Prizes and Feedback in Dynamic Contests*

Jeffrey C. Ely † George Georgiadis ‡ Luis Rayo §

May 7, 2026

Abstract

We study contests in which players privately choose when to exert costly effort toward a random success. A designer seeks to maximize total effort by choosing how to allocate a fixed prize and what feedback to provide during the contest. The optimal design is a "weighted egalitarian" contest, which resembles a raffle. The designer withholds feedback during an early phase, then reveals successes immediately thereafter. Players work continuously during the early phase, then continue only until success during the late one. Those succeeding early receive more raffle tickets, and hence a higher probability of winning, than those succeeding later.

Keywords: contests, principal-agent, moral hazard, information design

*We are grateful to Sina Moghadas Khorasani, as well as to participants at various seminars and conferences for helpful comments.

†Charles E. and Emma H. Morrison Professor of Economics, Northwestern University. Support from NSF grant SES-1851883 is gratefully acknowledged. jeff@jeffely.com

‡Associate Professor of Strategy, Northwestern Kellogg. g-georgiadis@kellogg.northwestern.edu

§Erwin P. Nemmers Professor of Strategy, Northwestern Kellogg. luis.rayo@northwestern.edu

1 Introduction

Managers, regulators, firms, and philanthropists often rely on contests when day-to-day effort is hard to monitor but breakthroughs matter. Innovation prizes, internal promotion tournaments, research competitions, sales contests, and time-bounded product-development races all share a common structure: players privately decide how long to work, success is uncertain, and the organizer must motivate effort using a limited prize. In such settings, the organizer has two natural instruments. The first is the rule by which the prize is allocated among those who succeed. The second is the information, or feedback, that players receive while the contest unfolds. This paper studies how these two instruments should be designed jointly.

Specifically, we consider a continuous-time contest among N symmetric players competing for a single indivisible prize. Each player privately chooses at every date whether to exert costly effort toward a binary success. Success can occur at most once for each player and is stochastic, with cumulative effort raising the probability of success. In the environments we study, additional effort becomes weakly less effective as accumulated effort grows. A contest consists of two objects: a reward rule, summarized by the expected probability that a player who succeeds at a given date receives the prize, and an information policy, represented as a recommendation policy that tells a player whether to continue working, depending on whether and when she has already succeeded. The contest designer wishes to maximize total effort subject to incentive compatibility and the fixed prize budget constraint.

The methodological difficulty is that the designer's choice of feedback and rewards affects incentives dynamically. A player who is asked to work at date t may be tempted not merely to stop forever, but to pause briefly and shift effort into the future. Such a pause changes the probability of success today, the distribution of future success times, and the likelihood that the player remains in states in which she is asked to work. We begin by adapting the local incentive approach of Ely, Georgiadis and Rayo (2025) to our contest environment. For any fixed recommendation policy, we derive the pointwise-minimal expected reward schedule that deters instantaneous pauses. That is, we compute the least expensive expected reward needed to make following the recommendation locally optimal. This step is agentwise: it depends on a player's own recommendation process, the success technology, and the expected reward from success, but not directly on the number of opponents or the ex post tie-breaking rule. This separation allows us to first compute the cheapest expected rewards that can support a desired effort path and only then ask whether those expected rewards can be implemented by an actual contest satisfying the single-prize budget.

Once the minimal local-incentive rewards are substituted into the budget constraint, the designer's problem becomes a choice over the induced effort path and the probability that an unsuccessful player is still asked to work. We solve this reduced problem using a Lagrangian relaxation. The resulting optimal effort path is bang-bang over two endogenous dates: up to a cutoff date τ, all players are asked to work continuously; from τ until a final deadline $\mathcal{T}$, only unsuccessful players are ask to work; after $\mathcal{T}$, no further effort is induced. We then verify that the contest implementing this effort path satisfies the original budget constraint and is globally incentive compatible, so that ruling out local pauses leads to the optimal contest.

We show that the optimal contest is a weighted egalitarian one. A player who succeeds early receives a larger expected share of the prize than a player who succeeds late. Players who do not succeed by $\mathcal{T}$ receive no prize. Equivalently, one can think of each early success as receiving weight one and each late success as receiving a smaller weight $\alpha \in [0, 1]$, with the prize assigned in proportion to these weights. The smaller weight on late successes reflects the information structure: after τ, players who are still asked to work know that they have not yet succeeded, and success immediately releases them from further effort, so a smaller expected prize is sufficient to motivate continuation until the deadline.

Although the formal implementation can be described as a weighted raffle, its practical interpretation can be broader. What matters for incentives is not that the organizer literally runs a lottery, but that agents understand how different success dates translate into different expected chances of receiving the prize. In many organizations, the final allocation of a prize, promotion, or award depends on additional considerations that are not controlled by the agents, not known to them in advance, or sufficiently opaque that they appear random from their perspective. The model's raffle can therefore be interpreted as a reduced-form description of such residual uncertainty. If early successes are placed in a more favorable category and the remaining selection depends on factors that agents view as beyond their control, or as effectively random from their point of view, then the agents face the same kind of expected reward schedule as in the weighted egalitarian contest.

A professional partnership provides a concrete example. During the period over which an associate is evaluated for promotion to partner, the firm may give extra credit to successes generated before an initial deadline, while still encouraging associates without such a success to keep pushing afterwards. The final decision may then depend on practice-area needs, client fit, partner support, or realized demand—factors associates cannot control or predict precisely.

Related Literature

Our paper contributes to several literatures. The first is the literature on agency problems under moral hazard, where a principal designs incentives to motivate unobservable effort. Canonical treatments include Holmström (1979) in a static environment and Sannikov (2008) in a dynamic one; see also Georgiadis (2024) for a review. Subsequent contributions study environments that are close to ours in their focus on stochastic breakthroughs. For example, Mason and Välimäki (2015) and Green and Taylor (2016) analyze settings with discrete stochastic successes, while Halac, Kartik and Liu (2016) study environments in which agents learn privately about project feasibility. Relative to this literature, we study moral hazard in a competitive environment with multiple agents and make feedback an explicit and fully-flexible design instrument. The designer in our model does not only choose how to reward success; she also chooses what agents learn, and hence when they are induced to continue working.

The paper is also related to the literature on contests. A central insight of this literature is that competition among agents can be used to strengthen incentives. Foundational contributions include Lazear and Rosen (1981), Green and Stokey (1983), Nalebuff and Stiglitz (1983), and Moldovanu and Sela (2001), who study how rank-based rewards and prize-sharing arrangements shape performance. More closely related are dynamic contest models for stochastic innovations, such as Taylor (1995) and Benkert and Letina (2020), which consider winner-takes-all prizes under fixed feedback rules. Halac, Kartik and Liu (2017) study experimentation contests with rank-monotonic prizes and symmetric deterministic disclosures, emphasizing the optimal timing of contest termination. Our contribution differs in that the principal simultaneously designs the feedback structure and the reward structure, and must do so subject to a fixed single-prize budget.

