Prizes and Feedback in Dynamic Contests*

Jeffrey C. Ely † George Georgiadis ‡ Luis Rayo §

September 9, 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, and firms often rely on contests when day-to-day effort is hard to monitor. Innovation contests, research competitions, and 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 rule by which the prize is allocated and the information, or feedback, that players receive while the contest unfolds. This paper studies how these instruments should be jointly designed.

We consider a continuous-time contest among N symmetric players competing for a fixed prize. At every date, each player privately chooses whether to exert costly effort toward a binary success. Conditional on effort, success arrives according to a Poisson process and can occur at most once for each player. This one-success structure is a simplification. Success may represent a verifiable technical qualification that renders an innovation useful to the designer and is readily contractible, whereas further effort adds value through subsequent improvements or innovations that are harder to contract upon.

The designer chooses a reward rule, which determines the expected probability that a player who succeeds at a given date receives the prize, and a feedback policy, which tells each player whether to continue working depending on whether and when she has succeeded. The designer values both total effort and the event that at least one contestant succeeds, subject to incentive compatibility and the fixed-prize budget constraint.

The methodological challenge is that rewards and feedback affect incentives dynamically. A player who is asked to work at date t may be tempted not only to stop permanently, but also to pause briefly and shift effort into the future. Such a pause changes the probability and timing of success and the likelihood that the player remains in states in which she is asked to work. We adapt the local-incentive approach of **?** to compute, for any fixed feedback policy, the least expensive expected rewards that deter such pauses. We then ask whether these rewards can be implemented by an actual contest satisfying the single-prize budget.

The optimal design is a weighted egalitarian contest. Up to a cutoff date τ, all players are asked to work continuously and receive no feedback about their own or others' successes. From τ until an exogenous deadline T, only unsuccessful players are asked to work, and success is disclosed immediately. Players who succeed before τ receive weight one, while those who succeed between τ and T receive a smaller weight; the prize is allocated among those who succeed in proportion to their weights.

The economics of this design is driven by the impact of keeping players uninformed combined with a "backward-compounding effect." The only way to induce effort beyond the moment of success is to keep players in the dark: once a player learns that she has succeeded, she has no reason to continue working. But silence is costly. An uninformed player may already have succeeded, so she must be compensated for this possibility with a higher reward. This cost is amplified by a backward-compounding effect. If the designer raises the reward at some later date, then rewards at all earlier dates must also rise to prevent players from pausing and instead working later on. Thus, the same amount of additional effort is more expensive to sustain later in the contest than earlier. As a result, silence is optimally concentrated at the start.

The same logic explains the unequal prize weights. The premium for early success (unit weight versus a smaller one) is exactly the compensation for working in the dark. Before τ, a successful player keeps exerting wasted effort without knowing it, so deterring pauses is expensive; after τ, a working player knows she has not yet succeeded and knows that success will release her, so a smaller reward suffices.

Although the formal implementation resembles a raffle, its practical interpretation is broader. What matters is that agents understand how different success dates translate into different expected chances of receiving the prize. In practice, the final allocation of a prize, promotion, or award may depend on additional considerations that agents do not control or observe in advance. The model's raffle can therefore be interpreted as a reduced form for such residual uncertainty: early successes are placed in a more favorable category, while the remaining selection depends on factors that agents view as beyond their control or effectively random. Section 5.1 applies the model to online data-science competitions.

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 **?** in a static environment and **?** in a dynamic one; see also **?** for a review. Subsequent contributions study environments that are close to ours in their focus on stochastic breakthroughs. For example, **?** and **?** analyze settings with discrete stochastic successes, and show that rewards optimally decrease over time. Our model shares this feature but, in addition to having multiple competing agents, our reward structure differs fundamentally as the principal chooses not only how to reward successes, but also what agents learn.

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 **?**, **?**, **?**, and **?**, who study how rank-based rewards and prize-sharing arrangements shape performance. More closely related are dynamic contest models for stochastic innovations, such as **?** and **?**, which consider winner-takes-all prizes under fixed feedback rules. Our setting sits between the two: as in **?** the designer values effort, but as in **?** the prize can be allocated only on the basis of a contractible output, which in our case is binary and arrives at most once per player. The reward instrument is therefore coarser than the objective it serves, and feedback becomes an additional lever for sustaining effort. **?** 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 **?** and **?** studies how a designer can influence actions by choosing what information agents receive. **?** and **?** 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 **?**, **?**, **?**, **?**, **?**, and **?**. 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, **?**, **?**, **?**, **?**, **?**, **?**, and **?** study how disclosure rules affect competitive effort. These papers typically restrict attention to relatively simple disclosure policies. More recently, **?** study optimal dynamic contests in continuous time and show that optimal feedback can take a more sophisticated form.

