FEEDBACK DESIGN IN DYNAMIC MORAL HAZARD

JEFFREY C. ELY
Department of Economics, Northwestern University

GEORGE GEORGIADIS
Kellogg School of Management, Northwestern University

LUIS RAYO
Kellogg School of Management, Northwestern University

We study the joint design of dynamic incentives and performance feedback for an environment with a coarse (all-or-nothing) measure of performance, and show that hiding information from the agent can be an optimal way to motivate effort. Using a novel approach to incentive compatibility, we derive a two-phase solution that begins with a "silent phase" where the agent is given no feedback and is asked to work non-stop, and ends with a "full-transparency phase" where the agent stops working as soon as a performance threshold is met. Hiding information leads to greater effort, but an ignorant agent is also more expensive to motivate. The two-phase solution—where the agent's ignorance is fully frontloaded—stems from a "backward compounding effect" that raises the cost of hiding information as time passes.

KEYWORDS: Principal-agent, moral hazard, information design.

1. INTRODUCTION

A KEY DESIGN component of many jobs, such as those in professional services firms, is the performance feedback offered to employees, as it allows them to adjust their behaviors and learn what their future rewards might look like. While some experts argue that a policy of full transparency is best—that is, keeping employees fully apprised of their performance—such practice is far from uniform as employers may see a strategic gain from hiding information or postponing its release; see, for example, Maister (1993).

Here we study the optimal joint design of performance feedback and monetary rewards in a dynamic-agency environment where the underlying monitoring technology is coarse. We argue that because of this coarseness, concealing information from the agent about their performance can be a profitable way to increase their effort, despite the fact that doing so requires paying them greater rewards.

In our baseline model, the principal wishes to maximize the agent's aggregate effort net of monetary rewards, and must discourage strategic pauses by the agent. The only performance measure available to the principal is an all-or-nothing signal of effort. The resulting problem is challenging because the (dynamic) feedback policy can in principle be highly complex, and interact in complicated ways with the chosen monetary rewards.

To solve it, we show that it suffices to restrict attention to policies that discourage instantaneous effort pauses by the agent—akin to a first-order approach. This allows us to show that the optimal contract consists of two phases, with the agent first kept fully in the dark and then kept fully apprised of their performance. This bang-bang solution arises from a desire to keep the agent working beyond the time when the performance measure records a success in combination with a "backward-compounding effect," which implies that hiding information is more costly if that hiding occurs farther in the future, and hence it is optimal to frontload the agent's ignorance.

Our finding that concealing information early on can be profitable for the principal may speak to the early stages of professional services jobs, where the firm's superiors sometimes keep performance information hidden from their associates to artificially extend the initial trial period of their careers. In tandem with that lack of transparency, associates are granted a finite time to secure promotion and may earn larger rewards the sooner they are promoted. These features have analogs in our model as well. (We expand on this discussion in Section 4.)

We have chosen a particularly stark baseline model for both tractability and transparency. This model also naturally lends itself to various extensions where a two-stage contract remains optimal. We consider, in particular, the possibility that players secure additional payoffs from a continued relationship, that the agent is able to directly learn about a success, and that the output technology features a "learning-by-doing" component. In all these cases, backward compounding again leads to a bang-bang solution that frontloads the agent's ignorance.

Finally, we consider an extension where the principal is able to monitor the agent both indirectly (via the success) and directly by observing effort itself (at a cost). Here we show that provided direct monitoring is sufficiently costly, withholding information by means of an initial silent phase dominates full disclosure once again.

Related Literature. We contribute first and foremost to the literature on dynamic agency models under moral hazard; see Sannikov (2008), and for an overview, Georgiadis (2024). Canonical models assume that incremental output at each instant (or each period) can take many values, whereas our work is closer to models where it is binary. In Mason and Välimäki (2015), for example, a principal designs a contract to motivate a Poisson "breakthrough," and in Green and Taylor (2016) two breakthroughs are required. Halac, Kartik, and Liu (2016) considered a setting where players learn about the feasibility of a breakthrough—which, in our language, is equivalent to a declining hazard rate—and the agent is privately informed about their ability.[1] These models, in contrast to ours, do not permit the principal to strategically provide feedback to the agent.

We also contribute to the literature on information design. Rayo and Segal (2010) and Kamenica and Gentzkow (2011) studied the optimal provision of information in static environments, and Ely (2017) and Renault, Solan, and Vieille (2017) extended these analyses to dynamic settings. The latter two papers consider a game between a receiver (e.g., an investor) and a sender (e.g., an advisor) where the sender observes a payoff-relevant state variable that evolves exogenously, and chooses a message policy to entice the receiver to take an action. Here the optimal policy is effectively a static one, with the sender behaving as if they were myopic.[2]

---

[1]Keller, Rady, and Cripps (2005) and Bonatti and Horner (2011) analyzed the equilibria of such good-news Poisson experimentation models.

[2]Hörner and Skrzypacz (2017) discussed conditions under which the optimal policy is not static.

Ely and Szydlowski (2020), Orlov, Skrzypacz, and Zryumov (2020), Ball (2023), Kaya (2023), and Smolin (2021) considered games between a principal (sender) and an agent (receiver) in which the agent decides when to stop supplying an action, while the principal monitors the evolution of a payoff-relevant state (which is independent of the action) and transmits messages to entice the agent not to stop. In all of these papers, the agent's action is observable (so there is no moral hazard) and the agent's monetary rewards are exogenously given or severely restricted.[3] These papers therefore exclusively focus on the design of feedback, whereas we study the interaction between feedback and monetary rewards and the optimal joint design of these, and do so in a moral-hazard environment where the agent can secretly manipulate that feedback.[4]

In a career-concerns framework, Hörner and Lambert (2021) studied the design of information provided to an external market to motivate the agent. Similarly, in Ostrizek (2022), all information is public and the focus is on the tradeoff between precise information, which enables higher-powered monetary incentives, and coarse information, which slows down learning and reduces future agency rents.[5]

2. BASELINE MODEL

We begin with the simplest version of the model, which we then extend in various directions. At each instant t and up to an exogenous deadline T, an agent privately chooses whether to spend effort toward producing a binary signal, which we call a "success." We take T to be large but finite. Success, which occurs at most once, is observed only by a principal and arrives stochastically as a function of the agent's accumulated effort. The principal, who enjoys commitment power, designs monetary rewards together with a real-time feedback policy on the basis of that signal, with the goal of maximizing the agent's total effort net of monetary payments. Both players are risk neutral and the agent is cash-constrained.[6]

At each moment, the agent either works or waits—that is, effort is binary—and when working, incurs a flow cost $c < 1$. The agent's probability of success at or prior to any point in time is given by the function $F : [0, T] \to [0, 1]$ of the cumulative time e that the agent has spent working. $F(\cdot)$ is weakly increasing and satisfies $F(0) = 0$ and $F(T) \leq 1$. A useful equivalent representation is to imagine that a hidden random *effort requirement* $z \geq 0$ is drawn according to the c.d.f. F and success occurs as soon as the agent's accumulated effort e reaches z.

We will assume that F has a differentiable density f and that both $f(e)$ and the hazard rate $\lambda(e) := f(e)/(1 - F(e))$ are weakly decreasing in e. Some of our results will also

---

[3]In Smolin (2021), the reward is a function of the feedback provided, but this function is exogenous and not part of the design. In Kaya (2023), the principal pays a share of profit that is fixed and not responsive to the agent's actions. In Ely and Szydlowski (2020), the feedback serves as a reward for observable effort, while in the present setting effort is unobservable.

[4]In a setting with multiple agents, Ely, Georgiadis, Khorasani, and Rayo (2023) also permitted the principal to flexibly choose the agents' reward schedule. However, characterizing the optimal contest hinges on having sufficiently many agents so that full rent-extraction is possible.

[5]Also related is the contracting literature where the principal acquires costly information about the agent's effort(s); see, for example, Georgiadis and Szentes (2020) and Orlov (2022). In these models, information matters only indirectly via its impact on the players' payments.

[6]While assuming a single success is stylized, our goal is to capture the notion that, in reality, performance measures may be coarse due to monitoring costs.

make use of the following function:

$$\Phi(t) := F(t)\frac{d}{dt}\frac{1}{f(t)} \tag{1}$$

whose significance will soon be clear. The problem of maximizing expected effort minus expected rewards will be well-behaved when Φ is weakly increasing, which we will assume.

These assumptions are satisfied if, for example, F arises from good-news Poisson experimentation (as studied, for instance, by Halac, Kartik, and Liu (2016)) where the project can be "good" or "bad" (or the agent's ability "high" or "low"), unknown to both players, and a success arrives with constant hazard rate only if the project is good and the agent is working. This type of technology may be (roughly) representative of a professional services firm where the ability of an associate is initially unknown. A special case is that of a constant hazard rate.

The principal offers the agent a monetary reward $R(t)$ if success arrives at t, and no reward (or punishment) when it never arrives.[7] So that the reward does not reveal any undesired information, we assume (without loss) that it is paid at T. In addition, the principal designs a feedback policy that specifies, for each point in time, a probability distribution over messages as a function of past messages and the time of past success, if it happened. By standard arguments, it is sufficient to consider direct feedback policies in which the principal recommends to the agent whether to work or wait at each instant. We shall be sure to discuss how the resulting effort recommendation policies can be implemented using a feedback policy about the agent's success.