A third connection is to information design. The classic persuasion framework of Rayo and Segal (2010) and Kamenica and Gentzkow (2011) studies how a designer can influence actions by choosing what information agents receive. Ely (2017) and Renault, Solan and Vieille (2017) extend related ideas to dynamic environments, though primarily in settings with exogenous actions or single-agent decision problems. More recent work brings information design closer to moral hazard by studying how feedback affects an agent's incentive to continue supplying costly effort. This includes Ely and Szydlowski (2020), Orlov, Skrzypacz and Zryumov (2020), Ball (2023), Kaya (2023), Smolin (2021), and Ely, Georgiadis and Rayo (2025). Relative to these single-agent analyses, our paper emphasizes the new issues that arise under competition. Feedback must manage each player's intertemporal incentives, but rewards must also be implementable as shares of a common indivisible prize.

There is also a growing literature on information design in contests. In two-period settings, Lizzeri, Meyer and Persico (2005), Yildirim (2005), Aoyagi (2010), Ederer (2010), Goltsman and Mukherjee (2011), and Mihm and Schlapp (2019) study how disclosure rules affect competitive effort. These papers typically restrict attention to relatively simple disclosure policies. More recently, Ely et al. (2023) study optimal dynamic contests in continuous time and show that optimal feedback can take a more sophisticated form. Like the present paper, Ely et al. (2023) characterize fully optimal feedback and incentive schemes in an environment with binary stochastic performance. They do so under simplifying assumptions (specifically, a constant hazard rate of success and a sufficiently large number of contestants) that guarantee that every optimal feedback policy keeps agents fully informed of their own successes while extracting all rents. We relax those assumptions by studying environments in which the hazard rate of success can decline and the number of contestants is small enough that they optimally earn rents. This makes the design problem considerably more challenging because the optimal policy may involve delaying feedback even about an agent's own success.

The closest single-agent antecedent is Ely, Georgiadis and Rayo (2025), henceforth EGR. EGR show that strategic delays in feedback are costly but can increase effort, and their approach to local incentive compatibility is an important building block for our analysis. Delayed feedback can also be useful in our setting. The difference is that our environment contains multiple agents competing for a fixed prize. This introduces a budget constraint that is absent in EGR and creates a nontrivial implementation problem: the expected rewards required for incentive compatibility must be generated by an actual contest rule. The optimal implementation in our model therefore takes the form of a weighted egalitarian contest, or equivalently a weighted raffle, rather than a single-agent bonus scheme.

2 Model

We consider a contest among N players who compete for a single, indivisible prize with value normalized to 1. The contest unfolds over a fixed time horizon $[0, T]$, where T is large but finite. At every moment in time, each player privately decides whether to exert effort, incurring a constant flow cost $c \in (0, 1)$ whenever she does so.

Output takes the form of an all-or-nothing "success". Each player can succeed at most once, and their instantaneous probability of success is a function of the cumulative effort they have invested up to that point. Specifically, if a player has spent total effort e, then the probability of success is characterized by a strictly increasing function $F(e)$, satisfying $F(0) = 0$ and $F(T) \leq 1$, with a differentiable and bounded density f. We further assume that both $f(e)$ and the hazard rate $\lambda(e) := f(e)/(1 - F(e))$ are weakly decreasing in e.

Throughout, we assume that $\lambda(e)$ is weakly decreasing and strictly larger than nc for all e. (Ely et al. (2023) characterize optimal contests that maximize total effort and extract all rents for the simpler case where $\lambda(t) \leq Nc$ and constant. In that case, all players are kept fully informed about their own successes.)

A technical object central to our analysis is the function

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

whose significance will soon become clear. The designer's problem will be well-behaved when Φ is weakly increasing and $\Phi'(t)/\Phi(t) \geq \lambda(t)$ for all t, which we will assume.

Designing the contest requires specifying two key components: a prize allocation rule, which determines who the prize is awarded to—as a function of the timing of each player's success, and a feedback policy. We will encode the prize allocation rule by a function $w(t)$, which gives the probability that a player who succeeds at time t wins the prize.[1] We focus on symmetric contests in which this reward function is the same for all

[1] Strictly speaking, the contest specifies the probability of winning as a function of the players' success

players.

The information policy determines what each contestant learns about the status of the contest at every point in time. Any information policy can equivalently be represented as a recommendation policy, where each player is told when to continue working and when to stop. We model such a policy with two non-increasing, integrable functions: $q(t|s)$, the probability that a player is recommended to work at time t conditional on having succeeded at $s \leq t$; and $r(t)$, the probability that a player is recommended to work at time t if he has not yet succeeded. Thus, the probability that a player continues to work at least until time t is

$$p(t) = r(t)[1 - F(t)] + \int_0^t r(s)f(s)q(t|s)ds. \tag{2}$$

A player's expected payoff from following the recommendations is

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

where $r(s)w(s)f(s)$ is the expected flow reward at time s, and the second term represents the total expected cost of effort. The designer's objective is to choose a reward structure and a recommendation policy to maximize expected total effort:

$$N \int_0^T p(s)ds$$

subject to an incentive compatibility constraint that it is always optimal for each player to follow the recommendations, and the budget constraint

$$N \times \int_0^T r(t)w(t)f(t)dt \leq 1 - \left[1 - \int_0^T r(t)f(t)dt\right]^N, \tag{3}$$

times. The reward function $w(t)$ summarizes only the marginal (expected) probability that a player who succeeds at time t wins the prize; it does not uniquely determine the underlying contest rules, as the ex post allocation may also depend on the timing of other players' successes.

which stipulates that the total probability the prize is awarded does not exceed the total probability that at least one player succeeds.[2]

Observe that we restrict attention to policies in which a player's recommendation depends on her own success history, but not on the realized successes of her opponents. That this restriction is without loss can be seen from a simple coarsening argument. Start from any symmetric direct mechanism in which, at date t, player i's recommendation may depend both on her own history and on information about the success histories of other players. Fix player i and her own history. If the original mechanism is incentive compatible, then following the recommendation is optimal after every realization of the additional information about opponents. Now pool these opponent-history realizations and give the player only the same work/stop recommendation, with the same conditional probability given her own history. The gain from obeying the pooled recommendation is the conditional expectation of the gains from obeying the corresponding recommendations in the original mechanism. Since each of those gains is nonnegative, their conditional expectation is nonnegative as well.[3]

3 Local Incentives and Minimal Reward Schedule

This section describes the least costly expected reward schedule that supports a given recommendation policy. Specifically, we provide the pointwise-minimal expected reward needed to make the policy locally incentive compatible—that is, needed to deter instantaneous effort pauses—and later show that this expected reward can be implemented by an

[2]More precisely, the left-hand side is the probability that the prize is allocated and a success has occurred. This cannot exceed the total probability that a success occurs. Allocating the prize to a player who has not succeeded would undermine incentives and reduce the principal's payoff.