The two papers most closely related to ours are **?** and **?**. The present paper connects these two analyses by applying the feedback-design tools developed in **?** to a multi-agent contest with a fixed prize.

The comparison with **?** is clearest relative to a benchmark in which the prize budget is just sufficient to contract separately with every player under full disclosure. At this benchmark, each player works until succeeding or until the deadline and is informed immediately upon succeeding. **?** consider a budget below this level. Because prize resources are scarce, sustaining effort requires reassuring contestants that not too many rivals have already succeeded. We instead consider a budget above the benchmark. Here, the designer can already induce every unsuccessful player to work until the deadline without providing any information about her opponents. The additional prize capacity is used instead to finance an initial silent phase, during which players continue working even after their own success.[1] Thus, unlike the peer-feedback designs in **?**, our mechanism more closely resembles individual contracting: each player's feedback depends only on her own success history, while the common prize is pooled through a shared raffle at the end.

The connection with **?** concerns both the structure of the optimal policy and the method used to derive it. In a single-agent environment with flexible rewards and feedback, they characterize the minimal rewards that deter strategic pauses and show how a backward-compounding effect produces a silent phase followed by a phase of immediate disclosure. We use those tools to determine the feedback and expected rewards required for each contestant. The new challenge is implementation in a multiple-player setting: rewards cannot be chosen independently across players, but must be generated by a common, fixed prize. We therefore jointly determine the feedback cutoff, contest deadline, and relative rewards for early and late successes, and show that, despite the fixed budget, the resulting reward schedule can be implemented by a weighted egalitarian contest.

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—a binary choice—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 while exerting effort, success arrives with a constant hazard rate λ. Thus, if a player has spent cumulative effort e, the probability of success is $F(e) = 1 - e^{-\lambda e}$, with corresponding density $f(e)$. Throughout we assume that $\lambda > Nc$.[2]

[1] The silent phase shrinks as the prize budget converges to the benchmark level.

[2] **?** characterize optimal contests that maximize total effort for the simpler case in which $\lambda \leq Nc$. In that case, all players are kept fully informed about their own successes and the contest designer can extract all rents.

Effort is privately chosen and is neither observable nor contractible. The designer observes the vector of realized successes and their dates, and these outcomes are contractible. Contestants, however, learn about their own successes only through the feedback selected by the designer. The designer commits ex ante to a mechanism, $\mathcal{M}$, comprising a prize-allocation rule and a feedback policy.[3]

The prize-allocation rule determines who receives the prize as a function of the vector of success times. We encode this rule by a function $w(t)$, which gives the probability that a player who succeeds at time t wins the prize.[4] We focus on symmetric contests in which this reward function is the same for all players. Implicit in this formulation is that an unsuccessful player receives nothing. Since players are risk-neutral, this is without loss because rewarding unsuccessful players would weaken incentives, thus necessitating a larger total reward to induce the same effort.

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 under which each player is told whether to continue working or to stop. We restrict to policies where a recommendation to stop is permanent. We summarize the recommendation policy using two integrable continuation probabilities. The function $r(t)$ is the probability that a player who has not succeeded remains recommended to work through date t, while $q(t \mid s)$ is the probability that a player who succeeded at date $s \leq t$ remains recommended to work through date t. Because a recommendation to stop is permanent, both $r(t)$ and $q(t \mid s)$, which are survival probabilities, are necessarily non-increasing in t.[5]

[3]Commitment is essential. Without it, after observing a success the designer could conceal it and retain the prize rather than carry out the promised disclosure and allocation.

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

[5]The recommendations may in principle depend on other players' histories and on auxiliary randomizations. The functions $r(t)$ and $q(t \mid s)$ are simply the corresponding continuation probabilities averaged over these additional variables, conditional on the player's own success status and success date.

The probability that a player is working at date t is

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

Hence, 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 and $cp(s)$ is the expected effort cost at time s.