Because pauses in effort have no impact on the players' payoffs and merely waste time, it is without loss for the principal to recommend (possibly stochastically) that the agent either continue working and await further recommendations or permanently quit, rather than temporarily pause. Such a recommendation policy can be represented by a function $q(s|t)$ denoting the probability that the agent is still asked to work at date s conditional on having succeeded at $t \le s$, together with a function $r(s)$ denoting the probability that the agent is still asked to work at date s conditional on *not* having succeeded before then. Note that since a recommendation to quit is a permanent one, both recommendation functions must be non-increasing in s (and so that the players' expected payoffs are well defined, we assume that they are integrable in all arguments).

Especially useful for the analysis is the total (ex ante) probability $p(s)$ that the agent is asked to work at least until date s. This is given by

$$p(s) = r(s)\left[1 - F(s)\right] + \int_0^s r(u)f(u)q(s|u)\,du, \tag{2}$$

where the first term is the probability that an agent who has not yet succeeded is asked to work at s, and the second term is the probability that an agent who has already succeeded is asked to work at s. Note that regardless of the recommendation policy, we have $p(s) \ge r(s)[1 - F(s)]$. Also useful is the function $Q(t) := \int_t^T q(u|t)\,du$, which measures the expected *future* work for an agent who succeeds at t and obeys all recommendations.

---

[7]Restricting to rewards that condition on time of success alone is without loss because, owing to the agent's risk neutrality, a reward schedule $R(t, x)$ that conditions on a second random variable x, such as past feedback, can be replaced by a reward function equal to $\mathbb{E}_x R(t, x)$ (where the expectation is taken at time 0) without altering the ex ante incentive constraint, and doing so prevents adding further incentive constraints in the future. Furthermore, rewarding the agent in the absence of a success merely hinders incentives to work, and imposing a punishment if a success never arrives is infeasible due to the agent being cash-constrained.

The agent's expected payoff from obeying the recommendations is

$$\underbrace{\int_0^T r(s)R(s)f(s)\,ds}_{\text{expected reward}} - c \times \underbrace{\int_0^T p(s)\,ds}_{\text{expected effort}},$$

where $r(s)R(s)f(s)$ is the ex ante expected reward at time s measured in flow terms. The principal chooses a reward schedule and a recommendation policy to maximize

$$\int_0^T p(s)\,ds - \int_0^T r(s)R(s)f(s)\,ds$$

subject to the incentive compatibility constraint that the agent always finds it optimal to obey all recommendations. Here we have assumed that success is only a signal of effort that delivers no direct benefits to the players, which we shall relax later on.[8]

3. INCENTIVE COMPATIBILITY

There are a variety of ways the agent can deviate from recommendations, including pausing and restarting at any time. Fortunately, as we will show, the optimal policy can be derived by focusing on a family of "local" incentive constraints.[9] We will derive necessary conditions for a policy to dissuade the agent from brief (instantaneous) pauses, which we can use to identify a candidate optimal policy. We will then verify that there exist no profitable global deviations.

3.1. *Instantaneous Pauses*

Because the hazard rate of success weakly decreases over time, the principal would like to promise, other things equal, greater rewards as time goes by. But this creates a challenge: the agent will be tempted to pause temporarily so they can succeed later, where rewards are greater, rather than now.[10] This leads us to conjecture that temporary pauses, rather than permanent ones, will be the hardest ones to deter. A natural place to start, then, is making sure that at least the shortest of such pauses—that is, instantaneous ones—are deterred.

The expected payoff earned by the agent from t onward if they obey all recommendations, computed from the standpoint of time 0, is

$$U(t) := \int_t^T r(s)R(s)f(s)\,ds - c \times \int_t^T p(s)\,ds. \tag{3}$$

Now suppose the agent obeys all recommendations before t and after $t + \Delta t$, but shirks during the interval in between. Such a deviation changes the arrival rate of success and

---

[8]We have also assumed for simplicity that players do not discount time; our results would remain qualitatively unchanged (as viewed from a time zero perspective) if they had a common discount rate.

[9]This is reminiscent of the sufficiency of the "no-postponed-setbacks" condition in Feng, Taylor, Westerfield, and Zhang (2024).

[10]Such temptation to pause is known as the "dynamic agency" effect; see, for example, Halac, Kartik, and Liu (2016). What is novel about this effect in our model is that it is impacted by the principal's feedback.

therefore changes the distribution of recommendations after $t + \Delta t$ as well. The agent's continuation payoff at t, considering that pause, is

$$\widetilde{U}(t, \Delta t) := \int_{t+\Delta t}^{T} r(s)R(s)f(s - \Delta t)\, ds - c \times \int_{t+\Delta t}^{T} p\big(s|\omega_t^{\Delta t}\big)\, ds,$$

where the integrals begin at $t + \Delta t$ because the agent's flow payoff during the pause is zero, and $p(s|\omega_t^{\Delta t})$ denotes the total probability that the agent continues to spend effort through $s \geq t + \Delta t$ following this deviation (we derive an expression for this term in the [Appendix](#)). Incentive compatibility requires that $U(t) \geq \widetilde{U}(t, \Delta t)$, or equivalently, upon subtracting $U(t + \Delta t)$ from both sides of this inequality,

$$U(t) - U(t + \Delta t) \geq \int_{t+\Delta t}^{T} r(s)R(s)\big[f(s - \Delta t) - f(s)\big]\, ds - c \times \int_{t+\Delta t}^{T} \big[p\big(s|\omega_t^{\Delta t}\big) - p(s)\big]\, ds.$$

Dividing through by Δt and taking the limit as it converges to zero allows us to establish a necessary condition for incentive compatibility.

PROPOSITION 1: *The recommendation policy $q(\cdot|\cdot)$, $r(\cdot)$ and reward schedule $R(\cdot)$ are incentive compatible only if, for all t,*

$$r(t)R(t)f(t) - cp(t) \geq cr(t)Q(t)f(t) + \int_{t}^{T} r(s)\big[R(s) - cQ(s)\big]\big|f'(s)\big|\, ds$$

$$- c\int_{t}^{T} r(s)f(s)\, ds, \tag{LIC}$$

where $p(t)$ is given in (2), and $Q(t) = \int_{t}^{T} q(u|t)\, du$ is the expected future work for an agent who succeeds at t.

We call this inequality the *local* incentive constraint, and we say that a policy (q, r, R) is locally incentive compatible if it meets it. The left-hand side measures the agent's on-path expected flow rents at t, that is, their expected reward minus instantaneous cost. These flow rents correspond to $-U'(t)$. The right-hand side is $\lim_{\Delta t \to 0}[\widetilde{U}(t, \Delta t) - U(t + \Delta t)]/\Delta t$, which represents the marginal impact of a pause at t on future rents. This pause has three effects (corresponding to each of the three terms on the right). First, it eliminates the agent's expected future effort cost $cQ(t)$ in the event they work and succeed at t. Second, it lowers the agent's accumulated effort and therefore raises the density of success $f(s)$ going forward, which in turn raises the chance of earning each of the net future rewards $r(s)[R(s) - cQ(s)]$. And, finally, it raises the chance that at each future date the agent has not yet succeeded and must therefore keep working.[11]

3.2. *Minimal Rewards*

Proposition [2](#) shows that for any given recommendation policy, there is a unique least-expensive way to deter all instantaneous pauses. This is achieved by meeting (LIC) with equality at all times, starting at T and working backward.

---

[11]The probability that the agent has not yet succeeded at s is $1 - F(s)$; a pause today (which lowers accumulated effort from s to $s - \Delta s$) raises that probability by $f(s)$, which, multiplied by $cr(s)$ and integrated over every future date, is equal to the third term on the right-hand side of (LIC).

PROPOSITION 2: *Given any recommendation policy $q(\cdot|\cdot)$, $r(\cdot)$, there exists a unique reward schedule $R(\cdot)$ that satisfies the local incentive constraint with equality at every t. It is given by*

$$r(t)R(t) = c\left[\frac{p(t)}{f(t)} - \int_t^T \frac{f'(s)}{f(s)^2}p(s)\,ds - \int_t^T r(s) - r(t)q(s|t)\,ds\right]. \qquad (4)$$

Moreover, this reward schedule is pointwise smaller than any other implementing reward schedule.

Intuitively, because $R(t)$ in the local incentive constraint is affected only by future rewards rather than past ones, it is possible by working backward from T to meet this constraint with equality at all times. Furthermore, it is desirable to do so: raising the reward schedule above that level over some interval of time would force the principal to raise rewards at all past times (to deter earlier pauses) which would needlessly inflate the principal's costs.

To interpret (4), it is useful to consider a contract with a *deterministic deadline* in which the principal selects a time $\tilde{T} \leq T$ after which all effort stops, and asks an unsuccessful agent to work with probability 1 until that time (i.e., $r(t) = 1$ for all $t \leq \tilde{T}$ and $r(t) = q(t) = 0$ thereafter). As it turns out, a contract in this class will be optimal. Equation (4) simplifies to

$$R(t) = c\left[\frac{p(t)}{f(t)} - \int_t^{\tilde{T}} \frac{f'(s)}{f(s)^2}p(s)\,ds - \int_t^{\tilde{T}} 1 - q(s|t)\,ds\right]. \qquad (5)$$

The first term on the right-hand side is the reward level that would give the agent zero flow rent at time t, measured from an ex ante perspective. This term grows with $p(t)$ because a higher work probability implies a larger ex ante cost for the agent, and falls with $f(t)$ because a greater density of success means the reward $R(t)$ is more likely to materialize. If there were no future dates, this zero-rent reward is all the principal would need to offer. The second term represents a "backward compounding" effect arising from the fact that a greater future reward requires a greater present reward, as otherwise the agent would pause. This backward compounding is modulated by $-f'/f^2$ (the speed at which $1/f$ grows) because the faster this ratio grows (i.e., the faster the likelihood of success drops), the greater the future rewards have to be.[12] The last term is an "information rebate" that the principal gets if they inform the agent sometime in the future about a success at t: the lower the $q(s|t)$, the lower the effort the agent is asked to exert in the future *after* succeeding at t, and hence the lower the promised reward $R(t)$ needs to be.