[3]This restriction concerns only real-time recommendations. The final allocation rule may still depend on the entire vector of success times. We summarize such an allocation rule by the marginal expected reward $w(t)$, the probability that a player who succeeds at date t receives the prize, where the expectation is taken over the success times of the other players and any randomization in the allocation rule. Since players are risk neutral and observe only the recommendation process described above, their incentives depend on the allocation rule only through these conditional expected rewards.

actual contest rule satisfying the prize budget, and that this contest is globally incentive compatible. The schedule in question is obtained from Ely, Georgiadis and Rayo (2025). The only difference is one of interpretation: here the reward is not a transfer chosen freely by a single-agent principal, but the expected share of a fixed contest prize that a player receives upon success.

Fix a recommendation policy $\big(q(\cdot \mid \cdot), r(\cdot)\big)$. Recall that $r(t)$ is the probability that a player who has not yet succeeded is recommended to work at time t, while $q(u \mid t)$ is the probability that a player who succeeded at date t is still recommended to work at the later date $u \geq t$. Define

$$Q(t) \;:=\; \int_t^T q(u \mid t)\, du.$$

Thus, $Q(t)$ is the expected amount of *future* work that the recommendation policy asks from a player who succeeds at date t. This object is central because a success need not immediately end the player's effort obligation.

Recall that $p(t)$ is the ex-ante probability that the player is working at date t, as defined in (2). We write the minimal locally incentive compatible reward schedule in the following form:

$$r(t)w^{\min}(t) \;=\; \underbrace{c\,\frac{p(t)}{f(t)}}_{\text{myopic reward}} \;+\; \underbrace{\Gamma(t; p, r, Q)}_{\text{dynamic adjustment}} \;. \tag{4}$$

The first term is the reward that would be needed if there were no future incentive effects. To see this, suppose date t were the last relevant instant. The agent bears an expected flow cost $c\,p(t)$ from working at date t. The expected flow reward from success at t is $r(t)w(t)f(t)$. Equating the two gives $r(t)w(t) \;=\; c\,p(t)/f(t)$. In a dynamic problem, however, a brief pause at date t changes the probability of future success, the probability that the player will remain in states in which she is asked to work, and the amount of work required after success. The term $\Gamma(t; p, r, Q)$ collects these dynamic effects.

The following proposition is a restatement, in our notation, of the minimal-reward formula in Ely, Georgiadis and Rayo (2025).

Proposition 1 (Dynamic adjustment in the minimal reward schedule)**.** *Fix a recommendation policy $\big(q(\cdot \mid \cdot), r(\cdot)\big)$ and let $p(\cdot)$ and $Q(\cdot)$ be defined as above. Among all reward schedules that deter instantaneous pauses, there is a unique pointwise-minimal schedule. It is given by* (4)*, with*

$$\Gamma(t; p, r, Q) = c\left[-\int_t^T \frac{f'(s)}{f(s)^2}\, p(s)\, ds \;-\; \int_t^T r(s)\, ds \;+\; r(t)Q(t)\right]. \tag{5}$$

To interpret Γ, consider a player who is supposed to work at date t and asks whether it is profitable to pause for an instant. Such a pause does not simply remove the current cost of effort. It also reshuffles the player's future prospects. The three terms in (5) correspond to the three resulting continuation effects.

First,

$$-c\int_t^T \frac{f'(s)}{f(s)^2}\, p(s)\, ds$$

is a *future-productivity effect*. Because f is weakly decreasing, we have $f'(s) \leq 0$, so this term is weakly positive. If the player pauses now, she reaches future dates with slightly less accumulated effort. Under declining density, this makes future effort more productive: the probability density of future success is higher than it would have been had she worked continuously. This gives the player an incentive to delay effort. The designer must therefore raise the reward for success at t to offset the attraction of shifting success probability into the future. The more important future work is, as measured by $p(s)$, and the more sharply the success density declines, as measured by $-f'(s)$, the larger this upward adjustment must be.

Second,

$$-c\int_t^T r(s)\, ds$$

is a *future-work-cost credit*. A pause at date t increases the chance that the player has not yet succeeded at later dates. When the player remains unsuccessful, the policy may continue to ask her to work, which is costly. This future cost makes pausing less attractive. Hence this term lowers the reward required at date t.

Third,

$$c\,r(t)Q(t)$$

is a *post-success-work premium*. If the player succeeds at date t, the recommendation policy may nevertheless keep her working in the future. The expected amount of such work is $Q(t)$. Because this work is costly, a success at t is less valuable to the player. The reward for success at t must therefore include compensation for that expected continuation cost. The factor $r(t)$ appears because this adjustment is relevant on the histories in which an as-yet-unsuccessful player is actually recommended to work at date t and can therefore generate a success at that date.

The minimal schedule is pinned down backward from the end of the contest. Local incentive compatibility at date t depends on rewards and recommendations from date t onward. Once the rewards required after t are fixed, the reward at t must be just high enough to remove the gain from a momentary pause. Setting the local constraint to bind at every date therefore gives the cheapest schedule. Any extra reward at a later date would make future success more attractive, which would tighten the no-pause constraint at earlier dates and force still higher earlier rewards.

Finally, note that this result is agentwise. For a fixed recommendation policy, the local incentive calculation depends on the player's own expected reward schedule and on the success technology (F, f), but not directly on the number of opponents or on the particular ex post tie-breaking rule used by the contest. This separation is what allows us to proceed in two steps. First, compute the minimal expected rewards required to support a desired effort path. Second, show how a weighted egalitarian contest can implement exactly those expected rewards while respecting the single-prize budget constraint.

4 Weighted Egalitarian Contest

In this section, we introduce and characterize a specific family of contests, termed *weighted egalitarian*, which we will later demonstrate to be optimal. Each contest within this family is defined by three parameters: two dates, τ and $\mathcal{T} \geq \tau$, and a weighting coefficient $\alpha \in [0,1]$. We partition the contest duration into two distinct phases: the interval $[0,\tau]$, termed the *early period*, and the subsequent interval $(\tau, \mathcal{T}]$, termed the *late period*.

