the timing of effort/compensation. We use the following relations. A work schedule $\alpha(\cdot)$ *postpones* effort relative to $\alpha'(\cdot)$ at a history h_t if, for all $\tau > t$,

$$\mathbb{E}\left(\int_{s=t}^{\tau} e^{-r_a(s-t)}\alpha(h_s) | h_t\right) \leq \mathbb{E}\left(\int_{s=t}^{\tau} e^{-r_a(s-t)}\alpha'(h_s) | h_t\right) \tag{1}$$

with an equality for $\tau = \infty$ and a strict inequality for some τ. Similarly, a compensation policy $w(\cdot)$ *postpones* compensation relative to $w'(\cdot)$ at h_t if, for all $\tau > t$,

$$\mathbb{E}\left(\int_{s=t}^{\tau} e^{-r_a(s-t)}w(h_s) | h_t\right) \leq \mathbb{E}\left(\int_{s=t}^{\tau} e^{-r_a(s-t)}w'(h_s) | h_t\right) \tag{2}$$

observable any deviation by the agent is detected, whereas when opportunities are concealable a deviation by the agent is detected only if he exerted a strictly positive effort other than $\alpha(h_t)$ on an available opportunity.

[7]We state the incentive compatibility constraints formally in the following sections.

with an equality for $\tau = \infty$ and a strict inequality for some τ. Analogous definitions for *expediting* effort and compensation are obtained by reversing the inequalities in (1) and (2).

Note that the above definitions use the agent's discount factor. In addition, observe that the principal-discounted marginal cost of providing the agent with an agent-discounted util at time t is $e^{-r_p t} \cdot \frac{1}{e^{-r_a t}} = e^{(r_a - r_p)t}$ and, similarly, the principal-discounted marginal benefit from the agent exerting one agent-discounted util at time t is $e^{-r_p t} \frac{B}{C} \frac{1}{e^{-r_a t}} = \frac{B}{C} e^{(r_a - r_p)t}$. Whether these expressions increase or decrease in t is fully pinned down by whether the principal is patient or impatient. The following observation is implied.

Observation 1.

1. *Expediting compensation and postponing effort are both profitable for a strictly patient principal.*

2. *Postponing compensation and expediting effort are both profitable for an impatient principal.*

3 Observable Opportunities

First, we consider the case where the principal observes the random arrival of opportunities. In this case, any deviation by the agent is detected, and since the agent can guarantee himself a continuation payoff of zero at any point in time, the incentive compatibility constraints are given by

$$-C\alpha(h_t) + \mathbb{E}\left(\int_{s=t}^{\infty} e^{-r_a(s-t)} \left(w(h_s) - \mu C \alpha(h_s)\right) ds | (h_t, O, \alpha(h_t))\right) \geq 0 \quad \forall h_t \in H, \tag{IC_{pub}}$$

where $(h_t, O, \alpha(h_t))$ is the event in which, before time t, play proceeds according to h_t, and, at time t, an opportunity arrives and the agent exerts an effort of $\alpha(h_t)$ on that opportunity. The structure of the optimal contract and its qualitative properties depend on the players' relative patience.

Patient principal: ***"Have you done anything for me lately?"***

When $r_p < r_a$, increasing the (average) lag between compensation and work in a manner that keeps the agent indifferent reduces the principal's profit. It is therefore easy to see that to obtain the maximal profit from the *first* opportunity that arrives the principal would have to pay the agent the maximal wage $w = 1$ immediately after his work, for an interval of time that is just long enough to compensate him for the cost of effort. However, were the principal to do so, it would not be possible for him to provide immediate compensation for any additional opportunities that might arise within that time interval. Hence, the cheapest form of compensation for the very first opportunity reduces the potential profit from further opportunities. In fact, the patient principal faces a complex optimization problem where he endeavors to provide timely compensation for the agent's effort on each of the randomly arriving opportunities. The main result of this section shows that the solution to this problem is simple and qualitatively appealing: under the optimal contract, the agent exerts the same effort α^* on all opportunities, and he receives a flow wage $w = 1$ if an opportunity has been available in the last S^* units of time.

This optimal compensation policy can be colloquially described as "have you done anything for me lately?"; i.e., the compensation is fully determined by whether or not work was performed by the agent within a recent time span (of fixed length) but does not depend on *how much* work was performed. A useful alternative interpretation of this form of compensation is that of "conditional promises"; i.e., following the agent's work on a given opportunity, the principal promises a fixed periodic payment for a given time interval, but this promise is nullified upon the arrival of the next opportunity. The complete nullification of the principal's obligations to the agent upon the arrival of a new opportunity frees incentivization resources precisely when they are needed, and enables the principal to incentivize effort on the currently available opportunity via a new conditional promise.

Our assumption that $C + \frac{\mu C}{r_a} < \frac{1}{r_a}$ implies that full implementation of all opportunities can be attained in an incentive-compatible contract. However, this need not be optimal for the principal. To see why this is the case, recall that the principal's cost of providing the agent with one agent-discounted util t units of time in the future is $e^{(r_a - r_p)t}$. As the principal's profit from a util worth of effort exerted by the agent is $\frac{B}{C}$, it follows that the *maximal profitable lag* between compensation and work is T^*, where T^* is defined implicitly by

$$e^{T^*(r_a - r_p)} = \frac{B}{C}$$

if $r_p < r_a$, and is given by $T^* = \infty$ if $r_a = r_p$. If a conditional promise of length T^* is insufficient to incentivize the agent to exert full effort, then $\alpha^* < 1$; i.e., the principal instructs the agent to forgo part of each opportunity.

To formally characterize the optimal contract, let $\sigma_{-1}(h_t)$ denote the supremum of opportunity arrival times along the history h_t, and let $\sigma_{-1}(h_t) = -\infty$ for histories in which no opportunity arrives. In all subsequent results, equalities and uniqueness statements should be interpreted as holding almost surely.

Proposition 1. *Assume that the principal observes the arrival of opportunities. If the principal is patient, then there exist $\alpha^* \in (0, 1]$ and $S^* \in (0, T^*]$ such that*

$$\alpha(h_t) = \alpha^* \quad ; \quad w(h_t) = \begin{cases} 1 & \text{if } t - \sigma_{-1}(h_t) \leq S^* \\ 0 & \text{if } t - \sigma_{-1}(h_t) > S^* \end{cases}$$

is an optimal contract. Moreover, this is the unique optimal contract if $r_p < r_a$.

When $r_a = r_p$ there may be multiple optimal contracts. Intuitively, in this case, postponing compensation does not alter the principal's profit or violate any of the agent's incentive compatibility constraints. Thus, any contract that results from postponing compensation relative to the optimal contract characterized in Proposition 1 is also an optimal contract.[8]

To study the impact of lumpiness, we compare settings where the frequency and magnitude of opportunities vary. Formally, we say that the opportunities represented by the parameters (C_1, B_1, μ_1) are lumpier than those represented by (C_0, B_0, μ_0) if $C_1 = \lambda C_0$, $B_1 = \lambda B_0$, and $\mu_1 = \frac{\mu_0}{\lambda}$ for some[9] $\lambda > 1$.

[8] The assumption that there is no fundamental shortage of resources implies that if $r_a = r_p$, then under the optimal contract $\alpha^* = 1$ and $S^* < \infty$. Hence, postponing compensation is feasible.

[9] Analogously, one could define *smoother* opportunities by considering $\lambda \in (0, 1)$.

A lumpier opportunity is "bigger" ($C_1 > C_0, B_1 > B_0$), and hence, when available, it requires a longer conditional promise to compensate the agent for his greater effort. On the other hand, since a lumpier opportunity is also "rarer" ($\mu_1 < \mu_0$), a conditional promise of any given length is likely to last longer, and so the value of any such promise is higher from the perspective of the agent. In the proof of the next result, we show that the second effect only partially compensates for the increase in the agent's cost of exerting full effort when opportunities are lumpier and hence, in general, lumpier opportunities require longer conditional promises.

The observations that a conditional promise of any given length lasts (on average) longer when opportunities become lumpier, and that lumpier opportunities require longer conditional promises, jointly imply that the lag between effort and compensation increases when opportunities become lumpier. Thus, a strictly patient principal is made strictly worse off when opportunities become lumpier (note that these changes have no impact if $r_a = r_p$). In fact, when opportunities are sufficiently lumpy, the length of the required conditional promise to incentivize full effort exceeds the maximal profitable lag T^*. Hence, if opportunities are sufficiently lumpy, it is suboptimal for a strictly patient principal to incentivize full effort. The following proposition formalizes the above discussion. We address the case where $r_a = r_p$, after we analyze the case of an impatient principal.