To illustrate, consider two polar opposite policies, each with a deterministic deadline: keeping the agent completely in the dark and keeping them fully informed. If the principal offers zero feedback, while asking the agent to work with probability 1 until $\tilde{T}$, we have $q(s|t) \equiv p(t) \equiv 1$ until that deadline. As a result, the minimal implementing reward schedule is

$$R^{\text{silence}} = \frac{c}{f(\tilde{T})}.$$

---

[12]The novelty of this backward compounding effect is that it incorporates the feedback policy; that is, it tells us how we need to adjust rewards so that the dynamic agency consideration (i.e., the temptation to delay effort) does not lead the agent to pause given how information is being revealed.

Intuitively, to minimize backward compounding, the principal wishes to grant zero rents at the terminal time; hence this flat schedule is the closest the principal can get to the entire zero-rent schedule $R(t) = c/f(t)$ (which is weakly increasing) without provoking any pauses. Because the agent receives no feedback, there is no rebate for the principal.

If instead the principal keeps the agent fully apprised—in which case the agent stops working as soon as they succeed—we have $q(s|t) \equiv 0$ and $p(t) = 1 - F(t)$ until the deadline $\widetilde{T}$. We term this design the *pronto* policy. The corresponding minimal reward schedule is

$$R^{\text{pronto}} = \frac{c}{\lambda(\widetilde{T})},$$

which is flat for a similar reason as before: the ideal schedule (in this case, the zero-rent schedule net of the rebate) is also increasing; thus, given that the principal wishes to grant zero rents at the terminal time, a flat schedule is the closest the principal can get to that ideal without causing the agent to pause.

Embedded in both these policies is a dynamic version of the classic rent/efficiency trade-off: the principal exactly internalizes the cost of effort at the margin (in this case, at the terminal date) but because $R(\widetilde{T})$ grows with $\widetilde{T}$, they must also pay greater infra-marginal rents if they seek to expand the overall gains from trade. What differs between the two policies is, on the one hand, the information rebate and, on the other, the maximum expected effort that can be asked of the agent. Because the pronto policy keeps agent rents to a minimum, it minimizes the principal's expected cost per unit of effort; but since the agent quits as soon as they succeed, it also creates an upper bound on the agent's expected total effort. Silence, in contrast, allows the agent to work until the deadline, but since there is no rebate, the cost for the principal could be very large.

In between these two examples, there are vastly many ways for the principal to offer less than immediate feedback. One example is a "delay mechanism" where the agent is informed of a success after a delay d. Here $q(s|t) = \mathbb{I}_{\{s \leq t+d\}}$ while $p(t) = 1$ up to time d and equal to $1 - F(t - d)$ thereafter. Thus, the minimal reward schedule takes a more complex form. Namely, $R(\widetilde{T}) = c[1 - F(\widetilde{T} - d)]/f(\widetilde{T})$ (so that rents are zero at the very end) and

$$R'(t)/c = \begin{cases} 1 & \text{if } t \in [0, d), \\ 1 - f(t-d)/f(t) & \text{if } t \in [d, \widetilde{T} - d], \\ -f(t-d)/f(t) & \text{if } t \in (\widetilde{T} - d, \widetilde{T}], \end{cases}$$

which means that the reward schedule first increases, then decreases, and finally decreases at an even faster rate.

Such delay mechanisms have been found to be optimal in different settings (e.g., Ely (2017)), but as we shall see, are suboptimal in ours. The optimal policy will instead be a simpler combination of the two polar opposites above.

4. OPTIMAL POLICY

Here we characterize the optimal policy by making use of the minimal reward schedule in Proposition 2. Our first step is to use that schedule to express the principal's objective solely in terms of $p(t)$ and $r(t)$.

LEMMA 1: *The principal's payoff evaluated at the minimal implementing reward schedule is*

$$\int_0^T p(t)\,dt - c\int_0^T \underbrace{p(t)\big(1 + \Phi(t)\big) - \big(r(t) - p(t)\big)}_{\text{virtual effort}}\,dt, \tag{Obj}$$

where $\Phi \equiv F \times (1/f)'$ satisfies $\Phi(0) = 0$ and by assumption is weakly increasing.

The first term in (Obj) is total effort. The second term, whose integrand we term "virtual effort," is the total remuneration, that is, true cost plus information rents for the agent due to the backward compounding of rewards. Observe that at time zero, virtual effort is equal to true effort (with both equal to 1), as backward compounding is not a factor then. The function $\Phi(t)$ captures the compounding effect: a greater $(1/f(t))'$ calls for greater future rewards and hence greater past ones, whereas a larger $F(t)$ means these past rewards are paid more often. The term $r(t) - p(t)$ captures the information rebate: when the principal pays a successful agent with information (which lowers $p(t)$ below $r(t)$), there is less need for monetary rewards.

We shall find the optimal policy by first solving a relaxed problem where, in addition to ignoring non-local deviations, the principal selects $p(t)$ and $r(t)$ directly, subject only to an upper and lower bound for each and a monotonicity constraint for $r(t)$, without worrying about the need to generate the function p with a suitable choice of q. In particular, we consider the program

$$\sup_{r(\cdot),\,p(\cdot)} \quad \int_0^T p(t)\big\{1 - 2c - c\Phi(t)\big\}\,dt + c\int_0^T r(t)\,dt \tag{P}$$

$$\text{s.t.} \quad r(t)\big[1 - F(t)\big] \le p(t) \le 1 \quad \text{for all } t, \tag{6}$$

$$0 \le r(t) \le 1 \quad \text{for all } t \text{ and non-increasing in } t, \tag{Feas}$$

where the objective is equal to (Obj) upon rearranging terms, the constraint $r(t)[1 - F(t)] \le p(t)$ captures the requirement that the total probability of work cannot fall below the probability of an unsuccessful agent working, the constraint $p(t) \le 1$ stems from the requirement that $p(t)$ is a probability, and the last constraint stems from $r(t)$ being a survival function.

Proposition 3 shows that this problem admits a "bang-bang" solution where $p(t)$ and $r(t)$ are always set equal to their upper or lower bounds.[13]

PROPOSITION 3: *Let t^* be the earliest time when $1 - 2c - c\Phi(t) \le 0$. Provided $T > t^*$, the relaxed problem* (P) *is solved by setting*

$$p(t) = \begin{cases} 1 & \text{if } t \in [0, t^*], \\ 1 - F(t) & \text{if } t \in (t^*, T^*], \\ 0 & \text{if } t \in (T^*, T], \end{cases} \quad \text{and} \quad r(t) = \begin{cases} 1 & \text{if } t \in [0, T^*], \\ 0 & \text{if } t \in (T^*, T], \end{cases}$$

for some $T^* \in (t^*, T]$.

---

[13]The problem can be solved in two steps. First, fix $r(t)$ and set $p(t)$ equal to either $r(t)[1 - F(t)]$ or 1, according to whether the expression in curly braces—which is either always negative or, once it becomes negative, it remains so throughout—is negative or positive. Second, substitute this value of $p(t)$ in the objective and find an optimal $r(t)$. Because the resulting objective is linear in $r(t)$, it is without loss to set this function equal to 0 or 1. See the Appendix for further details.

This is a contract with a deterministic deadline T^* and two phases. In phase 1, lasting until $t^* < T^*$, the agent is asked to work with probability 1. In phase 2, the agent is asked to work only if they have not yet succeeded. The schedule p is uniquely implemented by the recommendation policy that never tells the agent to stop before t^*, and after that, tells them to stop as soon as they succeed. A successful agent will never work longer than an unsuccessful one because the latter is cheaper to motivate. Notice, finally, that the cutoff $t^* > 0$ if and only if $c < 1/2$. The corresponding reward schedule is then obtained by substituting p, r, and the corresponding q into (4).

This solution will constitute the backbone of every optimal policy, as we show next.

THEOREM 1: *Suppose $T > t^*$. Every optimal policy consists of at most two phases*:
1. ***Silent phase***: $t \leq t^*$. *Here the principal asks the agent to work with probability 1 regardless of their success, and remains silent throughout—that is, $q(t|s) \equiv r(t) \equiv 1$. If the agent succeeds at any time during this phase, they earn reward*

$$c/\lambda(T^*) + cF(t^*)/f(t^*).$$

2. ***Pronto phase***: $t \in (t^*, T^*]$. *Here the principal asks the agent to quit as soon as they succeed while remaining otherwise silent—that is, $r(t) \equiv 1$ and $q(t|s) \equiv 0$. If the agent succeeds at any time during this phase, they earn reward*

$$c/\lambda(T^*).$$

If the agent does not succeed by T^*, then they are advised to quit and earn zero reward. Phase 2 always has positive length, whereas phase 1 has positive length if and only if $c < 1/2$.[14]

Intuitively, the pronto policy is used for some duration because even though it allows the agent to promptly quit upon success, it minimizes the principal's cost per unit of effort (i.e., it minimizes agent rents) as the agent is cheapest to motivate when kept fully informed.[15] If the principal wishes greater effort than a pronto phase alone can achieve, the agent must at least sometimes be kept in the dark. Moreover, because ignorance necessitates greater rewards and these get compounded backward, it is best that such ignorance is maximally frontloaded; hence the initial silent phase. Provided effort is sufficiently valuable (specifically when $c < 1/2$), this silent phase is worth having.[16] Notice that this two-phase solution combines two of our earlier examples (silence and pronto) with the modification that due to backward compounding, the prize during the initial silent phase needs to grow to compensate for any rents earned by the agent during the pronto phase, which would otherwise lead the agent to pause during the initial phase.