The designer chooses a symmetric mechanism $\mathcal{M}$ to maximize

$$N \int_0^T p(s)ds + \nu \times \mathcal{S}(\mathcal{M}), \tag{2}$$

where $\mathcal{S}(\mathcal{M})$ denotes the probability that at least one player succeeds under the mechanism, and $\nu \geq 0$ measures how much the designer values the success itself.

This maximization is subject to an incentive-compatibility constraint requiring that players find it optimal to follow all recommendations. It is also subject to the budget constraint

$$N \int_0^T r(t)w(t)f(t)dt \leq \mathcal{S}(\mathcal{M}), \tag{3}$$

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

In general, $\mathcal{S}(\mathcal{M})$ depends on the joint distribution of players' recommendation and success histories and therefore cannot be expressed solely in terms of the reduced-form marginal probabilities defined above. Of particular interest is the special case in which recommendations are independent across players (ruling out, for example, a contest where players are asked to stop once a critical mass of peers have succeeded). In this case, a given player succeeds with probability $\int_0^T r(t)f(t)\,dt$, and hence

$$\mathcal{S}(\mathcal{M}) = 1 - \left[1 - \int_0^T r(t)f(t)\,dt\right]^N,$$

where the term in square brackets is the probability that a given player does not succeed. This representation makes the designer's problem tractable and allows us to solve the corresponding maximization explicitly. Our proof strategy involves identifying the optimal mechanism under independent recommendations and showing that correlated recommendations cannot yield a higher payoff.

3 Local Incentives and Minimal Reward Schedule

Here we provide the pointwise-minimal expected reward needed to deter instantaneous effort pauses. We later show that this expected reward guarantees global incentive compatibility, and, crucially, that it can be implemented by an actual contest rule satisfying the prize budget.

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

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

which is the expected amount of *future* work that the policy asks from a player who succeeds at date t.

Recall that $p(t)$ is the ex-ante probability that the player is working at date t, as defined in (1). 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 needed if there were no future incentive effects. To see this, suppose date t were the last instant. At that date, a working agent bears an expected flow cost $c\,p(t)$ while receiving an expected flow reward $r(t)w(t)f(t)$. 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 **?**.

Proposition 1 (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*

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

Proof. Omitted. See the proofs of Propositions 1 and 2 of **?**. □

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*—this is weakly positive since $f'(s) \le 0$. 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.

Second,

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

is a *future-work 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 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 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.

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.

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.[6]

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

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

[6]In practice, the randomization could be generated by a privately observed component of the designer's preferences.

where $\underline{M}$ and $\overline{M}$ denote the number of peers who succeed during the early and the late period, respectively.[7]

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 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 it to (6), the minimal reward function is:

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

Of particular interest are contests that make use of the entire time budget (i.e., $\mathcal{T} = T$) as they allow effort to be maximal. The following lemma shows that there exists a unique weighted egalitarian contest with full duration such that the local incentive constraint binds at all times.

Lemma 1. *For sufficiently large T, there exists a unique pair $(\tau^*, \alpha^*) \in (0, T) \times (0, 1)$ such that*

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

To understand this result intuitively, recall that the lowest (locally) incentive com-

[7] We write these variables without an i subscript for simplicity. Since they affect the reward only through the expectations above and the contest is symmetric, the reward itself does not require such subscript.

patible reward has a baseline $c/\ OKlambda$ 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 because they face almost no early competitors ($\underline{M} \approx 0$). 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 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 a weighted egalitarian contest as optimal:

Theorem 1. *The "weighted egalitarian" contest with duration T and parameters τ^* and α^* is optimal.*

Intuitively, the use of the full horizon reflects the ample prize budget. The assumption $\lambda > Nc$ makes the prize large enough to contract separately with every player under full disclosure and induce each to work until success or T. The optimal mechanism therefore never asks an unsuccessful player to stop before T. Any prize capacity beyond what is needed for this benchmark is instead used to induce additional effort from successful players by keeping them uninformed during part of the contest.