Proposition 2. *Assume that opportunities are observable and that $r_p < r_a$. When opportunities become lumpier, the principal's value strictly decreases, and the agent's effort on each opportunity weakly decreases.*

In the "smooth limit" of our model, i.e., the limit specification $(\lambda C, \lambda B, \frac{\mu}{\lambda})$ for $\lambda \to 0$, the optimal contract converges to a flow of spot contracts: at every instant, the agent provides a flow effort of 1 and receives compensation on the spot. Since compensation is immediate, so long as the principal is more patient than the agent, the optimal contract does not depend on the exact values of the players' discount rates. On the other hand, the optimal contract away from the smooth limit does depend on the exact values of the players' discount rates because effort and compensation become asynchronous. Recall that under the optimal contract, the agent exerts the maximal effort for which he can be compensated either before the next opportunity arrives or before T^* units of time have passed. Since T^* is increasing in r_p, i.e., decreases when the principal becomes more patient, it follows that the agent's effort α^* weakly decreases when the principal becomes more patient (holding the agent's discount rate fixed).[10]

Impatient principal: *"(Almost) automatic promotion"*

By Observation 1, postponing compensation and expediting effort (according to the agent's discount factor) are both profitable for an impatient principal. The key observation underlying the characterization in this section is that neither postponing compensation nor expediting effort violates incentive compatibility. Hence, the optimal contract should front-load effort and back-load compensation.

To understand the structure of the optimal contract in more detail, fix a sequence of times $\Sigma = \{\sigma_j\}_{j\in\mathbb{N}}$, where $\sigma_1 = 0$, and suppose that it is commonly known that opportunities arrive according to Σ. From the discussion above, it is immediate that, under the optimal contract, $w = 1$ from some point τ^W_Σ onward, and the agent works on all opportunities that arrive before some τ^O_Σ. Moreover, on all but, perhaps, the very last opportunity the agent exerts maximal effort, and his time-zero discounted expected payoff from the contract is zero.

We now consider the hypothetical setting where opportunities arrive randomly but the whole sequence Σ is realized and revealed to the players at the beginning of the interaction. If the contract is proposed after the realized Σ becomes publicly known, each realized sequence induces a different pair $(\tau^O_\Sigma, \tau^W_\Sigma)$. On the other hand, if the players sign a binding contract prior to observing the realization of Σ, the principal strictly benefits (ex ante) from averaging out the variability in τ^O_Σ and[11] τ^W_Σ. However, a contract where τ^O and τ^W do not depend on the realization of Σ would require the agent's commitment power.

Assume now that Σ is revealed gradually over time (as in our actual contracting problem). This enables the principal to benefit from offering nonstochastic terms

[10] In fact, it can be shown that α^* strictly decreases if $S^* = T^*$.

[11] For example, variability in τ^W_Σ can be viewed as a failure to fully postpone compensation since compensation begins earlier in some paths of play than in others.

(τ^O, τ^W)—and thus fully expediting effort and fully postponing compensation—without relying on the agent's commitment power. This stems from two simple properties of the environment and the suggested contract. First, the Poisson arrival of opportunities guarantees that the agent doesn't learn useful information about the future from past arrivals; and second, because compensation is back-loaded, the only binding incentive compatibility constraint is at σ_1, which we assumed to be at time 0, namely,

$$C + \int_0^{\tau^O} \mu C e^{-r_a t} dt = \int_{\tau^W}^{\infty} e^{-r_a t} dt. \tag{3}$$

The relation between τ^O and τ^W can easily be identified by local considerations: by marginally increasing τ^O and decreasing τ^W, the principal can incentivize the agent to exert effort for which he will be compensated once $\tau^O - \tau^W$ units of time have passed. Thus, in optimum, $\tau^O - \tau^W$ is such that the net surplus generated by effort is offset by the principal's relative impatience over $\tau^O - \tau^W$ units of time,

$$\frac{B}{C} = e^{(r_p - r_a)(\tau^O - \tau^W)}, \tag{4}$$

which, in particular, implies that $\tau^W < \tau^O$ for any $r_p > r_a$.

In a setting that begins with an opportunity, the agent's lack of commitment doesn't affect the optimal contract. Our original problem, however, does not begin with an available opportunity. If the agent could commit, the optimal contract for the principal would still leave zero surplus to the agent and have the same threshold structure (though the thresholds would be different). In this contract, however, the agent's continuation utility drops below zero if the first opportunity arrives sufficiently early. Hence, such a contract cannot be the solution when the agent lacks commitment power, and so some stochasticity must remain in the contract. This, however, takes an extremely intuitive form: prior to the arrival of the first opportunity, the players just wait, and upon arrival of that opportunity, they set the clock to zero and use the aforementioned contract (τ^O, τ^W).

To characterize the optimal contract formally, let $\sigma_1(h_t)$ denote the infimum of the arrival times of opportunities along the history h_t, and let $\sigma_1(h_t) = \infty$ for histories in which no opportunity arrived.

Proposition 3. *Assume that the principal observes the arrival of opportunities. If this principal is impatient, then the unique optimal contract is*

$$\alpha(h_t) = \begin{cases} 1 \text{ if } t \leq \sigma_1(h_t) + \tau^O \\ 0 \text{ if } t > \sigma_1(h_t) + \tau^O \end{cases} \quad \text{and } w(h_t) = \begin{cases} 0 \text{ if } t \leq \sigma_1(h_t) + \tau^W \\ 1 \text{ if } t > \sigma_1(h_t) + \tau^W \end{cases} ,$$

where τ^W, τ^O are the unique solution to (3) *and* (4)*.*

Qualitatively, the optimal contract consists of three phases. Initially, the agent exerts effort but does not receive a wage; at some point, he begins to receive compensation but still has to work whenever an opportunity arrives; and finally, he is promoted to a position where he receives a wage but does not exert further effort. From the ex-ante perspective, the promotion dates are random, but all the uncertainty is resolved when the first opportunity arrives. Hence, as in the case where the principal is patient, there is very little correlation between the total effort and wage throughout the interaction.

Comparative statics with respect to both the principal's discount factor and the lumpiness of opportunities are readily available. First, as the principal becomes more impatient the agent's expected effort decreases. To see this, note that the optimality condition (4) can be written as

$$\frac{B}{C} e^{(r_a - r_p)\tau^O} = e^{(r_a - r_p)\tau^W}, \tag{5}$$

and since $\tau^O > \tau^W$, the derivative, with respect to r_p, of the LHS of (5) is less than that of the RHS. As (3) implies that τ^W is decreasing in τ^O, it follows that increasing r_p will lead to a decrease in τ^O (and an increase in τ^W).

Regarding the effect of lumpiness, on the one hand, when opportunities are larger the agent exerts more effort on the first opportunity to arrive. By (4), the middle phase (when the agent receives a wage and exerts effort) is not altered by the degree of lumpiness. Therefore, to compensate the agent for exerting more effort on the first opportunity to arrive, the principal will begin providing a wage earlier. Thus, from the arrival time of the first opportunity, the agent's expected effort increases when opportunities become lumpier. On the other hand, it will take more time for the first opportunity to arrive when opportunities become rarer. It turns out that the latter effect is stronger, and so when opportunities become lumpier the agent's expected discounted effort and, accordingly, the principal's value decrease.

Proposition 4. *Assume that opportunities are observable and that $r_p > r_a$. The principal's value strictly decreases when opportunities become lumpier.*

Propositions 2 and 4 jointly establish that making opportunities lumpier is detrimental for the principal whenever $r_a \neq r_p$. However, if $r_a = r_p$, then making opportunities lumpier does not impact the principal's value so long as the assumption that $C + \mu \frac{C}{r_a} < \frac{1}{r_a}$ continues to hold. This follows from three observations. First, since $r_a = r_p$ the assumption that $C + \mu \frac{C}{r_a} < \frac{1}{r_a}$ implies that under the contract characterized in Proposition 1 the agent exerts full effort on all opportunities. Second, the agent's expected utility from that contract is zero. Finally, $r_a = r_p$ implies that the timing of compensation does not affect the principal's cost of providing compensation.

4 Concealable Opportunities

In the optimal contracts derived in the previous section, the arrival of opportunities typically leads to an immediate decrease in the agent's continuation utility. In settings where the arrival of an opportunity is observed only by the agent, such contracts provide incentives for the agent to conceal opportunities from the principal. In order to provide incentives for the agent to reveal when opportunities become available, the arrival of opportunities must never be "bad news" for the agent. Hence, the incentive compatibility constraints when opportunities are concealable are

$$-C\alpha(h_t) + \mathbb{E}\left(\int_{s=t}^{\infty} e^{-r_a(s-t)} \left(w(h_s) - \mu C\alpha(h_s)\right) ds | (h_t, O, \alpha(h_t))\right) \geq$$

$$\mathbb{E}\left(\int_{s=t}^{\infty} e^{-r_a(s-t)}\left(w(h_s)-\mu C\alpha(h_s)\right) ds | (h_t, N)\right) \quad \forall h_t \in H, \qquad (IC_{priv})$$

where (h_t, N) denotes the event in which, before time t, play proceeds according to h_t, and, at time t, and an opportunity does not arrive; and $(h_t, O, \alpha(h_t))$, as in the case of observable opportunities, denotes the event in which before time t, play proceeds according to h_t, and, at time t, an opportunity arrives and the agent exerts an effort of $\alpha(h_t)$ on that opportunity.

It is well known in the dynamic contracting literature that if the environment is stationary, then the agent's continuation utility can be used as a state variable for deriving the optimal contract (see Spear and Srivastava 1987 and Thomas and Worrall 1990). The argument behind this useful result relies on the property that the continuation payoffs of an efficient contract must always lie on the constrained Pareto frontier. This, in turn, follows from two simple observations: first, if the agent receives a continuation utility u via an inefficient continuation contract, then the principal can increase his value by replacing that continuation contract with a different contract that provides the agent with the same agent-discounted continuation utility u; second, since the agent is indifferent between the original and modified continuations of the contract, this change has no impact on earlier incentive compatibility constraints. Notice that these observations do not depend on the assumption that the players share the same discount factor and, thus, they are valid in our setting where the players use different discount factors.