As it turns out, the optimal policy need not be unique because the relaxed problem (P) may admit more than one optimal cutoff between phases and more than one optimal terminal date. Such multiplicity, however, is non-generic as it would not survive a slight perturbation of the function F.

The only remaining loose end is the possibility that the agent benefits from a global deviation, that is, one involving pauses during more than one instant. Fortunately, the

---

[14]If $T < t^*$, the optimal contract consists only of a silent phase that lasts until T and pays reward $c/f(T)$.

[15]From (Obj), the difference between virtual and actual effort is $p(t)(1 + \Phi(t)) - r(t)$ which, for any given $r(t)$, is strictly increasing in $p(t)$. Thus, silence gives the highest rents to the agent.

[16]The 1/2 appears because extending the length of the silent phase from dt to $2\,dt$ units of time means that a higher reward must be promised over the second such interval (owing to the agent's ignorance) and because of backward compounding, this higher reward must be promised over the first interval as well.

simple rewards in the theorem discourage all such deviations—and make it easy to check that this is the case. Observe that these rewards are non-increasing, and because both $f(t)$ and $\lambda(t)$ are weakly decreasing, they always grant the agent non-negative flow rents.[17] This makes a pause of any nature undesirable on two fronts: it causes the agent to miss out on a portion of such rents and shifts their success probability from the present to the future, where rewards are no greater.[18]

The optimal terminal date may be equal to T regardless of how large this exogenous deadline is. This occurs, for instance, when the hazard rate is constant as this allows the principal to extend phase 2 without giving up any rents. T^* would be strictly smaller than the deadline, in contrast, if that deadline was sufficiently large and the agent faced a Poisson good news experimentation technology. In this case, the hazard rate asymptotes to zero, and hence expanding phase 2 requires expanding its reward without bound, which then gets compounded backward.

While in our model F is exogenous, we can ask what type of distribution would be most profitable for the principal. Note that if F was Poisson with an arbitrarily low hazard rate, the principal would be able to achieve profits arbitrarily close to first-best $(1 - c)T$ simply by using an always-pronto policy with time-invariant reward c/λ, as this ensures that the agent works until T with arbitrarily high probability while earning no rents. This requires promising a very high reward, but paying it with very low probability.[19] Another distribution that implements the first best is the one that jumps from 0 to 1 at exactly T, as in this case a "forcing contract" that pays cT for total effort T grants zero rents.

Our assumption of a decreasing hazard rate may rule out technologies that are important in practice. One example is a hump-shaped hazard rate, which could easily arise if the agent's ability ("high" or "low") is initially unknown, only a high-ability agent can succeed, and conditional on high ability, the hazard rate *increases* with cumulative effort. This would represent, for example, a case where a high-ability associate learns on the job. Fortunately, Theorem 1 would still hold with a hump-shaped $\lambda(\cdot)$ provided that $\lambda(0) \geq \lambda(T^*)$ (and F is concave). This inequality ensures that the agent always earns non-negative rents throughout the pronto phase, and hence they have a dominant strategy to work throughout that phase. That they always find it optimal to work during the first phase as well follows from the assumption that f is non-increasing as this ensures that a global deviation not only causes the agent to lose flow rents, but also shifts their success probability from the present to the future, where rewards are no greater.

It is also worth noting how the model would differ if, by assumption, all pauses were permanent (i.e., the agent faced a stopping problem). In this simpler environment, the backward-compounding force would vanish and the principal would need only to ensure that the agent's continuation payoff is non-negative. Hence the principal would readily achieve the first best by remaining silent throughout, asking the agent to always work with probability 1, and promising them the zero-rent schedule $R(t) = c/f(t)$.[20]

---

[17] The agent's expected flow rent is $c[f(t)/\lambda(T^*) + F(t^*)f(t)/f(t^*) - 1]$ during phase 1 because they have no information, and $c[\lambda(t)/\lambda(T^*) - 1]$ during phase 2 (conditional on not having already succeeded) as they are fully informed.

[18] Note also that since there is no further reward forthcoming after the agent is advised to stop working, it is always optimal to follow such a recommendation.

[19] Observe that due to backward compounding, it does not suffice that the hazard rate is low during the last part of the horizon and high early on, as in that case all rewards (not only later ones) would need to be high, and hence the agent would be able to secure a high early reward with high probability.

[20] This observation also illustrates a key difference between our model and that of Ely and Szydlowski (2020). There, quitting is irreversible, but since the reward is exogenous, there is still a reason to inform the agent after a suitable delay.

In closing, the main practical implication of our baseline model is that in the presence of a coarse performance signal, hiding performance information from the agent may be an optimal way to motivate them. Consistent with this prescription, Maister (1993) observed in his analysis of professional partnerships (where the prospect of promotion is the key motivator at the beginning of a career) that partners often withhold performance information from their associates in order to prolong their trial phase (see pp. 170 and 173).[21] A formal empirical test could rely on both the timing of feedback and the prediction that promotion times will vary across associates—with bunching occurring for early promotions (at the end of the silent phase), and with those promoted sooner earning a greater prize, such as earning a larger raise or being assigned better opportunities post promotion.[22]

5. EXTENSIONS

Here we consider a series of extensions showing that a two-stage contract remains optimal under several variations of the model. We also consider a setting where the principal can choose to monitor effort directly, and show that whenever this monitoring is sufficiently costly, withholding information by means of an initial silent phase dominates full disclosure. For expositional simplicity, we focus on deterministic deadline contracts (where an unsuccessful agent is asked to work with probability 1 until some terminal time $\widetilde{T}$, and after this time all effort stops).

5.1. *Continuation Payoffs*

A restrictive feature of our baseline model is that as soon as a success is announced, the relationship between the principal and agent effectively ends. Here we generalize the model by allowing each of them to gain an exogenous continuation payoff once that announcement takes place, with the principal receiving π and the agent v (both time-invariant), in addition to the endogenous reward $R(t)$.

These continuation payoffs could come from several places. For instance, the agent may be able to use the news of their success to secure an outside opportunity. Alternatively, the two parties may continue their relationship and each benefit from it, for example, the principal may switch the agent to a new set of tasks with their own monitoring technology.[23] The principal's continuation payoff may also originate from success being intrinsically valuable.

This model is very similar to the baseline. To solve it, define the agent's reward $R(t)$ as including both the monetary reward and the continuation payoff v. Because delivering

---

[21]Maister argued (informally) that this practice can discourage strong performers, and therefore advised against it. Our model captures a version of that cost through the fact that hiding information requires raising the prize—but as we have shown, this practice can actually be beneficial if paired with the right rewards, and not abused.

[22]Other models can also account for a negative relationship between promotion time and reward, which has in fact been documented in some large firms; see Gibbs (1995) for a model and empirical evidence, and Ariga, Ohkusa, and Brunello (1999) for evidence concerning a large Japanese manufacturer. What is unique to our model is the combination of that prediction with the information policy and the variation in promotion times.

[23]For this model to capture the possibility of multiple successes, we need to assume that the agent is informed of a given success before they have a chance of attaining the next one. This could stem from the agent having to switch tasks in order to attempt the next success.

$R(t)$ to the agent costs only $R(t) - v$, the principal's payoff is now

$$\int_0^{\widetilde{T}} p(s)\,ds - \int_0^{\widetilde{T}} \big[R(s) - v\big]f(s)\,ds + \pi F(\widetilde{T}),$$

where $\widetilde{T}$ is the endogenous deadline, and where we have included the principal's payoff π conditional on the success arriving, which occurs with probability $F(\widetilde{T})$. Note that upon rearranging terms, this payoff is equal to the baseline payoff plus a new term $(v + \pi)F(\widetilde{T})$.

Since $R(t)$ represents the agent's total reward, Proposition 2 (which describes the least costly incentive compatible reward schedule) still applies. Hence the principal's payoff can be expressed as

$$\int_0^{\widetilde{T}} p(t)\big[1 - 2c - c\Phi(t)\big]\,dt + c\widetilde{T} + (v + \pi)F(\widetilde{T}),$$

which resembles the original objective (P) but with the extra term $(v + \pi)F(\widetilde{T})$. Since a more distant deadline $\widetilde{T}$ raises the probability that the continuation payoffs are obtained, this term gives the principal an incentive to extend the deadline.

Given that the work probability $p(t)$ affects only the first term in the objective—which is the same as in the baseline model—the optimal policy has the same two-phase structure as before, as characterized in Theorem 1; moreover, since that first term does not contain v or π, the duration of the silent phase, t^*, is unchanged. The only thing that changes is therefore that the principal stretches out the second phase.

Finally, notice that even though the total reward $R(t)$ must be positive to induce the agent to work, the monetary component of that reward could in principle be negative when v is large. That possibility is ruled out whenever $v \leq c/\lambda(T)$ as this guarantees that the reward for each of the two phases is at least v.