Allocation rule. Each player is assigned a weight z_i, determining the probability of receiving the prize according to $z_i / \sum_j z_j$. The player's weight depends exclusively on the timing of their success. Specifically, a player succeeding during the early period receives a weight $z_i = 1$, whereas a player succeeding during the late period receives a reduced weight $z_i = \alpha$. Players who do not succeed by the contest deadline $\mathcal{T}$ receive no weight and thus no prize. By construction, the sum of prize probabilities equals one, ensuring that the budget constraint is inherently satisfied.

Thus, given parameters $(\tau, \mathcal{T}, \alpha)$, each player i faces the reward function

$$w^{weg}(t; \alpha, \tau, \mathcal{T}) = \begin{cases} \mathbb{E}\left[\frac{1}{1+\underline{M}+\alpha\overline{M}}\right] & \text{if } t < \tau \\ \mathbb{E}\left[\frac{\alpha}{\alpha+\underline{M}+\alpha\overline{M}}\right] & \text{if } \tau \leq t \leq \mathcal{T} \\ 0 & \text{otherwise,} \end{cases}$$

where $\underline{M}$ and $\overline{M}$ denotes the number of players *other than i* who succeed during the early and the late period, respectively.[4]

Recommendation policy. Players are recommended to exert continuous effort until at least τ. From τ until the contest concludes at $\mathcal{T}$, only participants who have not yet succeeded are advised to continue working. Players succeeding within the late period

[4]Assuming continuous effort until success or contest end at $\mathcal{T}$, we have $\underline{M} \sim Binom(n-1, F(\tau))$, and conditional on $\underline{M} = m$, $\overline{M} \sim Binom(n-m-1, F(\mathcal{T}) - F(\tau))$.

are immediately informed and advised to stop working. Formally:

$$q(s|t) = \begin{cases} 1 & \text{if } s < \tau \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad r(t) = \begin{cases} 1 & \text{if } t < \mathcal{T} \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Proposition 1 established a unique, minimal incentive compatible reward schedule for any given recommendation policy. Applying this lemma to (6), the minimal reward function is:

$$w^{min}(t; \tau, \mathcal{T}) = \begin{cases} \frac{c}{\lambda(\mathcal{T})} + \frac{cF(\tau)}{f(\tau)} & \text{if } t < \tau \\ \frac{c}{\lambda(\mathcal{T})} & \text{if } \tau \leq t \leq \mathcal{T} \end{cases}$$

The following lemma shows that for any contest duration $\mathcal{T}$, there exists a weighted egalitarian contest with this duration such that the local incentive constraint binds at all times.

Lemma 1. *For any $\mathcal{T} \leq T$, there exists a unique $\tau(\mathcal{T})$ and $\alpha(\mathcal{T}) \in [0,1]$ such that*

$$w^{weg}(t; \alpha(\mathcal{T}), \tau(\mathcal{T}), \mathcal{T}) \equiv w^{min}(t; \tau(\mathcal{T}), \mathcal{T}).$$

To understand this result intuitively, fix $\mathcal{T}$ and recall that the lowest (locally) incentive compatible reward has a baseline $c/\lambda(\mathcal{T})$ for late successes and an early premium $cF(\tau)/f(\tau)$ for successes before τ. When τ is small, the required premium is small, which pushes the weighted egalitarian rule toward $\alpha \approx 1$ so that early and late expected shares are almost equal. In that case, however, late successes earn *above* the baseline: with almost no early competitors ($\underline{M} \approx 0$), a late success' expected share is close to $\mathbb{E}[1/(1 + \overline{M})]$, which exceeds $c/\lambda(\mathcal{T})$ per our assumption that $1/N > c/\lambda(\mathcal{T})$. At the other extreme, when τ is large the required premium is large; achieving it requires α to be small (penalizing late successes), which then drives their expected share *below* the baseline. Since for fixed τ the late share is strictly increasing in α, and the required premium $cF(\tau)/f(\tau)$ is strictly increasing in τ, there is a unique cutoff τ at which the corresponding incentive compatible α (delivering the exact early premium) simultaneously delivers the correct late baseline reward.

5 Optimal Contest

We now present the central result of our analysis, which identifies the weighted egalitarian contest as optimal:

Theorem 1. *There exists a unique $\mathcal{T}^*$ such that the "weighted egalitarian" contest with duration $\mathcal{T}^*$ and parameters $\tau(\mathcal{T}^*)$ and $\alpha(\mathcal{T}^*)$ is optimal.*

By way of intuition, note that the contest implements a *bang–bang* effort path: $p(t)$ is held at its maximal feasible level up to a cutoff τ, and at its minimal feasible level thereafter. Because the designer's objective is linear in the effort path, extreme points of the feasible set are optimal for any fixed budget, hence the corners. Moreover, early effort is *less expensive* than late effort: raising $p(t)$ at a late date t forces the designer to grant rents at t and, due to the local incentive constraints, to raise rewards at *all earlier dates* to keep agents from briefly pausing—ratcheting up the entire schedule. By contrast, boosting $p(t)$ early requires only a small incremental "early rent." It follows that any partial disclosure before τ is inefficient (it lets some early winners quit when effort is cheapest), and any opacity after τ is inefficient (it keeps winners working when extra effort is most expensive). The weighted-egalitarian design complements this structure by implementing the bang–bang path with the smallest possible incentive-compatible rewards while using the entire prize budget (per Lemma 3), thus maximally translating prize dollars into effort.

For a heuristic derivation, begin by conjecturing that the weighted egalitarian contest derived in the previous section, with some parameters $(\tau^*, \mathcal{T}^*, \alpha^*)$, is optimal. To verify this conjecture, plug the minimal local-IC rewards (Proposition 1) into the budget constraint, assign multiplier η to that constraint, and *choose* η so that the effort path that maximizes a relaxed (Lagrangian) problem coincides with the conjectured weighted-egalitarian effort path. Finally, implement this effort path with the weighted-egalitarian rewards and show that it attains the corresponding dual bound.

Upon substituting for the minimal rewards, the Lagrangian is

$$L(\eta) = \max_{r,\, p \text{ decr.}} \underbrace{\int_0^T p(t)\Big[1 - \eta c\big(2 + \Phi(t)\big)\Big]dt}_{\text{value of effort path}} + \underbrace{\eta c \int_0^T r(t)\, dt}_{\text{eligibility mass}} + \underbrace{\frac{\eta}{N}\Big(1 - \big[1 - \int_0^T r(t) f(t)\, dt\big]^N\Big)}_{\text{prize release}},$$

where p must satisfy $r(t)[1 - F(t)] \leq p(t) \leq 1$. The coefficient on $p(t)$ is $1 - \eta c[2 + \Phi(t)]$, which is strictly decreasing in t. We *calibrate the multiplier to our conjectured* cutoff by setting

$$\eta^* \;:=\; \frac{1}{c\,[\,2 + \Phi(\tau^*)\,]},$$

so that $1 - \eta^* c[2 + \Phi(t)]$ switches sign exactly at $t = \tau^*$: it is positive for $t < \tau^*$, zero at τ^*, and negative thereafter. At this η^*, the Lagrangian (pointwise in t) is maximized by the bang–bang effort path

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

i.e., "work for sure" early and "work only if not yet successful" late—precisely the weighted egalitarian shape.