We denote, respectively, by $\alpha(u)$, $w(u)$, and $V(u)$ the Markovian work schedule, the Markovian compensation policy, and the principal's value as a function of the agent's continuation utility. Note that $u \in [0, \frac{1}{r_a}]$ as, from any point in time onward, the agent can guarantee himself a nonnegative payoff by exerting no effort, and the agent's value from receiving the maximal wage indefinitely is $\int_0^\infty 1 \cdot e^{-r_a t} dt = \frac{1}{r_a}$.

Lemma 1. *$V(u)$ is strictly decreasing and weakly concave.*

Lemma 1 has two important consequences. First, it directly implies that the agent's expected utility from an optimal contract is zero. Second, in contrast to the case where opportunities were observable, when opportunities are concealable the incentive compatibility constraint at every history is binding (regardless of the relative patience of the players).

Corollary 1. *Assume that opportunities are concealable. Under an optimal contract:*

1. *The agent's expected utility is zero.*

2. *All the incentive compatibility constraints are binding.*

When opportunities are concealable, it is convenient to describe the optimal contract via its Markovian representation. By Corollary 1, the agent's continuation utility at the beginning of the interaction is zero and after he exerts effort $\alpha(u)$ his continuation utility increases by $\alpha(u)c$. The drift in the agent's continuation utility while no opportunities arrive is

$$du = r_a u - w(u). \tag{6}$$

Hence, the optimal contract is characterized by the solution of the following HJB equation:

$$\sup_{w(u),\alpha(u)\in[0,1]} \{-r_p V(u) + V'(u)[r_a u - w(u)] - w(u) \\ +\mu\left(\alpha(u)B + V\left(u + C\alpha(u)\right) - V(u)\right)\} = 0, \quad (HJB)$$

subject to (6), where $V'(u)$ exists almost everywhere since $V(\cdot)$ is concave (Lemma 1). The following is the main result of this section.

Proposition 5. *When opportunities are concealable, the optimal contract is generically unique. Moreover, there exist thresholds $u^W, u^O \in [0, \frac{1}{r_a}]$, such that the optimal contract is given by*

$$\alpha(u) = \min\{1, \frac{u^O - u}{C}\} \quad ; \quad w(u) = \begin{cases} 1 & \text{if } u > u^W \\ r_a u^W & \text{if } u = u^W \\ 0 & \text{if } u < u^W \end{cases}.$$

A distinctive feature of the optimal contract when opportunities are concealable is the perfect bookkeeping of the agent's work. Path by path, wage and effort discounted according to the agent's discount rate are equal. Such bookkeeping implies that, as opportunities become lumpier, the asynchronicity of compensation and effort increases, which, in turn, reduces the principal's ability to use compensation resources effectively, regardless of whether he is more or less patient than the agent. Hence, again, lumpier opportunities are detrimental for the principal.

Proposition 6. *Assume that opportunities are concealable. The principal's value strictly decreases when opportunities become lumpier.*

4.1 The Dynamics of Optimal Contracts

The threshold u^O dictates the dynamics of the agent's work. The threshold value $u^O = \frac{1}{r_a}$ corresponds to a policy where the principal instructs the agent to fully exploit every opportunity until all of his compensation budget is exhausted. For lower values of u^O, the principal will initially instruct the agent to fully exploit every opportunity that arrives; however, once multiple opportunities arrive in quick succession, the principal will temporarily instruct the agent to exert low effort. Eventually, there will be a sufficiently long time span in which very few opportunities arrive, and following this event the principal will again instruct the agent to fully exploit opportunities. The optimal threshold depends on the relative patience of both players.

Proposition 7. *Fix B, C, μ, and r_a. The threshold u^O is increasing in r_p. Moreover, if $r_p < r_a$, then $u^O \in (0, \frac{1}{r_a})$, and whenever $r_p \geq r_a$, $u^O = \frac{1}{r_a}$.*

The dynamics of compensation is slightly more nuanced. The level of the compensation threshold u^W captures the degree of back/front-loading of compensation. So long as the agent's continuation utility is below u^W, compensation is deferred to the future. Setting the compensation threshold at the maximal possible value, $u^w = \frac{1}{r_a}$, corresponds to full back-loading: when the agent's continuation utility reaches that level, an indefinite payment of the maximal flow wage, $w = 1$, is necessary to provide the agent with his promised continuation utility. At the other extreme, the compensation threshold $u^W = 0$ corresponds to full front-loading because, in this case, the principal pays the maximal wage whenever his debt to the agent is positive.

For $u^W \in (0, \frac{1}{r_a})$, the wage dynamics consists of two phases. In the beginning, the *back-loading phase* takes place. So long as $u < u^W$, the agent works and accumulates promises for future compensation but does not receive any wage. When his continuation utility u attains (or exceeds) the threshold u^W, the *front-loading phase* begins. In this phase, the agent receives a permanent "base wage" of $r_a u^W$ and a temporary "bonus wage" of $1 - r_a u^W$ whenever $u > u^W$, and hence, in this case, compensation is provided faster than in any alternative compensation scheme.[12] Thus, the optimal arrangement of compensation combines back- and front-loading.

Postponing compensation impacts the principal in two ways. First, it alters the principal-discounted cost of providing compensation, and second, it reduces the principal's ability to incentivize the agent to exert effort in the future. Intuitively, the threshold u^W balances the cost and benefit from postponing compensation, and so this threshold is increasing in r_p. This is formalized in the next proposition, which also shows that full back-loading occurs only if the principal is "very" impatient, and that a moderately impatient principal will fully front-load compensation.

Proposition 8. *Fix B, C, μ, and r_a. If u^W is an optimal threshold for r_p and $\tilde{u}^W$ is an optimal threshold for $\tilde{r}_p > r_p$, then $\tilde{u}^W \geq u^W$. Moreover, there exists $\rho > 0$ such that*

- *$u^W = 0$ is uniquely optimal for all $r_p < r_a + \rho$,*
- *$u^W = \frac{1}{r_a}$ is uniquely optimal for all $r_p > r_a + (\frac{B}{C} - 1)\mu$, and*
- *u^W is interior and generically unique for all $r_p \in (r_a + \rho, r_a + (\frac{B}{C} - 1)\mu)$.*

Propositions 7 and 8 are visualized in Figure 1. In particular, the figure shows that arbitrarily small changes in r_p may lead to discrete jumps in u^W for an impatient principal.

[12] Observe that if the agent's continuation utility exactly equals u^W, then the base wage of $r_a u^W$ maintains the agent's continuation utility constant at that level. Hence, the agent's continuation utility again never drops below u^W.

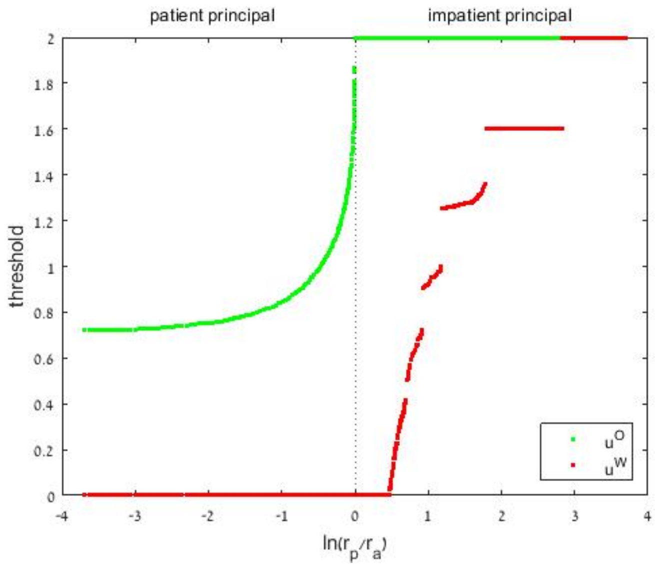

Figure 1: u^O (green) and u^W (red) as a function of $ln(\frac{r_p}{r_a})$, for $B = 2, c = \frac{2}{5}, r_a = \frac{1}{2}$, and $\mu = 2$. Note that the middle of the x-axis represents the case where[14] $r_p = r_a$.

To see the intuition for this discontinuity, suppose that the current state is $u = u^W$, and that the principal considers marginally reducing the threshold (which would require an increase in the current payment). Such a modification would increase the cost of paying the current debt but, on the other hand, would allow the principal to demand more effort in the future. A key determinant in evaluating the profitability of such a modification is the distribution of the *time* at which the additional effort will be provided. Consider the case where $u^W > \frac{1}{r_a} - C$. In this case, the principal will incentivize as much effort as possible upon the arrival of the next opportunity, and this *single* opportunity will exhaust the principal's ability to incentivize effort. Note that regardless of the exact value of u^W, the distribution of the time at which this additional effort will be provided is the same because the (lumpy) opportunity arrives all at once. Therefore, if, for a given r_p, it is profitable to marginally reduce the compensation threshold u^W it is also profitable to marginally reduce any other compensation threshold that is greater than

[14]The optimal thresholds were derived via Monte Carlo simulations. Note that on the extreme right of the figure the green dots are obscured by the red ones as $u^W = u^O = \frac{1}{r_a}$ when r_p is sufficiently large.