5.2. *More Informed Agent*

In our baseline model, the agent does not learn about a success unless the principal chooses to inform them. Here we instead allow the agent to learn about a success on their own, with some positive probability, independently of what the principal chooses to disclose. As we shall see, the optimal contract will be a two-phase policy that generalizes the baseline one by allowing for greater information for the agent during the initial phase.

Suppose in particular that the agent learns immediately of a success with some exogenous probability h even if the principal remains silent—or equivalently, that the principal must immediately inform them of a success with probability no lower than h, which means that $q(s|t) \leq 1 - h$.[24] Because the minimum reward schedule for a given schedule q is identical to that in the baseline model, Lemma 1 (which characterizes the principal's objective) remains valid. Thus, we can obtain the optimal policy in a similar way to before; namely, solve a modified relaxed problem that includes the new minimum disclosure requirement and then verify that the solution meets the remaining constraints.

---

[24]Because the agent stops working as soon as they observe a success, it is immaterial whether or not the principal knows what the agent has learned.

The relaxed problem is now

$$\sup_{\widetilde{T}\in[0,T],\,p(\cdot)} \int_0^{\widetilde{T}} p(t)\big[1 - 2c - c\Phi(t)\big]\,dt + c\widetilde{T} \tag{P$'$}$$

$$\text{s.t.} \quad 1 - F(t) \leq p(t) \leq 1 - hF(t) \quad \text{for all } t, \tag{Feas$'$}$$

where the new upper bound on $p(t)$ captures the fact that the agent cannot be asked to work at time t in the event that they have already succeeded by then and have learned about it, which, given the minimum disclosure requirement, occurs with probability no smaller than $hF(t)$.

Since the objective is identical to that in the baseline model when contracts have a deterministic deadline, and assuming again that $\widetilde{T} > t^*$, this relaxed problem has the following bang-bang solution:

$$p(t) = \begin{cases} 1 - hF(t) & \text{if } t \in [0, t^*], \\ 1 - F(t) & \text{if } t \in (t^*, \widetilde{T}]. \end{cases}$$

Note that the cutoff time t^* is the same as in the baseline model. The only difference relative to the baseline is therefore the lower work probability (i.e., greater information for the agent) during the first phase.

Upon following the same steps in the proof of Theorem 1 and assuming $\widetilde{T} > t^*$, we obtain the following:

COROLLARY 1: *In the extended model with a more informed agent, the optimal contract is similar to the baseline one except that during the first phase (which as before lasts until t^*), the agent is asked to work as long as they do not learn on their own that they have succeeded, and is offered a reward $c/\lambda(\widetilde{T}) + c(1 - h)F(t^*)/f(t^*)$ for a success during that phase, which is lower than before.*

Here the principal stays as close as possible to the original bang-bang solution by initially keeping the agent as uninformed as possible, and then proceeding to a pronto phase. Because the agent is now more informed during the first phase, the principal can get away with a lower reward during that phase—but since the agent also quits sooner in expectation, this is a worse outcome for the principal.

5.3. Upfront Effort Investment

Here we consider a simple learning-by-doing scenario where the agent must exert some minimum amount of cumulative effort S before they are able to succeed, and after that they succeed according to a non-increasing density f as in the baseline model. Thus, the hazard rate starts at zero and jumps up as soon as the required effort investment is complete. As we shall see, the optimal contract for this scenario will be very similar to the baseline one, differing at most in the lengths of the two phases and their associated rewards.

Let time run from $-S$ to $\widetilde{T}$, suppose S is small enough that the principal can induce the agent to work while still earning positive profits, and focus on contracts where the agent is asked to work continuously between $-S$ and 0 (which is without loss because any effort pauses before the investment is complete would merely waste time). So that

the agent is willing to make this investment, the principal must promise them an ex ante utility $U(-S)$ no lower than 0, or, equivalently, a continuation utility $U(0)$ no smaller than the total investment cost cS.[25]

The principal now maximizes the following weighted sum of profits and agent rents, where $\mu \geq 0$ denotes the appropriate multiplier for the new effort-investment constraint:

$$\int_0^{\widetilde{T}} p(s)(1 - \mu c)\, ds - \int_0^{\widetilde{T}} R(s)(1 - \mu)f(s)\, ds.$$

Notice that if $\mu = 1$, this objective would call for setting $\widetilde{T} = T$ and $p(s) = 1$ at all times, which, assuming T is sufficiently large, would impose losses on the principal. Moreover, if $\mu > 1$, the objective would call for making arbitrarily large transfers to the agent, which would also result in losses on the principal. Hence we focus on the case where $\mu < 1$.

As in the baseline model, the principal wishes to pay the smallest possible reward given the desired effort schedule, and hence adopts the minimal reward schedule in Proposition 2. Upon substituting for this minimal reward schedule, and manipulating terms, the objective becomes

$$\int_0^{\widetilde{T}} p(t)\big[1 - (2 - \mu)c - (1 - \mu)c\Phi(t)\big]\, dt + (1 - \mu)c\widetilde{T},$$

which differs from the baseline objective in (P) only in the modified coefficients for the effort cost c. It follows from the properties of Φ that the solution again begins with a silent phase (of possibly zero length) and ends with a pronto phase. The silent phase, however, will now have positive length for a (weakly) greater range of effort costs (i.e., $c < 1/(2 - \mu)$ rather than $c < 1/2$) and will grow with μ because keeping the agent in the dark longer has the added benefit of giving them greater rents. The pronto phase, in contrast, may potentially grow or shrink (or remain unaffected) depending on the details of F.

Notice that, beyond the current application, this same contract would be optimal any time the principal wishes to maximize a weighted sum of profits and agent rents provided the agent's weight is smaller than the principal's. Doing so would, for example, allow the principal to satisfy a standard ex ante participation constraint.

5.4. *Direct Monitoring*

An artificial feature of our baseline model is that despite effort always being valuable, it is impossible to monitor the agent beyond a single success. Here we consider a richer setting where at each instant the principal can monitor the agent both *indirectly* by watching for a success and promising a reward for it (as in the baseline model), and *directly* by paying a cost to observe the agent's effort, while offering them wages conditional on their chosen effort. One interpretation is that the "success" consists of the agent demonstrating

---

[25]The principal must also ensure that the agent does not withdraw effort during some subset of time before 0 and then makes up for the missing investment after that. But since the agent cannot earn any rewards before the investment is complete, such a deviation is at least weakly dominated by one where the agent works continuously until time zero and instead shirks, if they wish, after that, as this preserves the option to try to succeed as early as time 0.

a basic competence, after which it becomes more practical to directly observe effort than to look for more subtle signs of achievement.[26]

We assume that at every instant, the agent knows whether they are being directly monitored (e.g., because monitoring requires the presence of a supervisor), and this monitoring costs the principal $m > 0$ in flow terms, in addition to any wage. A justification for direct monitoring being more expensive than indirect monitoring is that it provides more accurate information. So that direct monitoring creates value, we assume that m is smaller than flow surplus $1 - c$. Provided the agent is working, their success arrives at a constant hazard rate $\lambda > 0$, and as in the baseline model, the principal observes this success for free.

Finding a fully optimal contract is challenging because the set of possible histories on which the principal can condition future monitoring (in additional to future messages and payments) is immense. We will be able to show, nonetheless, that whenever the monitoring cost is sufficiently large, the principal benefits from withholding information from the agent.

Specifically, we compare two policies, both of which induce the agent to work at all times. The first, which we term *always pronto*, entails keeping the agent fully informed of their success, and once a success arrives, directly monitoring them until T.[27] The second, which we call *early silence*, is like *always pronto* but with an initial silent phase and a higher reward for a success during that phase. This latter contract is very similar to the two-phase contract of the baseline model, but with a direct-monitoring phase added at the end.

These policies might represent (highly simplified) careers in professional partnerships with "lockstep" compensation, such as at Debevoise & Plimpton (an elite law firm). The announcement of a success would be the moment of promotion to partner, the reward would be the rent earned from the promotion, and the monitoring phase a reduced-form representation of a partner earning a profit share that depends only on seniority (the lockstep system) while being subject to a high-effort norm (see, e.g., Zehnder (2001)).[28]

PROPOSITION 4: *Consider the extended model with direct monitoring and a constant hazard rate of success. Suppose the monitoring cost $m > c$. Provided the silent phase is sufficiently short,* **early silence** *dominates* **always pronto**.[29]

Intuitively, the advantage of the *early silence* policy is that it saves on monitoring costs; its disadvantage is that due to the backward compounding of rewards caused by silence, it

---

[26]Some organizations rely heavily on peer monitoring among senior employees. We provide some examples below.

[27]Here the principal would offer reward c/λ for a success and would pay a flow wage equal to c for high effort (coupled with direct monitoring) after a success arrives. This policy is the best among all policies that keep the agent fully informed of their success because it grants zero rents while minimizing the amount of direct-monitoring time.

[28]Note that the lockstep system leads to close-to-flat compensation (i.e., close to independent of individual performance measures) whenever profits are split among a sizable number of partners (see Kandel and Lazear (1992)), as is arguably the case at large firms like Debevoise, with around 150 partners. A related example that goes beyond the professional services is the Mayo Clinic, which is famous for employing doctors who meet high standards in their own medical school (the "success") and then paying the doctors it hires a flat wage while using peer pressure to promote effort, which they call a "fishbowl effect" (the monitoring phase); see Berry and Seltman (2014).