With p thus pinned down, the Lagrangian reduces to choosing r. Writing $z := \int_0^T r(t) f(t)\, dt$ and

$$G(t) := [1 - F(t)]\big[1 - \eta^* c(2 + \Phi(t))\big] + \eta^* c,$$

the problem becomes

$$\max_{r \text{ decr.}} \int_{\tau^*}^{T} r(t)G(t)\,dt \;+\; \frac{\eta^*}{N}\Big(1-[1-z]^N\Big) \;+\; \text{constant.}$$

The second term is increasing and strictly concave in z; under our assumptions $G(t)$ is strictly decreasing in t. A standard "left-shift" argument then implies the unique maximizer is a *cutoff* $r(t) = \mathbf{1}\{t \le \mathcal{T}^*\}$, with $\mathcal{T}^*$ characterized by the first-order condition $G(\mathcal{T}^*) + \eta^*[1-F(\mathcal{T}^*)]^{N-1} f(\mathcal{T}^*) = 0$. Evaluating the minimal local-IC rewards on this (p, r) yields a two-level schedule (late baseline $c/\lambda(\mathcal{T}^*)$ plus early premium $cF(\tau^*)/f(\tau^*)$). By Lemma 1, there is a unique α^* so that the weighted egalitarian rule with $(\tau^*, \mathcal{T}^*, \alpha^*)$ matches these expected rewards pointwise. Because this contest is feasible and attains the dual value $L(\eta^*)$, the dual upper bound is achieved by a primal feasible policy—hence there is no duality gap. Complementary slackness is straightforward: the coefficient on $p(t)$ is zero at τ^* (so p is at its upper/lower bound on either side), while the derivative with respect to r is zero at $\mathcal{T}^*$ and has the appropriate sign elsewhere. Strict monotonicity (of Φ, G) and strict concavity (in z) deliver uniqueness of $(\tau^*, \mathcal{T}^*, \alpha^*)$.

The originality of this method lies in calibrating the shadow price to the conjectured cutoff: we do not solve the primal directly; instead, we pick η^* using the weighted egalitarian candidate, make the relaxed problem deliver exactly that cutoff structure, and then implement it with a minimal-rent raffle that saturates the dual bound.

5.1 Proof of Theorem 1

The proof proceeds by first characterizing the optimal parameters $\mathcal{T}^*$, $\tau^* = \tau(\mathcal{T}^*)$, and $\alpha^* = \alpha(\mathcal{T}^*)$. We then formulate the designer's relaxed problem, where she selects a recommendation policy to maximize her objective subject to the budget constraint, using the minimal reward function from Proposition 1. After considering a Lagrangian relaxation of this problem, we characterize the optimal solution for a carefully selected dual multiplier. Next, we show that this solution can be precisely implemented by a weighted egalitarian contest with the identified parameters $\{\mathcal{T}^*, \tau^*, \alpha^*\}$, confirming its feasibility for the designer's problem and therefore establishing its optimality. Finally, we verify that the derived contest is globally incentive compatible.

Step 1: The following lemma characterizes the parameters of the optimal contest.

Lemma 2. *There exists unique parameters $\alpha^* \in (0,1)$, $\mathcal{T}^* > 0$, and $\tau^* \in (0, \mathcal{T}^*)$ such that $w^{weg}(t; \alpha^*, \tau^*, \mathcal{T}^*) = w^{min}(t; \tau^*, \mathcal{T}^*)$, and*

$$\frac{2F(\mathcal{T}^*) - 1 - \Phi(\mathcal{T}^*)\big[1 - F(\mathcal{T}^*)\big]}{2 + \Phi(\tau^*)} + 1 - F(\mathcal{T}^*) + \left(1 - \big[1 - F(\mathcal{T}^*)\big]^{N-1}\right)\frac{f(\mathcal{T}^*)}{2 + \Phi(\tau^*)} = 0. \tag{7}$$

Step 2: Using Proposition 1, the definition of $\Phi(t)$, and integration by parts, the budget constraint given in (3) can be simplified as follows:

$$c \times \left(\int_0^T [2 + \Phi(t)] p(t) - r(t) dt \right) \leq \frac{1 - \left[1 - \int_0^T r(t) f(t) dt\right]^N}{N} \tag{BC}$$

Notice that the budget constraint is now expressed solely as a function of $p(t)$ and $r(t)$, and all references to $q(s|t)$ have conveniently disappeared.

Consider the following program:

$$\max_{r(\cdot),\, p(\cdot) \text{ decreasing}} \int_0^T p(t) dt \quad \text{s.t.} \quad (BC) \quad \text{and} \quad 0 \leq r(t)[1 - F(t)] \leq p(t) \leq 1. \tag{P}$$

That is, we seek decreasing functions $p(t)$ and $r(t)$ to maximize the designer's objective subject to the budget constraint, as well as the lower and upper bound on $p(t)$ that follows from (2).

Step 3: Dualizing (BC), we have the Lagrangian

$$L(\eta) = \max_{r,p \text{ decr.}} \int_0^T p(t)\Big(1 - \eta c[2 + \Phi(t)]\Big) dt + \eta c \int_0^T r(t) dt + \frac{\eta}{N}\left(1 - \left[1 - \int_0^T r(t) f(t) dt\right]^N\right)$$

Define

$$\eta^* := \frac{1}{c[2+\Phi(\tau^*)]},$$

where τ^* was implicitly characterized in Lemma 2, and consider $L(\eta^*)$. We maximize the Lagrangian in two steps: First, we maximize the objective with respect to p for fixed r. Since Φ is increasing and the objective is linear in $p(t)$, it is optimal to set $p(t) = 1$ for all $t \leq \tau^*$, and otherwise set $p(t)$ to its lower bound, $r(t)[1 - F(t)]$.[5]

Substituting this $p(t)$, the objective becomes

$$\max_{r(t)\text{ decr.}} \int_0^{\tau^*} (1 - \eta^* c[2+\Phi(t)])\, dt \tag{8}$$
$$+ \int_{\tau^*}^{T} r(t) \left\{[1-F(t)]\,(1-\eta^* c[2+\Phi(t)]) + \eta^* c\right\} dt + \frac{\eta^*}{N}\Big(1 - \Big[1 - \int_0^T r(t) f(t) dt\Big]^N\Big).$$

The following lemma shows that this objective is maximized when $r(t)$ is bang-bang, jumping from 1 to 0 at $T = \mathcal{T}^*$.