$\frac{1}{r_a} - C$. Moreover, the profitability of such a modification is strictly decreasing in r_p. Hence, there exists a critical value r_p^* such that if the principal's discount rate is below (above) r_p^*, marginally reducing the compensation threshold $u^W \in (\frac{1}{r_a} - C, \frac{1}{r_a}]$ makes the principal strictly better (worse) off.[15]

As can be seen in Figure 1, there may exist other values of $r_p < r_p^*$ that feature a similar discontinuity. However, when the threshold u^W is below $\frac{1}{r_a} - C$, characterizing the distribution of time at which the additional effort will be obtained is significantly more complex. To see this, note that the additional effort is not exerted on the next opportunity but, in principle, it may be obtained after any number of opportunities have arrived, depending on the specific realization.

5 Concluding Remarks

In this paper, we studied the consequences of lumpy work requirements. We characterized the optimal contract for any pair of the players' discount rates, both for the case where the randomly arriving opportunities are observed by the principal and for the case where they are privately observed by the agent. In particular, we find that in all specifications, apart from the case where $r_a = r_p$ and opportunities are observable, lumpiness reduces the potential gains from the interaction.

We conclude the paper by highlighting two observations about the impact of the observability of opportunities. First, if opportunities are concealable it is optimal to have perfect bookkeeping of the agents' work, whereas, if opportunities are observable, it is strictly optimal to avoid such bookkeeping and adopt a compensation scheme that features only a low correlation between work and compensation.

Second, our analysis reveals that expediting compensation may sometimes be an efficient tool to mitigate dynamic moral hazard frictions. This is seen, for example, in the dynamics of the optimal contract when the principal is impatient relative to the agent. If the arrival of opportunities is observed by both players, then compensation is fully back-loaded. On the other hand, unless the principal is extremely impatient, the optimal contract under asymmetric information regard-

[15] The value of this threshold is $r_p^* = r_a + (\frac{B}{C} - 1)\mu$ (Proposition 8).

ing the arrival of opportunities combines back-loading and front-loading of compensation. This observation contributes a new angle to the traditional view that suggests that deferring compensation mitigates moral hazard (e.g., Salop and Salop 1976; Lazear 1981; Carmichael 1983).

References

Bird, Daniel, and Alexander Frug. 2019. "Dynamic nonmonetary incentives." *American Economic Journal: Microeconomics*, 11(4): 111–150.

Bird, Daniel, and Alexander Frug. 2020. "Optimal contracts with randomly arriving tasks." *The Economic Journal*, Forthcoming.

Carmichael, Lorne. 1983. "Firm-specific human capital and promotion ladders." *The Bell Journal of Economics*, 14(1): 251–258.

Edmans, Alex, and Xavier Gabaix. 2009. "Is CEO pay really inefficient? A survey of new optimal contracting theories." *European Financial Management*, 15(3): 486–496.

Forand, Jean Guillaume, and Jan Zápal. 2020. "Production priorities in dynamic relationships." *Theoretical Economics*, 15(3): 861–889.

Frankel, Alexander. 2016. "Discounted quotas." *Journal of Economic Theory*, 166: 396–444.

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

Garrett, Daniel F., and Alessandro Pavan. 2015. "Dynamic managerial compensation: A variational approach." *Journal of Economic Theory*, 159: 775–818.

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

Hoffmann, Florian, Roman Inderst, and Marcus M. Opp. 2020. "Only time will tell: A theory of deferred compensation." *Review of Economic Studies*, Forthcoming.

Kalleberg, Arne L. 2011. *Good jobs, bad jobs: The rise of polarized and precarious employment systems in the United States, 1970s–2000s.* Russell Sage Foundation.

Krasikov, Ilia, Rohit Lamba, and Thomas Schacherer. 2020. "Implications of unequal discounting in dynamic contracting." *SSRN Electronic Journal*.

Lazear, Edward P. 1981. "Agency, earnings profiles, productivity, and hours restrictions." *American Economic Review*, 71(4): 606–620.

Li, Jin, Niko Matouschek, and Michael Powell. 2017. "Power dynamics in organizations." *American Economic Journal: Microeconomics*, 9(1): 217–241.

Lipnowski, Elliot, and João Ramos. 2020. "Repeated delegation." *Journal of Economic Theory*, 188: 105040.

Mas, Alexandre, and Amanda Pallais. 2017. "Valuing alternative work arrangements." *American Economic Review*, 107(12): 3722–3759.

Mas, Alexandre, and Amanda Pallais. 2020. "Alternative work arrangements." *Annual Review of Economics*, 12(1): 631–658.

Möbius, Markus. 2001. "Trading favors." Harvard University, Working Paper.

Opp, Marcus M., and John Y. Zhu. 2015. "Impatience versus incentives." *Econometrica*, 83(4): 1601–1617.

Pavan, Alessandro. 2017. "Dynamic mechanism design: Robustness and endogenous types." In *Advances in Economics and Econometrics*. , ed. Bo Honore, Ariel Pakes, Monika Piazzesi and Larry Samuelson, 1–62. Cambridge University Press.

Salop, Joanne, and Steven Salop. 1976. "Self-selection and turnover in the labor market." *Quarterly Journal of Economics*, 90(4): 619–627.

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." *Review of Economic Studies*, 75(3): 957–984.

Spear, Stephen E., and Sanjay Srivastava. 1987. "On repeated moral hazard with discounting." *Review of Economic Studies*, 54(4): 599–617.

Thomas, Jonathan, and Tim Worrall. 1990. "Income fluctuation and asymmetric information: An example of a repeated principal-agent problem." *Journal of Economic Theory*, 51(2): 367–390.

A Appendix: Proofs

Proof of Proposition 1. Note that setting $w(h_t) > 0$ at h_t if the agent has not exerted effort in the past is suboptimal. Moreover, by the definition of T^*, it follows that, in any optimal contract, $w(h_t) = 0$ if the agent did not exert effort in the last T^* units of time. Otherwise, the principal would benefit form reducing compensation at t and decreasing the required effort on the last opportunity along h_t on which effort was exerted. Hence, for the rest of this proof we focus on contracts under which $w(h_t) = 0$ if $t - \sigma_{-1}(h_t) \geq T^*$.

Assume that an opportunity is currently available and denote by σ the random time until the arrival of the next opportunity. Denote by

$$\alpha^* \equiv \min\{\frac{1}{C}\mathbb{E}\left(\int_0^{\min\{T^*,\sigma\}} e^{-r_a t} dt\right), 1\}$$

the maximal effort that the agent is willing to exert on an opportunity in exchange for a conditional promise of length T^* (i.e., setting $w = 1$ until either T^* units of time have passed or an opportunity arrives).

Assume by way of contradiction that under an optimal contract $\alpha(h_t) > \alpha^*$ for a set of histories with positive measure. Since $\alpha^* = 1$ if $r_a = r_p$, for this part of the proof we consider only the case where[16] $r_p < r_a$. Denote by ν the set of finite histories h (of various lengths) for which $\alpha(h) > \alpha^*$ and $\alpha(h') \leq \alpha^*$ for every h' that is a prefix of h. Note that the contract reaches a history in ν with positive probability and that if $h, \tilde{h} \in \nu$, then neither history is a prefix of the other. Thus, it is sufficient to construct a profitable modification of the continuation contract conditional on an opportunity arriving at every $h_t \in \nu$.

Fix $h_t \in \nu$ and assume that an opportunity is available. A conditional promise of length T^* is not enough to compensate the agent for exerting an effort of more than α^* at h_t. Recall that $w = 0$ if no opportunity arrived in the last T^* units of time. Hence, with positive probability some of the compensation for the effort exerted at h_t must be provided after the arrival of the next opportunity. Formally, there exists a set $\nu_1(h_t)$ of continuation histories of h_t (of various lengths) with positive measure such that for every $h_s \in \nu_1(h_t)$ opportunities do not arrive in (t, s), and the agent's continuation utility conditional on an opportunity arriving at h_s is strictly positive (prior to exerting effort).

If there exists $\tilde{\nu} \subset \nu_1(h_t)$ with positive measure (conditional on reaching h_t) such that $\alpha(h_s) < 1$ for every $h_s \in \tilde{\nu}$, then postponing effort (according to the agent's discount factor) from h_t to the histories in $\tilde{\nu}$ is strictly profitable (Observation 1) and does not violate incentive compatibility.

If, on the other hand, $\alpha(h_s) = 1$ for almost all $h_s \in \nu_1(h_t)$, then for every such h_s there exists a set of continuation histories (of various lengths) with positive measure (conditional on reaching h_s) $\nu_2(h_s)$ such that for every $h_{s'} \in \nu_2(h_s)$: 1) no opportunities arrive in (s, s'), and 2) the agent's continuation utility if an opportunity arrives at $h_{s'}$ is greater than his continuation utility at h_s by at least $C(1 - \alpha^*)$. To see why such histories exist, recall that $w = 0$ if no opportunity arrived in the last T^* units of time, and so the maximal agent-discounted expected compensation that can be provided between two successive opportunities is $\alpha^* C$. Hence, to compensate the agent for exerting an effort of 1 on the opportunity at h_s, it must be the case that at least $C(1 - \alpha^*)$ of this compensation is provided after the next opportunity arrives with positive probability.