[29]While finding the fully optimal contract is beyond our current reach, our conjecture is that it is either an *early silence* contract when $m > c$ or *always pronto* otherwise. This is because indirect monitoring is free and the agent is risk neutral, which suggests that direct monitoring should only be used after all indirect monitoring is over.

grants the agent rents. Recall from the baseline model that, at the margin, silence grants the agent c in rents, and these rents grow the longer the principal remains silent because rewards get compounded further. Thus, whenever m is greater than the marginal rent c, at least some amount of silence is desirable.

Key to this result is our assumption that monitoring is public. If the principal was instead able to credibly monitor the agent in secret, there would be an alternative way to save on monitoring costs: monitoring the agent's effort only randomly and promising a flow payment larger than c in the event that monitoring does take place and the agent happens to be working. Indeed, by monitoring effort with arbitrarily low probability and promising arbitrarily large rewards, the principal can approximate the first-best profit using this direct monitoring alone.

6. CONCLUSION

We have argued that when a principal uses a coarse performance measure, hiding information from the agent is expensive but may nonetheless be an optimal way to motivate them. We have uncovered a novel factor, which we term "backward-compounding" of rewards, that makes hiding information more costly if it happens farther in the future. When a single success is possible, this results in an optimal two-phase contract, with all ignorance frontloaded, that starts with a fully silent phase and ends with a phase of full transparency.

The key challenge was to find an optimal feedback policy among the vast set of potential ones, which we have done by showing that it suffices, under certain fairly general conditions, to focus on deterring instantaneous effort pauses. In future work, we shall attempt to apply the present methods to multi-agent (e.g., contest) settings, and to settings with more general monitoring technologies.

APPENDIX: PROOFS

A.1. *Proof of Proposition 1*

Suppose the agent obeys all recommendations before t and after $t + \Delta t$, but shirks in the interval in between. Then the total probability that they continue to work through $s \geq t + \Delta t$ equals

$$p\left(s|\omega_t^{\Delta t}\right) = r(s)\left[1 - F(s - \Delta t)\right] + \int_0^t r(u)f(u)q(s|u)\,du + \int_{t+\Delta t}^s r(u)f(u - \Delta t)q(s|u)\,du.$$

Next, we characterize the following limit:

$$
\begin{aligned}
\dot{p}(s|t) = \lim_{\Delta t \to 0} & \frac{p\left(s|\omega_t^{\Delta t}\right) - p(s)}{\Delta t} \\
= \lim_{\Delta t \to 0} & \left\{ r(s)\frac{F(s) - F(s - \Delta t)}{\Delta t} \right. \\
& - \frac{1}{\Delta t}\int_t^{t+\Delta t} r(u)f(u)q(s|u)\,du \\
& \left. - \int_{t+\Delta t}^s r(u)\frac{f(u) - f(u - \Delta t)}{\Delta t}q(s|u)\,du \right\},
\end{aligned}
$$

which represents the marginal change in the work probability $p(s)$ following an infinitesimal deviation at t. Since F is differentiable, the first term is $r(s)f(s)$. Applying L'Hopital's rule and the Lebesgue differentiation theorem, the limit of the second term (which exists almost everywhere) is simply the integrand evaluated at t, namely, $-r(t)f(t)q(s|t)$. As for the third term, note that

$$-\lim_{\Delta t \to 0} \int_{t+\Delta t}^{s} r(u) \frac{f(u) - f(u - \Delta t)}{\Delta t} q(s|u)\, du$$

$$= -\lim_{\Delta t \to 0} \int_{t}^{s} r(u) \frac{f(u) - f(u - \Delta t)}{\Delta t} q(s|u)\, du$$

$$= -\int_{t}^{s} r(u) \lim_{\Delta t \to 0} \frac{f(u) - f(u - \Delta t)}{\Delta t} q(s|u)\, du = -\int_{t}^{s} r(u)f'(u)q(s|u)\, du,$$

where the third line follows from dominated convergence. Therefore, almost everywhere,

$$\dot{p}(s|t) = r(s)f(s) - r(t)f(t)q(s|t) - \int_{t}^{s} r(u)f'(u)q(s|u)\, du. \tag{7}$$

Next, recall that incentive compatibility requires that

$$U(t + \Delta t) - U(t) \le \int_{t+\Delta t}^{T} r(s)\big[f(s) - f(s - \Delta t)\big]R(s)\, ds + c \times \int_{t+\Delta t}^{T} \big[p(s|\omega_t^{\Delta t}) - p(s)\big]\, ds.$$

Dividing both sides by Δt and taking the limit as $\Delta t \to 0$, we have

$$U'(t) \le \int_{t}^{T} r(s)f'(s)R(s)\, ds + c \lim_{\Delta t \to 0} \int_{t+\Delta t}^{T} \frac{p(s|\omega_t^{\Delta t}) - p(s)}{\Delta t}\, ds$$

$$\Rightarrow \quad cp(t) - r(t)f(t)R(t) \le \int_{t}^{T} r(s)f'(s)R(s)\, ds + c \int_{t}^{T} \dot{p}(s|t)\, ds. \tag{8}$$

Finally, using (7) and the definition $Q(t) := \int_{t}^{T} q(u|t)\, du$ and rearranging terms yields (LIC).   Q.E.D.

A.2. *Proof of Proposition 2*

Fix a recommendation policy $\{q(\cdot|\cdot), r(\cdot)\}$, and hence $p(\cdot)$ and $\dot{p}(\cdot, \cdot)$. The first step is to show that any reward schedule R that meets (LIC) must be non-negative. If R meets (LIC), then

$$r(t)\big[R(t) - cQ(t)\big]f(t) \ge cp(t) - \int_{t}^{T} r(s)\big[R(s) - cQ(s)\big]f'(s)\, ds - c \int_{t}^{T} r(s)f(s)\, ds$$

$$\ge cr(t)\big(1 - F(t)\big) - \int_{t}^{T} r(s)\big[R(s) - cQ(s)\big]f'(s)\, ds$$

$$- cr(t) \int_{t}^{T} f(s)\, ds$$

$$\ge -\int_{t}^{T} r(s)\big[R(s) - cQ(s)\big]f'(s)\, ds,$$

where the second line follows from the definition of p in (2) and the monotonicity of r.

We will show that this inequality implies $r(t)[R(t) - cQ(t)]f(t) \geq 0$ for all t. From this, it will follow that $R(t) \geq 0$ for all t because all of the other terms are non-negative. Without loss of generality, we can assume that $R(\cdot)$ is continuous almost everywhere because by defining $\widetilde{R}$ as follows:

$$r(t)[\widetilde{R}(t) - cQ(t)]f(t) = -\int_t^T r(s)[R(s) - cQ(s)]f'(s)\, ds,$$

we obtain an a.e. continuous function which is pointwise weakly smaller than R and satisfies the same inequality:

$$r(t)[\widetilde{R}(t) - cQ(t)]f(t) \geq -\int_t^T r(s)[\widetilde{R}(s) - cQ(s)]f'(s)\, ds.$$

Defining $B(t) := -\int_t^T r(s)[R(s) - cQ(s)]f'(s)\, ds$, we can rewrite the inequality as

$$B'(t) - d(t)B(t) \leq 0,$$

where $d(t) := f'(t)/f(t)$. (Recall that $f'(t) \leq 0$.)

Consider the function

$$z(t) = e^{-D(t)}B(t),$$

where D is an anti-derivative of d. The derivative of $z(t)$ is

$$z'(t) = e^{-D(t)}B'(t) - e^{-D(t)}d(t)B(t) = e^{-D(t)}\big[B'(t) - d(t)B(t)\big],$$

which is non-positive. Since $z(T) = 0$, we conclude that $z(t) \geq 0$ for all $t \leq T$ and thus, by the definition of z, also that $B(t) \geq 0$ for all $t \leq T$. Using our original inequality, we conclude

$$r(t)[R(t) - cQ(t)] \geq \frac{B(t)}{f(t)} \geq 0.$$

To complete the proof of Proposition 2, let $R(t)$ be any incentive compatible reward schedule. From (8), it satisfies

$$r(t)R(t) \geq \underbrace{\frac{1}{f(t)}\left[c \cdot p(t) - c\int_t^T \dot{p}(s|t)\, ds - \int_t^T r(s)f'(s)R(s)\, ds\right]}_{=:Z^1(t)}.$$

Because $f'(\cdot) \leq 0$ and $Z^1(t) \leq r(t)R(t)$, we have

$$Z^1(t) \geq \frac{1}{f(t)}\left[c \cdot p(t) - c\int_t^T \dot{p}(s|t)\, ds - \int_t^T Z^1(s)f'(s)\, ds\right];$$

in other words, $Z^1(t)/r(t)$ is an incentive compatible reward schedule and therefore $Z^1(t) \geq 0$ for all t. Continuing in this manner, define, for all $k \geq 2$, the function Z^k by

$$Z^k(t) = \frac{1}{f(t)}\left[c \cdot p(t) - c\int_t^T \dot{p}(s|t)\, ds - \int_t^T Z^{k-1}(s)f'(s)\, ds\right].$$

By induction, we have that $0 \leq Z^k(t) \leq Z^{k-1}(t)$ for all t. We have thus constructed a pointwise decreasing sequence of non-negative functions. Let Z be the pointwise limit. By the dominated convergence theorem, we have

$$Z(t) = \frac{1}{f(t)}\left[c \cdot p(t) - c\int_t^T \dot{p}(s|t)\,ds - \int_t^T Z(s)f'(s)\,ds\right].$$

Define a new reward schedule R^* by $r(t)R^*(t) = Z(t)$. Note that R^* satisfies (8) with equality, and moreover, it is weakly pointwise lower than the original R. We will next show that the schedule defined in (4) is in fact the unique schedule satisfying the incentive constraint with equality and it is therefore pointwise lower than any incentive compatible schedule.