Lemma 3. *The objective in (8) is maximized by setting $r(t) = 1$ for all $t \leq \mathcal{T}^*$, and $r(t) = 0$ otherwise.*

It follows that when $\eta = \eta^*$ the Lagrangian is maximized when

$$p(t) = \begin{cases} 1 & \text{if } t \leq \tau^* \\ 1 - F(t) & \text{if } \tau^* < t \leq \mathcal{T}^* \\ 0 & \text{otherwise,} \end{cases}, \quad r(t) = \begin{cases} 1 & \text{if } t \leq \mathcal{T}^* \\ 0 & \text{otherwise.} \end{cases} \quad \text{and} \quad q(s|t) = \begin{cases} 1 & \text{if } t \leq \tau^* \\ 0 & \text{otherwise,} \end{cases}$$

Step 4: Let Π^* denote the designer's objective at the optimal solution. Since the Lagrangian represents a relaxation, we have $L(\eta) \geq \Pi^*$ for all η. By construction, the pointwise-smallest reward function that is locally incentive compatible is $w^{min}(t; \tau^*, \mathcal{T}^*)$, which by Lemma 1, is implemented by a weighted egalitarian contest with weight α^*.

[5] Note that the multiplier η^* was chosen precisely so that the term that multiplies $p(t)$ turns negative at τ^*.

Notice that the weighted egalitarian contest satisfies the budget constraint; hence the solution associated with η^* is feasible for (P), and so $L(\eta^*) \leq \Pi^*$. Therefore, we have $L(\eta^*) = \Pi^*$, which implies the weighted egalitarian contest with parameters $(\alpha^*, \tau^*, \mathcal{T}^*)$ solves the designer's problem.

Step 5: We conclude the proof by ruling out global deviations. As in Ely, Georgiadis and Rayo (2025), due to the simple reward structure, and because rewards are non-increasing and both $f(t)$ and $\lambda(t)$ are weakly decreasing, the players earn non-negative flow rents at every moment. Any pause is therefore doubly costly: it causes the player to forgo a portion of these rents and shifts the probability of success toward dates at which the reward is no higher. $\square$

6 Discussion

This paper analyzed the optimal design of incentives and feedback in dynamic contests where agents exert unobservable effort and compete for an indivisible prize under noisy, all-or-nothing success signals. We introduced a family of *weighted egalitarian* contests, consisting of a two-phase structure: an early period in which successes are rewarded more heavily, and a late period where prize weights are scaled down. We showed that, for appropriately chosen parameters, there exists a unique contest within this class that maximizes total effort, subject to incentive and budget feasibility. The optimal policy is implemented through a minimal feedback regime: contestants are provided with no information about their own or others' success until a designated cutoff, after which they receive immediate feedback upon succeeding and are instructed to stop exerting effort. This design achieves a balance between incentivizing early effort—by rewarding prompt success—and preserving resources to motivate laggards in the latter phase of the contest. Our approach leverages a dynamic programming argument and duality-based relaxation.

References

Aoyagi, Masaki. 2010. "Information feedback in a dynamic tournament." *Games and Economic Behavior*, 70(2): 242–260.

Ball, Ian. 2023. "Dynamic information provision: Rewarding the past and guiding the future." *Econometrica*, 91(4): 1363–1391.

Benkert, Jean-Michel, and Igor Letina. 2020. "Designing dynamic research contests." *American Economic Journal: Microeconomics*, 12(4): 270–89.

Ederer, Florian. 2010. "Feedback and Motivation in Dynamic Tournaments." *Journal of Economics & Management Strategy*, 19(3): 733–769.

Ely, Jeffrey C. 2017. "Beeps." *American Economic Review*, 107(1): 31–53.

Ely, Jeffrey C, and Martin Szydlowski. 2020. "Moving the goalposts." *Journal of Political Economy*, 128(2): 468–506.

Ely, Jeffrey C, George Georgiadis, and Luis Rayo. 2025. "Feedback design in dynamic moral hazard." *Econometrica*, 93(2): 597–621.

Ely, Jeffrey C, George Georgiadis, Sina Khorasani, and Luis Rayo. 2023. "Optimal Feedback in Contests." *The Review of Economic Studies*, 90(5): 2370–2394.

Georgiadis, George. 2024. "Contracting with moral hazard." *Elgar Encyclopedia on the Economics of Competition, Regulation and Antitrust*, 24–37.

Goltsman, Maria, and Arijit Mukherjee. 2011. "Interim Performance Feedback in Multistage Tournaments: The Optimality of Partial Disclosure." *Journal of Labor Economics*, 29(2): 229–265.

Green, Brett, and Curtis R. Taylor. 2016. "Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects." *American Economic Review*, 106(12): 3660–3699.

Green, Jerry R, and Nancy L Stokey. 1983. "A comparison of tournaments and contracts." *Journal of Political Economy*, 91(3): 349–364.

Halac, Marina, Navin Kartik, and Qingmin Liu. 2016. "Optimal Contracts for Experimentation." *The Review of Economic Studies*, 83(3): 1040–1091.

Halac, Marina, Navin Kartik, and Qingmin Liu. 2017. "Contests for Experimentation." *Journal of Political Economy*, 125(5): 1523–1569.

Holmström, Bengt. 1979. "Moral hazard and observability." *The Bell journal of economics*, 74–91.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review*, 101(6): 2590–2615.

Kaya, Ayça. 2023. "Paying with information." *Theoretical Economics*, 18(2): 669–706.

Lazear, Edward P, and Sherwin Rosen. 1981. "Rank-order tournaments as optimum labor contracts." *Journal of Political Economy*, 89(5): 841–864.

Lizzeri, Alessandro, Margaret A Meyer, and Nicola Persico. 2005. "Interim Evaluations in Dynamic Tournaments: the Effects of Midterm Exams."

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.

Mihm, Jürgen, and Jochen Schlapp. 2019. "Sourcing innovation: On feedback in contests." *Management science*, 65(2): 559–576.

Moldovanu, Benny, and Aner Sela. 2001. "The optimal allocation of prizes in contests." *American Economic Review*, 91(3): 542–558.

Nalebuff, Barry J, and Joseph E Stiglitz. 1983. "Prizes and incentives: towards a general theory of compensation and competition." *The Bell Journal of Economics*, 21–43.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov. 2020. "Persuading the principal to wait." *Journal of Political Economy*, 128(7): 2542–2578.