[16]If $r_a = r_p$, then $T^* = \infty$, and so the assumption that $C + \frac{\mu C}{r_a} < \frac{1}{r_a}$ implies that it is possible to incentivize full effort on all opportunities via an infinite conditional promise.

If there exists $\tilde{\nu} \subset \nu_2(h_s)$ with positive measure (conditional on reaching h_s) such that $\alpha(h_{s'}) < 1$ for every $h_{s'} \in \tilde{\nu}$, then postponing effort from h_s to the histories in $\tilde{\nu}$ does not violate incentive compatibility. Moreover, if effort can be postponed in this manner from a subset of $\nu_1(h_t)$ with positive measure, then doing so increases the principal's profit at h_t. Otherwise, for almost all $h_s \in \nu_1(h_t)$, it holds that $\alpha(h_{s'}) = 1$ for almost all $h_{s'} \in \nu_2(h_s)$. In this case, with strictly positive probability the agent's continuation utility at $h_{s'}$ is greater than his continuation utility at h_t by at least $2C(1 - \alpha^*)$. Continuing in an iterative manner shows that profit can be increased by postponing effort, as otherwise the agent's continuation utility increases without bound with positive probability (which cannot be the case as it is bounded from above by $\frac{1}{r_a}$).

For the rest of the proof we restrict attention to contracts for which $\alpha(h_t) \leq \alpha^*$ for all h_t. Denote by S^* the length of the conditional promise needed to provide the agent with $\alpha^* C$ discounted utils, i.e.,

$$\mathbb{E}\left(\int_0^{\min\{S^*,\sigma_1\}} e^{-r_a t} dt\right) = \alpha^* C.$$

Note that $S^* \leq T^*$ due to the definition of α^*.

Assume by way of contradiction that under an optimal contract $\alpha(h_t) < \alpha^*$ on a set of histories with positive measure. Let n be the minimal element of $\mathbb{N}$ for which the agent's effort on the n^{th} opportunity to arrive is strictly less than α^* with positive probability. Denote by ν the set of histories (of various lengths) after which the agent's effort on the n^{th} opportunity to arrive is strictly less than α^*.

Suppose that an opportunity is available at h_t that is a prefix of a history in ν. The *No-Opportunity-Continuation of length S^* at (h_t, O)*, denoted by NOC-$S^*(h_t, O)$, is the continuation of (h_t, O) of length S^* along which opportunities do not arrive.[17]

First, consider the case where $w = 1$ in almost all of NOC-$S^*(h_t, O)$ for every (h_t, O) that is a (proper or improper) prefix of a history that belongs in ν. Setting $w = 1$ in NOC-$S^*(h_t, O)$ of the first n opportunities to arrive is enough to compen-

[17] Recall that (h_t, O) denotes the event that play proceeds according to h_t and an opportunity arrives at h_t.

sate the agent for exerting an effort of α^* on each of these opportunities. Thus, in this case, it is possible to increase the agent's effort on opportunities that arrive at every $h_t \in \nu$ without altering the compensation plan.

Second, consider the complementary case where there exists a maximal $m \in \{1, 2, \ldots, n\}$ for which there exists a set of histories (of various lengths) with strictly positive measure, ν', such that for every $h_t \in \nu'$: 1) $m - 1$ opportunities arrived along h_t, 2) (h_t, O) is a prefix of a history in ν, and 3) $w < 1$ with positive probability in NOC-$S^*(h_t, O)$. Fix $h_t \in \nu'$ and assume an opportunity is available. By a similar argument to the one used above, the incentive compatibility constraint is not binding at any history that is both a continuation of (h_t, O) and a prefix of a history in ν. Hence, at h_t, the principal can increase his continuation value by increasing the agent's wage in NOC-$S^*(h_t, O)$ and increasing $\alpha(h_s)$ at every $h_s \in \nu$ such that h_s is a continuation of (h_t, O). This is profitable, as the NOC-S^* of the m^{th} opportunity to arrive ends before S^* units of time have passed after the arrival of the m^{th} opportunity. Performing this modification at every $h_t \in \nu'$ increases the principal's value at time zero.

To conclude the proof, denote by $\mathcal{C}^*$ the contract described in the proposition, and note that all the incentive compatibility constraints hold with equality under $\mathcal{C}^*$. Consider an arbitrary incentive-compatible contract $\mathcal{C}$ for which $\alpha(h_t) = \alpha^*$, but the compensation policy is not equivalent to that of $\mathcal{C}^*$. If, under $\mathcal{C}$, $w = 1$ in almost all of NOC-$S^*(h_t, O)$ for almost every finite history h_t, then the agent's expected wage under $\mathcal{C}$ is greater than under $\mathcal{C}^*$. Otherwise, for every finite history h_t for which, with positive probability, $w < 1$ in NOC-$S^*(h_t, O)$ under $\mathcal{C}$, the incentive compatibility constraint of the next opportunity to arrive after (h_t, O) is nonbinding with strictly positive probability. Thus, it is incentive compatible and weakly (strictly) profitable to expedite compensation if $r_p = r_a$ ($r_p < r_a$). $\square$

Proof of Proposition 2. We start this proof by establishing the comparative statics of α^* with respect to λ. If the agent's expected utility from a conditional promise of length T^* is strictly greater than C, then the agent exerts full effort on all opportunities. Moreover, this will remain the case if opportunities become marginally lumpier. Thus, we focus on the case where the agent's expected utility from such a promise is at most C.

The agent's expected utility from a conditional promise of length T^*, as a function of the parameter λ, is $\frac{1}{r_a+\mu/\lambda}(1-e^{-T^*(r_a+\mu/\lambda)})$. The marginal increase in the value of such a promise from making opportunities lumpier is $\frac{\mu}{(\mu+r_a)^2}(1-e^{-T^*(\mu+r_a)}T^*(\mu+r_a))$. Thus, to establish that the α^* is decreasing in λ it is enough to show that making opportunities marginally lumpier has a larger impact on the cost of exerting the required effort than on the value of a conditional promise of length T^*.

Under the assumption that a conditional promise of length T^* does not provide excess compensation, it holds that $C \geq \frac{1-e^{-T^*(\mu+r_a)}}{\mu+r_a}$. Hence, it is sufficient to show that

$$\frac{\mu}{(\mu+r_a)^2}(1-e^{-T^*(\mu+r_a)}T^*(\mu+r_a)) < \frac{1-e^{-T^*(\mu+r_a)}}{\mu+r_a}.$$

Observe that when $\mu+r_a$ is kept constant, this inequality is harder to satisfy for higher values of μ. Thus, it is sufficient to show that it holds for $r_a=0$ i.e., to show that

$$1-e^{-\mu T^*}(1+\mu T^*) < 1-e^{-\mu T^*},$$

which is an inequality that is true for any $\mu T^* > 0$. Note that the above calculation does not depend on the value of T^*. Hence, the same calculation shows that when it is possible to induce full effort, S^* increases when opportunities become marginally lumpier.

Next, we show that making opportunities lumpier is detrimental for the principal. If $\alpha^*=1$, this is an immediate consequence of S^* being increasing in λ. Assume that $\alpha^* < 1$ and let $f(r,\lambda) = \frac{1-e^{-T^*(r+\frac{\mu}{\lambda})}}{r\lambda+\mu}$ denote the $r-$discounted wage that is provided via a conditional promise of length T^* as a function of λ. Note that the average cost of providing a util of compensation is $\frac{f(r_p,\lambda)}{f(r_a,\lambda)}$.

The cross-derivative of $f(r,\lambda)$ evaluated at $\lambda=1$ is equal to

$$\frac{\partial^2 f(r,\lambda)}{\partial r \partial \lambda}|_{\lambda=1} = \frac{\mu e^{-T(\mu+r)}\left(T^*(\mu+r)(T^*(\mu+r)+2)-2e^{T^*(\mu+r)}+2\right)}{(\mu+r)^3}.$$

The sign of this cross-derivative is the sign of $x(x+2)+2-2e^x$, where $x = T^*(r+\mu)$. As this sign is negative, the cross-derivative is negative. As $f(r,\lambda)$ is positive and decreasing in λ, it follows that the average cost of compensating the agent for his effort is increasing in λ (recall that $r_a > r_p$). As the agent's total effort is also decreasing in λ, we can conclude that making opportunities lumpier reduces the principal's value. $\square$

Proof of Proposition 3. When opportunities are observable, the principal's problem can be solved separately for each possible arrival time of the first opportunity. This is because the principal will not provide compensation prior to the first opportunity, and the agent must have a nonnegative continuation utility at all times.

Consider an arbitrary first arrival time σ_1. We assumed that $C + \frac{\mu C}{r_a} < \frac{1}{r_a}$ and so the principal can incentivize the agent to exert maximal effort on the opportunity at σ_1. As the principal is impatient, he will do so in an optimal contract. Moreover, as we established in the main text, the principal will use a threshold structure. Hence, the principal's objective function (conditional on σ_1) is

$$\max_{\tau^O(\sigma_1),\tau^W(\sigma_1)} e^{-r_p\sigma_1}B + \int_{\sigma_1}^{\sigma_1+\tau^O(\sigma_1)} \mu B e^{-r_p t}dt - \int_{\sigma_1+\tau^W(\sigma_1)}^{\infty} e^{-r_p t}dt$$

$$\text{s.t. } C + \int_{\sigma_1}^{\sigma_1+\tau^O(\sigma_1)} \mu C e^{-r_a t}dt = \int_{\sigma_1+\tau^W(\sigma_1)}^{\infty} e^{-r_a t}dt. \tag{7}$$

This constraint (7) states that the agent's continuation utility at σ_1 is zero.