Let $G(t) := \int_t^T r(s)f'(s)R^*(s)\,ds$ and $H(t) := c \cdot p(t) - c\int_t^T \dot{p}(s|t)\,ds$. Notice that $G'(t) = -r(t)f'(t)R^*(t)$ almost everywhere.[30] We can then rewrite (8) as

$$G'(t) = \frac{f'(t)}{f(t)}[G(t) - H(t)].$$

This is a linear differential equation with boundary condition $\lim_{T\to T} G(T) = 0$, and admits the following unique solution:

$$G(t) = f(t)\int_t^T \frac{f'(s)H(s)}{f(s)^2}\,ds.$$

Note that

$$r(t)R(t) = -\frac{G'(t)}{f'(t)} = \frac{H(t)}{f(t)} - \int_t^T \frac{f'(s)H(s)}{f(s)^2}\,ds. \qquad (9)$$

Letting

$$H_1(t) = \int_t^T r(s)f(s)\,ds,$$

$$H_2(t) = -r(t)f(t)\int_t^T q(s|t)\,ds, \quad \text{and}$$

$$H_3(t) = -\int_t^T \int_t^s r(u)f'(u)q(s|u)\,du\,ds,$$

and using (7), we have $H(t) = cp(t) - c[H_1(t) + H_2(t) + H_3(t)]$. Notice that $H_3'(t) = -f'(t)H_2(t)/f(t)$, and by integrating by parts, we have

$$\frac{H_3(t)}{f(t)} - \int_t^T \frac{f'(s)H_3(s)}{f(s)^2}\,ds = \frac{H_3(t)}{f(t)} + \int_t^T \left(\frac{1}{f(s)}\right)' H_3(s)\,ds = \int_t^T \frac{f'(s)}{f(s)^2}H_2(s)\,ds.$$

---

[30]We can apply the Fundamental Theorem of Calculus at almost every t. The reward schedule $R^*(\cdot)$ is continuous almost everywhere because it satisfies the incentive constraint with equality and the right-hand side of the incentive constraint is continuous almost everywhere.

Using integration by parts again, we have

$$\frac{H_1(t)}{f(t)} - \int_t^T \frac{f'(s)H_1(s)}{f(s)^2}\,ds = \int_t^T r(s)\,ds.$$

Then (9) can be rewritten as

$$r(t)R(t) = c\frac{p(t) - H_1(t) - H_2(t) - H_3(t)}{f(t)}$$

$$- c\int_t^T \frac{f'(s)}{f(s)^2}\big[p(s) - H_1(s) - H_2(s) - H_3(s)\big]\,ds$$

$$= c\bigg[\frac{p(t)}{f(t)} - \int_t^T \frac{f'(s)}{f(s)^2}p(s)\,ds - \int_t^T r(s)\,ds + r(t)\int_t^T q(s|t)\,ds\bigg]. \quad (10)$$

$$\textit{Q.E.D.}$$

A.3. *Proof of Lemma 1*

By substituting the minimal implementing reward schedule characterized in (4), we can rewrite the principal's objective as

$$\int_0^T (1-c)p(t) + cf(t)\int_t^T \frac{f'(s)}{f(s)^2}p(s)\,ds + cf(t)\int_t^T r(s)\,ds - cr(t)f(t)\int_t^T q(s|t)\,ds\,dt.$$

To simplify the double integrals, we use that $\int_0^T a(t)\int_t^T b(s)\,ds\,dt = \int_0^T b(t)\int_0^t a(s)\,ds\,dt$ for any integrable functions $a(t)$ and $b(t)$. In particular, we have

$$\int_0^T f(t)\int_t^T \frac{f'(s)}{f(s)^2}p(s)\,ds = \int_0^T \frac{f'(t)F(t)}{f(t)^2}p(t)\,dt = -\int_0^T p(t)\Phi(t)\,dt,$$

$$\int_0^T f(t)\int_t^T r(s)\,ds = \int_0^T r(t)F(t)\,dt, \quad \text{and}$$

$$\int_0^T r(t)f(t)\int_t^T q(s|t)\,ds\,dt = \int_0^T \int_0^t r(u)f(u)q(t|u)\,du\,dt$$

$$= \int_0^T p(t) - r(t) + r(t)F(t)\,dt,$$

where the last equality follows from (2). Using these identities, we can simplify the principal's objective as

$$\int_0^T p(t)\,dt - c\int_0^T p(t)\big(1 + \Phi(t)\big) - \big(r(t) - p(t)\big)\,dt. \qquad \textit{Q.E.D.}$$

A.4. *Proof of Proposition 3*

Consider the principal's relaxed problem given in (P). Fixing an $r(\cdot)$, we can maximize the objective pointwise for each t. Noting that $\Phi(t)$ is weakly increasing by assumption, it

follows that the bracketed expression is either always negative, or, once it becomes negative, it remains so throughout. Therefore, the optimal p is the non-increasing function

$$p(t) = \begin{cases} 1 & \text{if } t \le t^*, \\ r(t)[1 - F(t)] & \text{if } t > t^*, \end{cases}$$

where t^* has been defined as the smallest t such that $1 - 2c - c\Phi(t) \le 0$.

Turning to the choice of $r(\cdot)$, substituting the above $p(\cdot)$ into the objective, we have

$$\max_{0 \le r(\cdot) \le 1} \int_0^{t^*} \left[1 - 2c - c\Phi(t) + cr(t)\right] dt + \int_{t^*}^T r(t)\left\{[1 - F(t)][1 - 2c - c\Phi(t)] + c\right\} dt,$$

with the additional constraint that $r(\cdot)$ is non-increasing. Observe that for all $t \le t^*$, the objective increases in $r(t)$ (at rate c). For $t > t^*$, the objective increases in $r(t)$ if and only if $[1 - F(t)][1 - c - c\Phi(t)] + c \ge 0$. Because the objective is linear in $r(\cdot)$, it is without loss to set $r(t) \in \{0, 1\}$, and hence there exists an optimal $r(t)$ with the following form:[31]

$$r(t) = \begin{cases} 1 & \text{if } t \le T^*, \\ 0 & \text{otherwise}, \end{cases}$$

for some $T^* \in (t^*, T]$. This implies that $p(t) = 1$ for $t \le t^*$, $p(t) = 1 - F(t)$ for $t \in (t^*, T^*]$, and $p(t) = 0$ for $t > T^*$, which is non-increasing as desired. Finally, note that the optimal $r(\cdot)$ is equivalent to choosing a deterministic terminal date T^* after which the relationship is dissolved.   *Q.E.D.*

A.5. *Proof of Theorem 1*

We establish this theorem in three steps.

First, we note from (2) that the optimal $p(\cdot)$ characterized in Proposition 3 is implemented by the recommendation policy that sets $q(s|t) = 1$ for all $s \le t^*$ and any t, and otherwise sets $q(s|t) = 0$. In other words, it comprises a silent phase $[0, t^*]$, during which the principal asks the agent to work with probability 1 regardless of their success, and a pronto phase $(t^*, T^*]$, during which the agent is advised to quit upon succeeding. Finally, if they have not succeeded by T^*, they are advised to stop without reward. Note that $q(\cdot|t)$ is non-increasing for any t as required.

Second, we substitute the above recommendation policy into (4) to obtain the optimal reward schedule. For $t \in (t^*, T^*]$, we have

$$R(t) = c\left[\frac{1 - F(t)}{f(t)} - \int_t^{T^*} \frac{f'(s)}{f(s)^2}[1 - F(s)] ds - (T^* - t)\right] = \frac{c}{\lambda(T^*)},$$

---

[31]We can obtain this solution as follows. Begin by ignoring the monotonicity constraint and find a relaxed optimal r with values contained in $\{0, 1\}$. Next, suppose that relaxed solution violates the monotonicity constraint and select the earliest adjacent intervals such that $r = 0$ over the first one (denoted A) and $r = 1$ over the second (denoted B). Observe that there is a constrained solution such that r is constant over $A \cup B$ (as otherwise we could weakly raise the objective by setting r constant and equal to whichever value the alternative conjectured constrained optimal had taken at time $\sup A$). Now set r equal to either 0 or 1 over $A \cup B$ and whenever possible equal to 1 (which is again without loss). Finally, continue repeating the same process until r is monotone throughout.

where the last equality follows by integrating by parts and $\lambda(T^*) = f(T^*)/[1 - F(T^*)]$. Next, for $t \in [0, t^*]$, we have

$$
\begin{aligned}
R(t) &= c\left[\frac{1}{f(t)} - \int_t^{t^*} \frac{f'(s)}{f(s)^2}\,ds - \int_{t^*}^{T^*} \frac{f'(s)}{f(s)^2}\big[1 - F(s)\big]\,ds - (T^* - t^*)\right] \\
&= c\left[\frac{F(t^*)}{f(t^*)} + \frac{1}{\lambda(T^*)}\right],
\end{aligned}
$$

where the last equality again follows by integrating by parts and $\lambda(T^*) = f(T^*)/[1 - F(T^*)]$.