Rayo, Luis, and Ilya Segal. 2010. "Optimal information disclosure." *Journal of Political Economy*, 118(5): 949–987.

Renault, Jérôme, Eilon Solan, and Nicolas Vieille. 2017. "Optimal dynamic information provision." *Games and Economic Behavior*, 104: 329–349.

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

Smolin, Alex. 2021. "Dynamic evaluation design." *American Economic Journal: Microeconomics*, 13(4): 300–331.

Taylor, Curtis R. 1995. "Digging for golden carrots: An analysis of research tournaments." *The American Economic Review*, 872–890.

Yildirim, Huseyin. 2005. "Contests with multiple rounds." *Games and Economic Behavior*, 51(1): 213–227.

A Omitted Proofs

A.1 Proofs of Proposition 1.

Omitted. See the proofs of Proposition 1 and 2 of Ely, Georgiadis and Rayo (2025).

A.2 Proof of Lemma 1.

Recall that

$$\underline{M} \sim \text{Binom}\big(N-1, F(\tau)\big), \qquad \overline{M} \mid \underline{M} = m \sim \text{Binom}\big(N-m-1, F(\mathcal{T}) - F(\tau)\big),$$

denotes the number of players other than i who succeed during the early and the late period, respectively. Define

$$G_\tau(\alpha) \;:=\; \mathbb{E}\Big[\tfrac{\alpha}{\alpha+\underline{M}+\alpha\overline{M}}\Big], \qquad H(\tau) \;:=\; \mathbb{E}\Big[\tfrac{1}{1+\underline{M}+\alpha(\tau)\overline{M}}\Big] - \Big[\tfrac{c}{\lambda(\mathcal{T})} + \tfrac{cF(\tau)}{f(\tau)}\Big].$$

Step 1: $\alpha(\tau)$ exists and is unique for every τ.

Fix τ and differentiate G_τ *inside* the expectation. For any realisation $(\underline{m}, \overline{m})$ write

$$g_{\underline{m},\overline{m}}(\alpha) := \frac{\alpha}{\alpha + \underline{m} + \alpha\overline{m}} = \frac{\alpha}{\underline{m} + \alpha(\overline{m}+1)}.$$

Differentiating yields

$$g'_{\underline{m},\overline{m}}(\alpha) = \frac{\underline{m}}{\big(\underline{m} + \alpha(\overline{m}+1)\big)^2} \geq 0,$$

with strict inequality whenever $\underline{m} > 0$. Because $\Pr(\underline{M} > 0) = 1 - (1 - F(\tau))^{n-1} > 0$ for any $\tau > 0$, we have $G'_\tau(\alpha) = \mathbb{E}\big[g'_{\underline{M},\overline{M}}(\alpha)\big] > 0$ for all $\alpha \in [0, 1)$, so $G_\tau(\alpha)$ is *strictly* increasing. Moreover $G_\tau(0) = 0$ and, since $1 + \underline{M} + \overline{M} \leq n$, $G_\tau(1) \geq 1/N$. Because

$c/\lambda(\mathcal{T}) \in (0, 1/N)$ by assumption, there is a *unique*

$$\alpha(\tau) \in (0,1) \;\text{ such that }\; G_\tau\big(\alpha(\tau)\big) = \frac{c}{\lambda(\mathcal{T})}.$$

Moreover, the mapping $\tau \mapsto \alpha(\tau)$ is continuous by the Implicit Function Theorem.

Step 2: $H(\tau)$ is strictly decreasing and changes sign.

First note that both $\underline{M}$ and $\overline{M}$ depend on τ only through $F(\tau)$; hence H is continuous. To see that it is monotone, increase τ by $d\tau > 0$.

- $F(\tau)$ rises, so $1 + \underline{M}$ in the denominator of H (computed at $\tau + d\tau$) stochastically increases, strictly lowering the expectation in H.
- The bracketed component $c/\lambda(\mathcal{T}) + cF(\tau)/f(\tau)$ *increases* in τ because F and F/f are weakly increasing by assumption (f and λ are decreasing).

Both effects lower H, so $H'(\tau) < 0$ wherever the derivative exists.

Finally, evaluate H at the endpoints:

$$H(0) = \mathbb{E}\Big[\tfrac{1}{1+\alpha(0)\overline{M}}\Big] - \frac{c}{\lambda(\mathcal{T})} > 0 \text{ and } H(\mathcal{T}) = \mathbb{E}\Big[\tfrac{1}{1+\underline{M}}\Big] - \Big[\frac{c}{\lambda(\mathcal{T})} + \frac{cF(\mathcal{T})}{f(\mathcal{T})}\Big] < 0.$$

The strict signs follow because in each case the second bracketed term exceeds the first. By continuity and strict monotonicity there is a unique $\tau(\mathcal{T}) \in (0, \mathcal{T})$ with $H\big(\tau(\mathcal{T})\big) = 0$.

Step 3: Concluding the proof.

Set $\alpha(\mathcal{T}) := \alpha\big(\tau(\mathcal{T})\big)$. By Step 1 the late-period reward equals $c/\lambda(\mathcal{T})$, and by $H(\tau(\mathcal{T})) = 0$ the early-period reward equals $c/\lambda(\mathcal{T}) + cF\big(\tau(\mathcal{T})\big)/f\big(\tau(\mathcal{T})\big)$, so the minimal schedule $\overline{\overline{w}}(\cdot; \tau(\mathcal{T}), \mathcal{T})$ is matched pointwise by the weighted-egalitarian reward $\overline{w}(\cdot; \alpha(\mathcal{T}), \tau(\mathcal{T}), \mathcal{T})$. Uniqueness of both components follows from the strict monotonicity of G_τ and H. Finally, it follows from the Implicit Function theorem that $\alpha(\mathcal{T})$ is *strictly decreasing* while $\tau(\mathcal{T})$ is *strictly increasing* in $\mathcal{T}$. $\square$

A.3 Proof of Lemma 2.

From Lemma 1, we know that for any $\mathcal{T}$, there exists a unique pair $\{\tau(\mathcal{T}), \alpha(\mathcal{T})\}$ such that $\overline{w}(\cdot;\alpha(\mathcal{T}),\tau(\mathcal{T}),\mathcal{T}) \equiv \overline{\overline{w}}(\cdot;\tau(\mathcal{T}),\mathcal{T})$, and moreover, $\alpha(\mathcal{T})$ is *strictly decreasing* and $\tau(\mathcal{T})$ is *strictly increasing* in $\mathcal{T}$.