The assumption that $C + \frac{\mu C}{r_a} < \frac{1}{r_a}$ implies that, in optimum, both $\tau^O(\sigma_1)$ and $\tau^W(\sigma_1)$ are interior. To see this, note that the constraint (7) is violated if $\tau^W(\sigma_1) = \infty$ or $\tau^W(\sigma_1) = 0$. Furthermore, setting $\tau^O(\sigma_1) = 0$ implies that $\tau^W(\sigma_1) > \tau^O(\sigma_1)$, and so by slightly increasing $\tau^O(\sigma_1)$ (and decreasing $\tau^W(\sigma_1)$ to maintain incentive compatibility) the agent will exert more effort on opportunities for which he will receive compensation after he has exerted effort. As the principal is impatient, this change is profitable. Finally, setting $\tau^O(\sigma_1) = \infty$ implies that the agent continues exerting effort for an arbitrarily long period of time after he begins receiving compensation. Because the principal is impatient, slightly increasing $\tau^W(\sigma_1)$ (and decreasing $\tau^O(\sigma_1)$ to maintain incentive compatibility) is profitable.

The above discussion implies that the optimal thresholds are given by the FOC of the Lagrangian that corresponds to the above (concave) maximization problem.

The first-order conditions with respect to $\tau^O(\sigma_1)$ and $\tau^W(\sigma_1)$ are, respectively,

$$\begin{aligned}&\mu B e^{-r_p \tau^O(\sigma_1)} - \gamma(\sigma_1) e^{-r_a \tau^O(\sigma_1)} \mu C = 0\\&e^{-r_p \tau^W(\sigma_1)} - \gamma(\sigma_1) e^{-r_a \tau^W(\sigma_1)} = 0,\end{aligned}$$

where γ is the Lagrange multiplier. It follows that $\frac{B}{c} = e^{(r_p - r_a)(\tau^O(\sigma_1) - \tau^W(\sigma_1))}$. Hence, $\tau^O(\sigma_1) = \tau^W(\sigma_1) + K$, where $K > 0$ is a constant that does not depend on σ_1. This relation implies that the LHS of (7) is increasing in $\tau^O(\sigma_1)$ while the RHS is decreasing in $\tau^O(\sigma_1)$. Hence, there is a unique optimal solution that does not depend on σ_1. □

Proof of Proposition 4. Solving the optimal thresholds, τ^O, τ^W, as a function of λ gives

$$\begin{aligned}\tau^O(\lambda) &= \frac{\ln\left(c\mu\left(\frac{B}{C}\right)^{\frac{r_a}{r_a - r_p}} + 1\right) - \ln(c(\mu + \lambda r_a))}{r_a} + \frac{\ln\left(\frac{B}{C}\right)}{r_p - r_a},\\ \tau^W(\lambda) &= \frac{\ln\left(c\mu\left(\frac{B}{C}\right)^{\frac{r_a}{r_a - r_p}} + 1\right) - \ln(c(\mu + \lambda r_a))}{r_a}.\end{aligned}$$

Recall that the principal's value is

$$\mathbb{E}\left(e^{-r_p \sigma_1}\left(B\lambda + \frac{B\mu(1 - e^{-r_p \tau^O(\lambda)})}{r_p} - \frac{e^{-r_p \tau^W(\lambda)}}{r_p}\right)\right).$$

Plugging in the expressions for the optimal thresholds, differentiating with respect to λ, and evaluating at $\lambda = 1$ shows that the principal's value is decreasing in λ:

$$(r_a - r_p)\frac{\mu\left(\frac{B}{C}\right)^{-\frac{r_p}{r_a - r_p}}\left(B\mu\left(\frac{B}{C}\right)^{\frac{r_p}{r_a - r_p}} + 1\right) e^{-r_p\left(\frac{\ln\left(c\mu\left(\frac{B}{C}\right)^{\frac{r_a}{r_a - r_p}} + 1\right) - \ln(c(\mu + r_a))}{r_a} + \frac{\ln\left(\frac{B}{C}\right)}{r_p - r_a}\right)}}{(\mu + r_a)(\mu + r_p)^2},$$

where this expression is negative as $(r_a - r_p) < 0$ and the ratio is positive. □

In the analysis that follows we use a technical lemma that states that for every incentive-compatible contract for which $u > 0$, there exists another incentive-compatible contract that implements the same work schedule via a compensation policy that is (pointwise) weakly lower.

Lemma A.1. *Assume that opportunities are concealable. Moreover, assume that under an incentive-compatible contract the continuation contract at h_t, $\langle \alpha(\cdot), w(\cdot) \rangle$, is such that the agent's continuation utility is $u > 0$. There exists $\tilde{u} < u$ such that for every $u' \in (\tilde{u}, u)$ there exists an incentive-compatible contract $\langle \alpha'(\cdot), w'(\cdot) \rangle$ that provides the agent with a continuation value of u', and for which $w'(h_s) \leq w(h_s)$ and $\alpha'(h_s) = \alpha(h_s)$ at every h_s that is a continuation of h_t.*

Proof of Lemma A.1. Consider an incentive-compatible contract $\langle \alpha(\cdot), w(\cdot) \rangle$ under which the agent's continuation utility is $u > 0$ and normalize the current time to zero. If the agent's expected discounted wage along the histories in which there are no binding incentive compatibility constraints is positive, then the agent's continuation utility at time zero can be decreased by reducing his wage along those histories. If this is not the case, then the agent's wage is almost surely zero prior to a binding incentive compatibility constraint. Hence, concealing all opportunities with probability one is a best response for the agent. However, as $w = 0$ before the agent exerts effort, this best response provides a payoff of $0 < u$. $\square$

Proof of Lemma 1. Let $\langle \hat{\alpha}(\cdot), \hat{w}(\cdot) \rangle$ be an incentive-compatible (continuation) contract under which the agent's expected discounted payoff is $u > 0$. From Lemma A.1 it follows that there exists $\tilde{u} < u$ such that if the agent's continuation utility is in $(\tilde{u}, u)$, then the principal can induce the same work schedule for a lower wage. Thus, there is an open neighborhood to the left of u for which the principal can obtain a value strictly greater than $V(u)$. The strict monotonicity of $V(\cdot)$ follows from the fact that the choice of u is arbitrary.

Next, we show that $V(u)$ is weakly concave. Let $u_1 < u_2$ such that $u_1, u_2 \in [0, \frac{1}{r_a}]$. One (unnatural) way the principal can deliver a promise of $\frac{u_1+u_2}{2}$ is to fictitiously split all opportunities and compensation in half and create two (perfectly correlated) fictitious worlds, each of which contains half of the wage flow and half of each opportunity. Observe that scaling all payoffs by $\frac{1}{2}$ multiplies the players' discounted payoffs by half in any contract, and so any optimal contract in the original non-scaled world is also an optimal contract in each fictitious world. The principal can then use the continuation contract that supports $V(u_1)$ in the non-scaled world to provide the agent with a continuation utility of $\frac{u_1}{2}$ in fictitious world 1, and the continuation contract that supports $V(u_2)$ in the non-scaled world to provide the agent with a continuation utility of $\frac{u_2}{2}$ in fictitious world 2. Since using these continuation contracts cannot increase the principal's payoff, it follows that $V(\frac{1}{2}(u_1 + u_2)) \geq \frac{1}{2}V(u_1) + \frac{1}{2}V(u_2)$, which establishes the concavity of $V(\cdot)$. $\square$

Proof of Proposition 5. We establish this proposition separately for the case where the principal is patient and the case where he is impatient. In each case, we first derive one part of the optimal contract (the work schedule when the principal is impatient, and the compensation policy when he is impatient), and then use the HJB equation to fully derive the optimal contract and show that it is, generically, unique.

Case 1: impatient principal ($r_p > r_a$). The first step of the proof is to show that under any optimal contract the work schedule is $\overline{\alpha}(u) = \min\{1, \frac{1/r_a - u}{C}\}$.

Assume by way of contradiction that $\alpha(\hat{u}) < \min\{1, \frac{1/r_a - \hat{u}}{C}\}$ for some $\hat{u} \in [0, \frac{1}{r_a}]$. Suppose that the current state is $\hat{u}$ and that an opportunity is currently available. If the agent's expected discounted future effort is zero, then it is possible to increase α and increase the agent's compensation in the future without changing his continuation utility. This change is profitable because $B > C$, the principal is impatient, and compensation is provided in the future. If, on the other hand, the agent's expected discounted future effort is positive, then the principal can expedite effort (in the non-Markovian representation of the contract) without altering the compensation policy. By Observation 1 it is profitable for the principal to expedite effort, and, since he does so according to the agent's discount factor, it also relaxes all incentive-compatibility constraints.