Finally, we verify that the above recommendation policy and reward schedule pair is globally incentive compatible. Following the recommendations during the second phase, when the agent is asked to work only if they have yet to succeed is a dominant strategy, because the prize is time-invariant and the agent earns rents (owing to the non-increasing hazard rate).[32] Turning to the first phase, if the agent shirks for Δ units of time and otherwise follows the recommendations, then their ex ante payoff is

$$
\begin{aligned}
\frac{\widetilde{U}(0, \Delta)}{c} &= \left[\frac{1}{\lambda(T^*)} + \frac{F(t^*)}{f(t^*)}\right]F(t^* - \Delta) + \frac{F(T^* - \Delta) - F(t^* - \Delta)}{\lambda(T^*)} \\
&\quad - (T^* - \Delta) + \int_{t^*}^{T^*} F(t - \Delta)\,dt.
\end{aligned}
$$

Using the concavity of F, it is straightforward to show that $\widetilde{U}(0, \Delta)$ decreases in Δ, and so the agent prefers to follow the recommendations throughout the first phase as well.

Q.E.D.

A.6. *Proof of Proposition 4*

Under the *always pronto* policy, incentivizing effort throughout the $[0, T]$ time interval costs the principal in expectation

$$
\int_0^T \lambda e^{-\lambda t}\left[\frac{c}{\lambda} + (c + m)(T - t)\right]dt = T(c + m) - \frac{m}{\lambda}(1 - e^{-\lambda T}). \tag{11}
$$

That is, when the agent succeeds (which they do with density $\lambda e^{-\lambda t}$), the principal pays them c/λ, and thereafter, the principal pays them a flow wage c while bearing flow monitoring cost m.

Next, we consider the *early silence* policy. Under this policy, the principal remains silent until some time $\bar{t}$, then reveals a success as it arrives, and thereafter directly monitors the agent at flow cost m. Incentive compatibility requires that during the monitoring phase the agent is paid a flow wage c provided they work (and 0 otherwise), and during the pronto phase they are paid c/λ upon succeeding. During the silent phase, for effort to be incentive compatible, per Theorem 1, the agent must be paid $c/\lambda + cF(\bar{t})/f(\bar{t}) = ce^{\lambda \bar{t}}/\lambda$

---

[32]Note that this is the only place in the proof where we use that the hazard rate $\lambda(t)$ is non-increasing. In fact, it suffices that $\lambda(T^*) \leq \lambda(0)$ as we explain following Theorem 1.

if they succeed before $\bar{t}$. Therefore, the principal's expected cost is

$$
\int_0^{\bar{t}} \lambda e^{-\lambda t} \left[\frac{c}{\lambda} e^{\lambda \bar{t}} + (c+m)(T-\bar{t}) \right] dt + \int_{\bar{t}}^T \lambda e^{-\lambda t} \left[\frac{c}{\lambda} + (c+m)(T-t) \right] dt
$$

$$
= (c+m)(T-\bar{t}) + \frac{c}{\lambda}(e^{\lambda \bar{t}} - 1) - \frac{m}{\lambda}(e^{-\lambda \bar{t}} - e^{-\lambda T}). \tag{12}
$$

Early silence dominates *always pronto* whenever (12) is smaller than (11). This is the case if and only if

$$
\Delta(\bar{t}) := \frac{c}{c+m} \times \frac{e^{\lambda \bar{t}} - 1}{\lambda \bar{t}} + \frac{m}{c+m} \times \frac{1 - e^{-\lambda \bar{t}}}{\lambda \bar{t}} < 1.
$$

Notice that $\Delta(\bar{t})$ is continuous, it converges to 1 as $\bar{t} \to 0$, and to ∞ as $\bar{t} \to \infty$. Letting $\alpha := c/(c+m)$, we have

$$
\Delta'(\bar{t}) = \frac{(\lambda \bar{t} + 1)(1-\alpha)e^{-\lambda \bar{t}} + \alpha(\lambda \bar{t} - 1)e^{\lambda \bar{t}} + 2\alpha - 1}{\lambda \bar{t}^2},
$$

and applying L'Hopital's rule, we get that $\Delta'(\bar{t})$ converges to $2\alpha - 1$ as $\bar{t} \to 0$. Therefore, $\Delta(\bar{t})$ is (strictly) smaller than 1 for $\bar{t}$ sufficiently small if and only if $2\alpha - 1 < 0$, or equivalently, $m > c$. Hence early silence dominates always pronto whenever $m > c$ and $\bar{t}$ is sufficiently short.   *Q.E.D.*

REFERENCES

ARIGA, KENN, YASUSHI OHKUSA, AND GIORGIO BRUNELLO (1999): "Fast Track: Is It in the Genes? The Promotion Policy of a Large Japanese Firm," *Journal of Economic Behavior & Organization*, 38 (4), 385–402. [0608]

BALL, IAN (2023): "Dynamic Information Provision: Rewarding the Past and Guiding the Future," *Econometrica*, 91 (4), 1363–1391. [0599]

BERRY, LEONARD L., AND KENT D. SELTMAN (2014): "The Enduring Culture of Mayo Clinic," *Mayo Clinic Proceedings*, 89, 144–147. [0612]

BONATTI, ALESSANDRO, AND JOHANNES HORNER (2011): "Collaborating," *American Economic Review*, 101 (2), 632–663. [0598]

ELY, JEFFREY C. (2017): "Beeps," *American Economic Review*, 107 (1), 31–53. [0598,0604]

ELY, JEFFREY C., AND MARTIN SZYDLOWSKI (2020): "Moving the Goalposts," *Journal of Political Economy*, 128 (2), 468–506. [0598,0599,0607]

ELY, JEFFREY C., GEORGE GEORGIADIS, SINA KHORASANI, AND LUIS RAYO (2023): "Optimal Feedback in Contests," *The Review of Economic Studies*, 90 (5), 2370–2394. [0599]

FENG, FELIX ZHIYU, CURTIS R. TAYLOR, MARK M. WESTERFIELD, AND FEIFAN ZHANG (2024): "Setbacks, Shutdowns, and Overruns," *Econometrica*, 92 (3), 815–847. [0601]

GEORGIADIS, GEORGE (2024): "Contracting With Moral Hazard," in *Elgar Encyclopedia on the Economics of Competition, Regulation and Antitrust*. Cheltenham, UK: Edward Elgar Publishing. [0598]

GEORGIADIS, GEORGE, AND BALAZS SZENTES (2020): "Optimal Monitoring Design," *Econometrica*, 88 (5), 2075–2107. [0599]

GIBBS, MICHAEL (1995): "Incentive Compensation in a Corporate Hierarchy," *Journal of Accounting and Economics*, 19 (2–3), 247–277. [0608]

GREEN, BRETT, AND CURTIS R. TAYLOR (2016): "Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects," *American Economic Review*, 106 (12), 3660–3699. [0598]

HALAC, MARINA, NAVIN KARTIK, AND QINGMIN LIU (2016): "Optimal Contracts for Experimentation," *The Review of Economic Studies*, 83 (3), 1040–1091. [0598,0600,0601]

HÖRNER, JOHANNES, AND NICOLAS S. LAMBERT (2021): "Motivational Ratings," *The Review of Economic Studies*, 88 (4), 1892–1935. [0599]

HÖRNER, JOHANNES, AND ANDRZEJ SKRZYPACZ (2017): "Learning, Experimentation and Information Design," *Advances in Economics and Econometrics*, 1, 63–98. [0598]

KAMENICA, EMIR, AND MATTHEW GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101 (6), 2590–2615. [0598]

KANDEL, EUGENE, AND EDWARD P. LAZEAR (1992): "Peer Pressure and Partnerships," *Journal of political Economy*, 100 (4), 801–817. [0612]

KAYA, AYÇA (2023): "Paying With Information," *Theoretical Economics*, 18 (2), 669–706. [0599]

KELLER, GODFREY, SVEN RADY, AND MARTIN CRIPPS (2005): "Strategic Experimentation With Exponential Bandits," *Econometrica*, 73 (1), 39–68. [0598]

MAISTER, DAVID (1993): "Quality Work Doesn't Mean Quality Service," in *Managing the Professional Service Firm*, Part II, 69–78. [0597,0608]

MASON, ROBIN, AND JUUSO VÄLIMÄKI (2015): "Getting It Done: Dynamic Incentives to Complete a Project," *Journal of the European Economic Association*, 13 (1), 62–97. [0598]

ORLOV, DMITRY (2022): "Frequent Monitoring in Dynamic Contracts," *Journal of Economic Theory*, 206, 105550. [0599]

ORLOV, DMITRY, ANDRZEJ SKRZYPACZ, AND PAVEL ZRYUMOV (2020): "Persuading the Principal to Wait," *Journal of Political Economy*, 128 (7), 2542–2578. [0599]

OSTRIZEK, FRANZ (2022): "Vague by Design: Performance Evaluation and Learning From Wages." [0599]

RAYO, LUIS, AND ILYA SEGAL (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118 (5), 949–987. [0598]

RENAULT, JÉRÔME, EILON SOLAN, AND NICOLAS VIEILLE (2017): "Optimal Dynamic Information Provision," *Games and Economic Behavior*, 104, 329–349. [0598]

SANNIKOV, YULIY (2008): "A Continuous-Time Version of the Principal-Agent Problem," *The Review of Economic Studies*, 75 (3), 957–984. [0598]

SMOLIN, ALEX (2021): "Dynamic Evaluation Design," *American Economic Journal: Microeconomics*, 13 (4), 300–331. [0599]

ZEHNDER, EGON (2001): "A Simpler Way to Pay," *Harvard Business Review*, 79 (4), 53–56. [0612]