Step 1: Define for every $\tau > 0$ the function

$$\Delta_\tau(t) \ := \ \frac{2F(t) - 1 - \Phi(t)\big[1 - F(t)\big]}{2 + \Phi(\tau)} + 1 - F(t) + \Big[1 - \big(1 - F(t)\big)^{N-1}\Big]\frac{f(t)}{2 + \Phi(\tau)}.$$

Since F is concave, Φ is weakly increasing, and $\Phi'(t)/\Phi(t) > \lambda(t))$ we obtain

$$\Delta'_\tau(t) = -\frac{2f(t) + \Phi'(t)\,[1 - F(t)] - f(t)\,\Phi(t)}{2 + \Phi(\tau)} - f(t) - (N-1)f(t)\big(1 - F(t)\big)^{N-2}\frac{f(t)}{2 + \Phi(\tau)} \ < \ 0,$$

so Δ_τ is *strictly decreasing*. Because $F(0) = 0$ and $F(T) \leq 1$, one has

$$\Delta_\tau(0) = \frac{1 - \Phi(0)}{2 + \Phi(\tau)} + 1 > 0,$$

whereas $\Delta_\tau(T) < 0$. Hence for each τ there is a unique $\mathcal{T}(\tau)$ satisfying $\Delta_\tau\big(\mathcal{T}(\tau)\big) = 0$, and $\Delta_\tau(t) > 0$ (< 0) for $t < \mathcal{T}(\tau)$ ($t > \mathcal{T}(\tau)$). Because $\Delta_\tau(t)$ is strictly decreasing *in t* while $\Phi(\tau)$ enters only multiplicatively, it follows from the Implicit Function theorem that $\mathcal{T}(\tau)$ is *strictly decreasing* and continuous in τ. Observe that (7) is precisely the condition $\Delta_\tau(t) = 0$ with $t = \mathcal{T}$.

Step 2: Define

$$\Psi(\tau) \ := \ \tau\big(\mathcal{T}(\tau)\big).$$

Because $\tau(\cdot)$ is strictly increasing while $\mathcal{T}(\cdot)$ is strictly decreasing, Ψ is strictly *decreasing* and continuous on $(0, T)$, with

$$\Psi(\tau) - \tau \begin{cases} > 0, & \text{as } \tau \downarrow 0, \\ < 0, & \text{as } \tau \uparrow T. \end{cases}$$

Hence Ψ has a unique fixed point $\tau^* \in (0, T)$. Set $\mathcal{T}^* := \mathcal{T}(\tau^*)$ and $\alpha^* := \alpha(\mathcal{T}^*)$.

By construction, $\mathcal{T}^*$ satisfies (7), and by Lemma 1, $\overline{w}(t; \alpha^*, \tau^*, \mathcal{T}^*) = \overline{\overline{w}}(t; \tau^*, \mathcal{T}^*)$. Uniqueness follows from the strict monotonicities. □

A.4 Proof of Lemma 3

We can rewrite the objective in (8) as

$$J[r] = C_0 + \int_{\tau^*}^{T} r(t)G(t)\,dt + B\big(z(r)\big),$$

where

$$C_0 := \int_0^{\tau^*} (1 - \eta^* c[2 + \Phi(t)])\,dt, \qquad G(t) := [1 - F(t)]\big[1 - \eta^* c(2 + \Phi(t))\big] + \eta^* c,$$
$$B(z) := \frac{\eta^*}{N}\Big(1 - \big[1 - z\big]^N\Big), \text{ and} \qquad z(r) := \int_0^T r(t) f(t) dt.$$

Notice that B is strictly increasing and strictly concave for all $z \in (0, 1)$.

Step 1: $G(t)$ is strictly decreasing on $[\tau^*, T]$.

Write $A(t) := 1 - \eta^* c\,[\,2 + \Phi(t)\,]$, so $G(t) = (1 - F(t))A(t) + \eta^* c$. Because Φ is increasing and $\Phi(\tau^*)$ solves $A(\tau^*) = 0$, we have $A(t) \leq 0$ for every $t \geq \tau^*$. Differentiating,

$$G'(t) = -f(t)A(t) - \eta^* c\,[\,1 - F(t)\,]\Phi'(t) = f(t)\big[\eta^* c\,[\,2 + \Phi(t)\,] - 1\big] - \eta^* c\,[\,1 - F(t)\,]\Phi'(t).$$

Use the assumption $\Phi'(t)/\Phi(t) \geq f(t)/(1-F(t))$ to obtain $(1-F)\Phi' \geq f\Phi$ and hence

$$G'(t) \;\leq\; f(t)\Big[\eta^* c\,[\,2+\Phi(t)\,] - 1 - \eta^* c\,\Phi(t)\Big] \;=\; f(t)\,[\,2\eta^* c - 1\,].$$

Because $A(\tau^*) = 0$ implies $\eta^* c = 1/(2+\Phi(\tau^*)) < \frac{1}{2}$, the coefficient $(2\eta^* c - 1)$ is strictly negative. Since $f(t) > 0$, we have $G'(t) < 0$, so G is strictly decreasing on $[\tau^*, T]$.

Step 2: Optimal $r(t)$ is bang-bang.

Let r be any feasible non-increasing function. Suppose that there exist $t_1 < t_2$ with $r(t_1) < 1$ and $r(t_2) > 0$. Move a small mass $\varepsilon > 0$ of r from t_2 to t_1, yielding $\tilde{r}$. Because G and f are decreasing,

$$\int_{\tau^*}^{T} \tilde{r}(t)G(t)\,dt - \int_{\tau^*}^{T} r(t)G(t)\,dt \;=\; \varepsilon[\,G(t_1) - G(t_2)\,] \;>\; 0,$$

and

$$z(\tilde{r}) - z(r) = \varepsilon[\,f(t_1) - f(t_2)\,] > 0.$$

Since $B'(\cdot) > 0$, the change also increases $B(z(r))$. Hence $J[\tilde{r}] > J[r]$. Iterating this "left shift" until no such pair (t_1, t_2) exists converts any optimal r into a *cut-off rule*

$$r_s(t) = \mathbf{1}\{t \leq s\} \quad \text{for some } s \in [\tau^*, T].$$

Step 3: Choosing the optimal cut-off.

Define

$$J(s) := C_0 + \int_{\tau^*}^{s} G(t)\,dt + B\big(F(s)\big), \qquad s \in [\tau^*, T].$$

Because B is strictly concave and G is continuous, J is continuously differentiable with

$$J'(s) = G(s) + \eta^*[\,1 - F(s)\,]^{N-1} f(s).$$

Using $G'(t) < 0$ and $f'(t) \leq 0$, one verifies that J' is strictly decreasing. Lemma 2 tells us that $J'(\mathcal{T}^*) = 0$; hence $\mathcal{T}^*$ is the unique maximizer of $J(s)$, and the corresponding policy r defined in the statement of the lemma is uniquely optimal. $\square$