The above claim enables us to simplify the HJB equation to

$$(HJB_{imp}) \qquad \sup_{w(u)\in[0,1]} \{-r_p V(u) + V'(u)[r_a u - w(u)] - w(u)$$
$$+\mu\left(\overline{\alpha}(u)B + V\left(u + C\overline{\alpha}(u)\right) - V(u)\right)\} = 0.$$

From the FOC of (HJB_{imp}) it follows that $w(u) = 1$ if $V'(u) < -1$ and that $w(u) = 0$ if $V'(u) > -1$. Since $V(\cdot)$ is weakly concave (Lemma 1), there is a (possibly degenerate) interval $I \subset [0, \frac{1}{r_a}]$ over which $V'(u) = -1$. Note that for any $u^\dagger \in I$ the compensation policy given by

$$w_{u^\dagger}(u) = \begin{cases} 1 & \text{if } u > u^\dagger \\ r_a u^\dagger & \text{if } u = u^\dagger \\ 0 & \text{if } u < u^\dagger \end{cases}$$

is an optimal compensation policy.

Next, we show that, generically, I is degenerate. Fix B, C, μ, and r_a, and let $I(r_p)$ denote the interval (or point) for which $V'(\cdot) = 1$ for a principal with discount rate r_p. To establish the generic uniqueness of optimal contracts we will show that if there exist $\tilde{r}_p < \hat{r}_p$ such that both $I(\tilde{r}_p)$ and $I(\hat{r}_p)$ have a positive measure, then these intervals have a disjoint interior. The result then follows from a standard argument about the density of rational numbers.

Assume by way of contradiction that for some $\tilde{r}_p < \hat{r}_p$, the set $I^* \equiv I(\tilde{r}_p) \cap I(\hat{r}_p)$ has a nonempty interior. Select u^* and $\epsilon > 0$ such that $u^*, u^* - \epsilon \in int(I^*)$.

Fix the optimal compensation policy $w_{u^*}(\cdot)$, and let $\Delta w_s = \mathbb{E}(w_s|u_0 = u^* - \epsilon) - \mathbb{E}(w_s|u_0 = u^*)$ and $\Delta\alpha_s = E(\alpha_s|u_0 = u^* - \epsilon) - \mathbb{E}(\alpha_s|u_0 = u^*)$. Since the chosen compensation policy, $w_{u^*}(\cdot)$, is optimal, we have

$$V(u^* - \epsilon) - V(u^*) = \int_0^\infty e^{-r_p s}(\mu B \Delta\alpha_s - \Delta w_s) ds. \tag{8}$$

As path by path u_s is monotone in u_0, and the work schedule and compensation policies are threshold policies, it follows that $\mu B\Delta\alpha_s - \Delta w_s \geq 0$ for all s, with a strict inequality on a set of times with strictly positive measure. Hence, differentiating the RHS of (8) with respect to r_p shows that the RHS of (8) is decreasing in r_p. However, as $V'(u) = -1$ for all $u \in I^*$ it holds that $V(u^* - \epsilon) - V(u^*) = \epsilon$. Hence, (8) can be satisfied for at most one r_p and so the interior of I^* is empty.

It follows that if the interior of $I(r_p)$ is nonempty, then the compensation policies corresponding to elements of $int(I(r_p))$ are suboptimal for any $r'_p \neq r_p$. Thus, we can index every r_p for which the optimal contract is not unique by a rational number from the interior of $I(r_p)$. Hence, the set of principal-discount factors for which the optimal contract is not unique is countable.

Case 2: patient principal ($r_p \leq r_a$). We begin by showing that the optimal compensation policy is

$$\overline{w}(u) = \begin{cases} 1 & \text{if } u > 0 \\ 0 & \text{if } u = 0. \end{cases}$$

To do so, we show that if $u(h_t) > 0$ then in an optimal contract $w(h_t) = 1$ in the next dt units of time conditional on no opportunity arriving in that interval. If $u(h_t) = \frac{1}{r_a}$, this is immediate. Assume by way of contradiction that $u(h_t) \in (0, \frac{1}{r_a})$, and that the agent does not receive the maximal wage with probability 1 in the next dt units of time conditional on no opportunity arriving in that interval. By arguments analogous to those used in the proof of Lemma A.1, it is possible to expedite wage into the interval $[t, t + dt]$ (conditional on no opportunity arriving) without violating the incentive compatibility constraints in any history that is a continuation of h_t. If the principal is strictly patient, then expediting wage is profitable (Observation 1). If, on the other hand, $r_p = r_a$, expediting wage is profitable as it enables the principal to require more effort in the future.

The above claim enables us to simplify the HJB equation to

$$(HJB_p) \quad V(u) = \sup_{\alpha(u) \in [0, \min\{1, \frac{\frac{1}{r_a} - u}{C}\}]} \{-r_p V(u) + V'(u)[r_a u - \overline{w}(u)] - \overline{w}(u) + \mu\left(\alpha(u)B + V\left(u + C\alpha(u)\right) - V(u)\right)\} = 0.$$

The FOC of HJB_p with respect to $\alpha(u)$ is $B + V'(u + \alpha(u)C)\, C = 0$. Thus, to show that there is a unique optimal contract, it is sufficient to show that $V(\cdot)$ is strictly concave. To do so, we return to the construction used to establish the weak concavity in Lemma 1 and further the analysis by utilizing the structure of $\bar{w}(\cdot)$.

If no opportunity arrives for T units of time, where T is given implicitly by $u_1 = \frac{1-e^{-r_a T}}{r_a}$, an event with positive probability, then the agent's continuation utility in fictitious world 1 is zero while his continuation utility in fictitious world 2 is strictly positive. At this point, the principal can temporarily merge the two fictitious worlds and expedite the owed compensation in world 2 by using the flow wage in world 1. By Observation 1 this modification is profitable for a strictly patient principal and, hence, $V(u)$ is strictly concave if $r_p < r_a$. If $r_a = r_p$, then merging the fictitious worlds increases the discounted effort the agent can be incentivized to exert in the future, which also increases the principal's profit. $\square$

Proof of Proposition 6. To establish this proposition, it is convenient to think of each opportunity as being composed of many "small opportunities." We will show that making opportunities lumpier in the original model is equivalent to a certain change in the correlation structure of these small opportunities.

First, we consider the case where opportunities become lumpier by a rational factor. Assume that opportunities become lumpier by $\frac{N}{M} > 1$, where $N, M \in \mathbb{N}$. We analyze this change by considering an auxiliary representation of the model in which there are $M \times N$ Poisson processes, each with an arrival rate of $\frac{\mu}{N}$, that govern the arrival of the small opportunities. Moreover, we assume that the payoff from exerting full effort on each small opportunity is $(-\frac{C}{M}, \frac{B}{M})$. Both the original and the lumpy versions of the model correspond to appropriately defined correlation structures of the arrival processes in the auxiliary representation.

To map the auxiliary representation to the original model, divide the Poisson processes into N groups of M processes each, such that within a group the processes are perfectly correlated, and across groups the processes are independent. To see why this correlation structure represents the original model, note that when a group of opportunities is available the payoff vector from exerting full effort on all opportunities in the group is $M \times (-\frac{C}{M}, \frac{B}{M}) = (-C, B)$, which is exactly the payoff vector from exerting full effort on a single opportunity in the original model.

Moreover, the probability that a given group arrives in an (infinitesimal) interval dt is $\frac{\mu}{N}dt$, and since the groups are independent, the probability that some group arrives in that interval is $\sum_{i=1}^{N} \frac{\mu}{N}dt = \mu dt$.

To map the auxiliary representation to the lumpy model, divide the processes into M groups of N processes each, such that within a group the processes are perfectly correlated, and across groups the processes are independent. For this correlation structure, the payoff from exerting full effort on all opportunities in a group is $N \times (-\frac{C}{M}, \frac{B}{M}) = (-\frac{N}{M}C, \frac{N}{M}B)$ and the probability that some group of opportunities is available in an interval dt is $\frac{\mu}{N/M}dt$.

Next, we construct a sequence of modifications that begins with the lumpy representation and ends with the original one, such that the first two modifications do not impact the principal's value, and the third modification strictly increases it.

Consider the lumpy representation. The first modification utilizes the idea of splitting the interaction into fictitious worlds introduced in Lemma 1. In particular, we create N fictitious worlds, denoted by $(1, \ldots, N)$, that each contain $\frac{1}{N}$ of the flow wage and M arrival processes, one from each group. We denote the processes in fictitious world n by $(P_1^n, \ldots, P_M^n)$. Note that the arrival processes within each fictitious world are independent of one another, and so each fictitious world is a scaled version of the lumpy representation. Hence, by the argument used in Lemma 1, the sum of the principal's values across all fictitious worlds is equal to his value in the lumpy representation.

The second modification is to the correlation structure of the processes across fictitious worlds. Changing the correlation structure of two arrival processes that are assigned to *different* fictitious worlds does no impact the principal's value in either fictitious world. Hence, so long as the processes within each fictitious world are independent of one another, the correlation across fictitious worlds is immaterial. Thus, we can replace the original correlation structure with the following correlation structure: P_m^n and $P_{m'}^{n'}$ are perfectly correlated if $m - n \overset{(\text{mod} N)}{=} m' - n'$, and independent otherwise. This modification maintains the independence of the processes within each fictitious world. To see this note that for any $n \leq N$ and $m, m' \leq M$, such that $m' \neq m$, the fact that $M < N$ implies that $m - n \overset{(\text{mod} N)}{\neq} m' - n$.