Because the designer's objective is linear in the effort path, feedback is pushed to the corners of either full silence (which maximizes expected effort) or full disclosure (which minimizes it). Moreover, early effort is less expensive than late effort: raising effort at a late date t requires the designer not only to grant rents at t, but also to raise rewards at all earlier dates to prevent players from pausing and shifting effort forward. It follows that any partial disclosure before τ is inefficient, because it allows some early winners to stop when effort is cheapest, while any opacity after τ is inefficient, because it keeps winners working when additional effort is most expensive. The weighted-egalitarian design implements this bang–bang path with the smallest possible incentive-compatible rewards while exhausting the prize budget, thus maximally translating prize capacity into effort.

The use of the full horizon explains why correlation across players' recommendations cannot improve on the solution. No mechanism can make success more likely than

$$1 - \left[1 - F(T)\right]^N,$$

the probability obtained when every unsuccessful player works through T. Our independent policy attains this upper bound and, because the raffle awards the entire prize whenever at least one player succeeds, also attains the largest possible right-hand side of the budget constraint. Thus, a correlated policy cannot generate a higher probability of success or a larger expected reward budget. Nor can it raise expected effort: our policy already induces all the cheaper effort before success and spends the remaining budget on post-success effort at the earliest, least expensive dates.

Finally, the constant-hazard assumption is not essential. In the Appendix, we prove Theorem 1 under more general success distributions $F(e)$ with a weakly decreasing hazard rate (capturing a "low-hanging-fruit-first" technology) subject to some additional conditions—including that the hazard rate does not decline too quickly (see Assumption 1). This ensures that every player who remains unsuccessful continues working through the full horizon.

5.1 Application: Online Data-Science Competitions

Online data-science platforms host competitions sponsored by firms, government agencies, and research organizations for algorithms and related technical solutions, including demand forecasting, fraud detection, and medical-risk prediction. Sponsors release labeled training data and test data whose outcomes are withheld. Teams repeatedly submit predictions and use scores generated by a preannounced metric to refine their data processing, features, models, and tuning parameters.

A useful abstraction is that submissions fall into two broad groups: those with little realistic chance of winning and those at the competitive frontier. Several teams may reach this frontier, with measured performance so close that small differences in final rankings may partly reflect sampling variation rather than economically meaningful differences in quality. We therefore interpret success in our model as first developing an algorithm that places a team in serious contention. Successful teams are effectively tied, with residual uncertainty determining which receives the prize.

Unlike in our model, platforms typically provide ongoing scores through a public provisional leaderboard while reserving a private evaluation for the final ranking.[8] In practice, a team's own scores may help it assess whether its approach is promising and direct further effort, while other teams' scores may reveal what is attainable. Neither channel is present in our model as the success technology is already common knowledge.

The model nonetheless highlights an information-design margin that organizers may wish to consider, provided their prize budget is plentiful: an initial silent phase. By delaying feedback, an organizer can keep teams working beyond the moment at which they become serious contenders. This silence is most valuable early in the contest because inducing effort at a later date requires not only making success at that date sufficiently rewarding, but also raising the rewards attached to all earlier success dates. This

[8] Kaggle's competitions feature a public leaderboard visible during the contest and a private leaderboard used for final ranking. Related public/private leaderboard designs are used by Zindi, DrivenData, and AIcrowd.

"backward-compounding" logic makes late effort disproportionately expensive to sustain.

6 Conclusion

This paper jointly studies prize allocation and feedback in dynamic contests with unobservable effort and a fixed, indivisible prize. The optimal design is a *weighted egalitarian* contest: early successes receive larger prize weights than late successes, while feedback is withheld initially and then provided immediately upon success. Contestants therefore work throughout the early phase—even after succeeding—and, during the late phase, only until success or the deadline. Together, these instruments make the most of a fixed prize by concentrating effort beyond success at the early dates when it is cheapest to induce.

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.

Daley, Brendan, and Ruoyu Wang. 2018. "When to release feedback in a dynamic tournament." *Decision Analysis*, 15(1): 11–26.

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. 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

For added generality, we show that the weighted egalitarian contest characterized in Theorem 1 is optimal under distributions F that satisfy the following assumption (which is met when F has a constant hazard rate).

Assumption 1. *The distribution F satisfies the following properties:*

i. The hazard rate $\lambda(t) := f(t)/[1-F(t)]$ is weakly decreasing and satisfies $\lambda(T) > Nc$ as well as $-\lambda'(t)/\lambda(t) \leq f(t)/F(t)$ for all $t \leq T$.

ii. The function $\Phi(t) := -F(t)f'(t)/f(t)^2$ is weakly increasing for all $t \leq T$.

A.1 Proof of Lemma 1.

Let $\overline{\mathcal{S}} := 1 - \left[1 - F(T)\right]^N$ denote the probability that at least one player succeeds when every unsuccessful player works through T. Under Assumption 1, applying Proposition 1 to the recommendation policy (6) with $\mathcal{T} = T$ shows that the minimal reward schedule is $w^{min}(t;\tau,T) = c/\lambda(T) + cF(\tau)/f(\tau)$ for $t < \tau$ and $c/\lambda(T)$ for $\tau \leq t \leq T$; that is, the constant λ of Section 4 is replaced by the hazard rate at the horizon, $\lambda(T)$.

We first determine the cutoff τ^*. If the weighted egalitarian contest implements the minimal reward schedule, its expected payment to each player must equal

$$\int_0^T w^{min}(t;\tau,T)f(t)\,dt = \int_0^\tau \left[\frac{c}{\lambda(T)} + \frac{cF(\tau)}{f(\tau)}\right] f(t)\,dt + \int_\tau^T \frac{c}{\lambda(T)} f(t)\,dt = \frac{cF(T)}{\lambda(T)} + \frac{cF(\tau)^2}{f(\tau)}.$$

Because the raffle awards the entire prize whenever at least one player succeeds, symmetry requires this payment to equal $\overline{\mathcal{S}}/N$. Thus, τ^* must satisfy

$$\frac{cF(T)}{\lambda(T)} + \frac{cF(\tau^*)^2}{f(\tau^*)} = \frac{\overline{\mathcal{S}}}{N}. \tag{7}$$

The left-hand side is continuous and strictly increasing in τ, since F is increasing and f is decreasing (which follows from the hazard rate being non-increasing). At $\tau = 0$ it equals $cF(T)/\lambda(T)$, which is strictly smaller than $\overline{\mathcal{S}}/N$ because $c/\lambda(T) < 1/N < \overline{\mathcal{S}}/[NF(T)]$. At $\tau = T$ it equals $cF(T)/f(T)$, using $1/\lambda(T) = [1-F(T)]/f(T)$, which exceeds $\overline{\mathcal{S}}/N$ for sufficiently large T: $f(T) = \lambda(T)[1-F(T)]$ vanishes as $F(T) \to 1$, while $\lambda(T) \leq \lambda(0)$. Therefore, for sufficiently large T, there exists a unique $\tau^* \in (0, T)$ satisfying (7).

We next determine α^*. Fix $\tau = \tau^*$ and let $\underline{M}$ and $\overline{M}$ denote the numbers of peers who succeed during the early and late periods, so that $(\underline{M}, \overline{M}, N-1-\underline{M}-\overline{M})$ is multinomial with parameters $N-1$ and $\big(F(\tau^*), F(T)-F(\tau^*), 1-F(T)\big)$. The expected reward from a late success is $G_{\tau^*}(\alpha) := \mathbb{E}\big[\alpha/(\alpha+\underline{M}+\alpha\overline{M})\big]$. Its integrand is increasing in α, strictly so when $\underline{M} > 0$, an event of positive probability; hence G_{τ^*} is continuous and strictly increasing on $(0, 1]$. At $\alpha = 1$ all successful players receive equal weight, so $G_{\tau^*}(1) = \overline{\mathcal{S}}/[NF(T)] > 1/N > c/\lambda(T)$. At the other endpoint a late success wins only if no peer succeeded early, and it is straightforward to verify that

$$G_{\tau^*}(0+) := \lim_{\alpha \downarrow 0} G_{\tau^*}(\alpha) = \mathbb{E}\left[\frac{\mathbf{1}\{\underline{M}=0\}}{1+\overline{M}}\right] = \frac{[1-F(\tau^*)]^N - [1-F(T)]^N}{N\,[F(T)-F(\tau^*)]}.$$