The third modification is to re-merge the fictitious worlds. Note that under the correlation structure created in the second modification, there are N groups of M processes each, such that within a group the processes are perfectly correlated, and across groups the processes are independent. Thus, merging these fictitious worlds creates the auxiliary representation of the original model. Regardless of the relative patience, there are instances in which the principal benefits from merging two fictitious worlds: if $r_p < r_a$ this occurs when in fictitious world i the agent's continuation is positive while in fictitious world j it is zero, whereas if $r_p \geq r_a$ this occurs when in fictitious world i an opportunity is (partially) forgone while in fictitious world j the agent's continuation utility is below its maximal level. It follows that the sum of the principal's values across all fictitious worlds is strictly less than his value in the original model.

Finally, consider the case where $\lambda \notin \mathbb{Q}$. The principal's value is continuous in λ as i) the distribution of arrival times is continuous in λ, and ii) if opportunities are made slightly lumpier then the principal can instruct the agent to incur the same *cost* of effort on every opportunity that arrives by using the same compensation policy. As the set of rational numbers is dense, this establishes the proposition. $\square$

Proof of Proposition 7. In the proof of Proposition 5 we showed that if the principal is impatient, then $u^O = \frac{1}{r_a}$. Moreover, this is also the optimal threshold if $r_a = r_p$. To see this assume that $u = u^O < \frac{1}{r_a}$ and that an opportunity is available. Consider the following alteration to the contract: require some effort on the available opportunity, and then proceed according to the same Markovian contract. If requiring effort forces the principal to reduce the required effort on future opportunities, this alteration essentially expedites effort and does not change the principal's value. However, as the principal front-loads compensation, with positive probability, the agent's continuation utility upon the arrival of the next opportunity is zero. In this case, this alteration strictly increases the principal's value.

Thus, it is left to establish the proposition for the case where $r_p < r_a$, for which, as established in the proof of Proposition 5, $u^W = 0$. Under this compensation policy, if the agent's continuation utility is u, then he receives the maximal wage in all continuation histories for the next $T(u)$ units of time, where $u = \int_0^{T(u)} e^{-r_a t} dt$. Thus, $T(u)$ can be thought of as the *promised duration of compensation* in state u. To establish the proposition it is sufficient to show that the maximal promised duration of compensation, $T(u^O)$, increases with r_p.

Fix T_1 and $T_2 > T_1$ and denote by $\mathcal{X}_n$ the contract in which the principal requires the maximal incentive-compatible effort under the following compensation policy: 1) the promised duration of compensation can never exceed T_2, and 2) the promised duration of compensation can be *increased* to a value in $(T_1, T_2]$ at most n times. Denote the value from contract $\mathcal{X}_n$ by $W(\mathcal{X}_n)$. Observe that $\mathcal{X}_0$ is the contract in which the agent's maximal promised duration of compensation is T_1, and that $\mathcal{X}_\infty$ is the contract in which the promised duration of compensation is T_2.

For any $n \geq 0$, the contracts $\mathcal{X}_{n+1}$ and $\mathcal{X}_n$ induce an identical path of play until the $(n+1)^{\text{th}}$ time that the promised duration of compensation exceeds T_1 under $\mathcal{X}_{n+1}$. At that point, the agent exerts more effort under $\mathcal{X}_{n+1}$ than under $\mathcal{X}_n$. Moreover, from that point onward, the continuation play under both $\mathcal{X}_n$ and $\mathcal{X}_{n+1}$ is given by the same Markovian contract, with different initial states. Note that the net gain, at that point, from incentivizing additional effort by increasing the state is increasing in r_p. It follows, that if a principal with discount factor r_p prefers $\mathcal{X}_{n+1}$ to $\mathcal{X}_n$, then a principal with $r'_p > r_p$ strictly prefers the former contract to the latter. Moreover, note that the principal's preference between $\mathcal{X}_{n+1}$ and $\mathcal{X}_n$ is independent of n. Since

$$W(\mathcal{X}_\infty) - W(\mathcal{X}_0) = \sum_{n=0}^{\infty} (W(\mathcal{X}_{n+1}) - W(\mathcal{X}_n)),$$

it follows that if a principal with discount factor r_p prefers $\mathcal{X}_\infty$ to $\mathcal{X}_0$, then a principal with $r'_p > r_p$ strictly prefers the former contract to the latter.

Finally, note that setting $u^O = 0$ is suboptimal as doing so generates a value of zero, whereas setting a maximal duration of compensation equal to T^* generates a strictly positive profit. $\square$

Proof of Proposition 8. In the proof of Proposition 5 we showed that if $r_p \leq r_a$, then the unique optimal compensation threshold is $u^W = 0$. Thus, it is left to establish the proposition for the case where $r_p > r_a$. Recall that for this case $u^O = \frac{1}{r_a}$ under any optimal contract and, thus, throughout the proof, we assume that $u^O = \frac{1}{r_a}$.

Consider two contracts, $\mathcal{C}_1$ and $\mathcal{C}_2$, that differ in their compensation threshold, $u_2^W > u_1^W$. Denote by $u_{t,i}$ the agent's continuation utility at time t under contract $\mathcal{C}_i$ and let

$$\overline{\tau} = \sup\{t : u_{t,2} \leq u_2^W\}.$$

That is, $\overline{\tau}$ is the latest time at which the agent's continuation utility is weakly lower than u_2^W under $\mathcal{C}_2$. Note that $\overline{\tau}$ is finite (almost surely) by the Borel–Cantelli lemma.

Observe that since $u_2^W > u_1^W$, it holds that $u_{t,1} \leq u_{t,2}$ for all t. This implies that the agent exerts the same effort on every opportunity that arrives before $\overline{\tau}$ under both $\mathcal{C}_1$ and $\mathcal{C}_2$. In addition, it implies that the compensation for effort exerted on those opportunities is postponed under $\mathcal{C}_2$ relative to $\mathcal{C}_1$. Denote by $g(r_p)$ the gain from this postponement as a function of r_p. Moreover, from $\overline{\tau}$ onwards, under $\mathcal{C}_2$ the agent exerts weakly less effort than he does under $\mathcal{C}_1$, and he receives a wage of $w_t = 1$ at all times. Let $d(r_p)$ denote the difference in the time-zero discounted continuation value from $\overline{\tau}$ onward between $\mathcal{C}_1$ and $\mathcal{C}_2$. The net gain from replacing $\mathcal{C}_1$ with $\mathcal{C}_2$ is $g(r_p) - d(r_p)$. Note that $g(\cdot)$ is increasing and $d(\cdot)$ is decreasing. Hence, whenever $g(r_p) \geq d(r_p)$, we also have $g(r'_p) > d(r'_p)$ for all $r'_p > r_p$, which establishes the monotonicity of u^W.

Next, we show that $u^W = \frac{1}{r_a}$ is an optimal threshold if and only if $r_p \geq r_a + (\frac{B}{C} - 1)\mu$. The monotonicity established in the first part of the proof will then imply that $u^W = \frac{1}{r_a}$ is the unique optimal threshold if $r_p > r_a + (\frac{B}{C} - 1)\mu$. An upper bound on the principal's marginal net gain from providing the agent with a util at present is attained by the agent exerting an agent-discounted util on the first opportunity to arrive. Note that if $u^W > \frac{1}{r_a} - C$ and $u = u^W$, then this upper bound is attained. The value of this upper bound is given by

$$\int_0^\infty \mu e^{-\mu t} \frac{B}{C} e^{(r_a - r_p)t} dt - 1 = \frac{\mu}{\mu + r_p - r_a} \frac{B}{C} - 1.$$

Since the principal is impatient, it is straightforward to show that this expression is positive if and only if $r_a + (\frac{B}{C} - 1)\mu - r_p > 0$. It follows that if $r_p > r_a + (\frac{B}{C} - 1)\mu$, then providing compensation while $u < \frac{1}{r_a}$ is suboptimal. On the other hand, it is strictly suboptimal to set $u^W = \frac{1}{r_a}$ if $r_p < r_a + (\frac{B}{C} - 1)\mu$, as for such discount rates it is profitable to pay the agent when $u > \frac{1}{r_a} - C$ (for such values the aforementioned upper bound is attained). The monotonicity established in the first part of the proof shows that $u^W = \frac{1}{r_a}$ is the unique optimal threshold if $r_p > r_a + (\frac{B}{C} - 1)\mu$.

Finally, we establish that there exists $r_p > r_a$ for which $u^W = 0$. Let τ denote the random arrival time of the first opportunity on which the agent will not exert full effort under the contract with $u^W = 0$. Note that under any other contract, the agent will not exert full effort (weakly) earlier. It follows that the marginal value from decreasing the agent's continuation utility is at least $E(e^{-r_p \tau})\frac{B}{C} > 0$. If $r_a = r_p$ the timing of compensation does not affect the principal's cost of providing compensation. This, in turn, implies that when $r_a = r_p$ the principal's marginal gain from providing compensation is at least $E(e^{-r_p \tau})\frac{B}{C} > 0$ for all u. By the continuity of payoffs in r_p, it follows that there exists $\rho > 0$ such that the principal strictly benefits from full front-loading of compensation if $r_p < r_a + \rho$. $\square$