We claim that $G_{\tau^*}(0+) < c/\lambda(T)$. Because λ is decreasing, $f(\tau^*) = \lambda(\tau^*)[1 - F(\tau^*)] \geq \lambda(T)[1-F(\tau^*)]$, so (7) implies $\lambda(T)\,\overline{\mathcal{S}}/(Nc) \leq F(T) + F(\tau^*)^2/[1-F(\tau^*)]$. It therefore suffices to show that $N\,G_{\tau^*}(0+)\,\big\{F(T) + F(\tau^*)^2/[1-F(\tau^*)]\big\} < \overline{\mathcal{S}}$. Substituting the expression for $G_{\tau^*}(0+)$ and rearranging, straightforward algebra shows that this inequality is equivalent to

$$\sum_{j=1}^{N} [1-F(\tau^*)]^j \;>\; \sum_{j=0}^{N-1} [1-F(\tau^*)]^j\,[1-F(T)]^{N-1-j}.$$

For $N = 2$ this reads $[1-F(\tau^*)]^2 > 1-F(T)$, which holds for sufficiently large T: by (7), $F(\tau^*)^2/f(\tau^*) \leq 1/(Nc)$, so τ^* is bounded uniformly in T while $1-F(T) \to 0$. For $N \geq 3$ no restriction on T is needed. Indeed, the right-hand side is increasing in $1-F(T) \leq 1-F(\tau^*)$ and is therefore at most $N[1-F(\tau^*)]^{N-1}$, whereas the left-hand side equals

$[1-F(\tau^*)]\sum_{j=0}^{N-1}[1-F(\tau^*)]^j$, and $\sum_{j=0}^{N-1} x^j > Nx^{(N-1)/2} \geq Nx^{N-2}$ for $x \in (0,1)$ and $N \geq 3$ by the arithmetic–geometric mean inequality. This proves the claim. It follows that there is a unique $\alpha^* \in (0,1)$ with $G_{\tau^*}(\alpha^*) = c/\lambda(T)$: the expected reward from a late success equals the minimal late reward.

It remains to verify the reward from an early success. Since the raffle awards the full prize whenever at least one player succeeds, each player's expected payout is $\overline{S}/N$, i.e.,

$$F(\tau^*)\,\mathbb{E}\Big[\frac{1}{1+\underline{M}+\alpha^*\overline{M}}\Big] + \big[F(T)-F(\tau^*)\big]\,G_{\tau^*}(\alpha^*) = \frac{\overline{S}}{N}.$$

Substituting $G_{\tau^*}(\alpha^*) = c/\lambda(T)$ and (7), and dividing by $F(\tau^*) > 0$, gives $\mathbb{E}\big[1/(1+\underline{M}+\alpha^*\overline{M})\big] = c/\lambda(T) + cF(\tau^*)/f(\tau^*)$, the minimal early reward. Hence $w^{weg}(t;\alpha^*,\tau^*,T) \equiv w^{min}(t;\tau^*,T)$. Finally, any implementing pair must satisfy (7), which pins down τ^*, and then $G_{\tau^*}(\alpha) = c/\lambda(T)$, which pins down α^*; the implementing contest is therefore unique. □

A.2 Proof of Theorem 1

Again write $\overline{S} := 1-[1-F(T)]^N$. It is also convenient to represent the success technology by thresholds: player i succeeds when her cumulative effort reaches θ_i, where $\theta_1,\ldots,\theta_N \overset{\text{i.i.d}}{\sim} F$ and independent of any randomization by the designer.

The proof proceeds as follows. Step 1 shows that under *any* mechanism, including one with correlated recommendations, the reduced-form probabilities $p(\cdot)$ and $r(\cdot)$ satisfy a budget constraint whose right-hand side is the constant $\overline{S}/N$; this yields a relaxed program whose value bounds the designer's payoff under every mechanism. Step 2 solves the relaxed program by a Lagrangian argument. Step 3 shows that its solution is implemented by the weighted egalitarian contest with parameters (α^*,τ^*,T), which attains the bound and is therefore optimal. Step 4 verifies global incentive compatibility.

Step 1: By its construction in the proof of Lemma 1, τ^* satisfies the budget identity (7).

Consider an arbitrary symmetric mechanism $\mathcal{M}$, whose recommendations may depend on the entire vector of success histories and on arbitrary common or private randomization. For a given player, collect all such variables other than her own success history in X, and let $r(t)$, $q(t \mid s)$ and $w(t)$ be her continuation probabilities and expected prize share averaged over X, as in Section 2. With these averages, (1) continues to give her probability $p(t)$ of working at t: the first term is the probability that she is unsuccessful and still working, and $r(s)f(s)ds$ is the probability that she succeeds in $[s, s + ds]$. Because a recommendation to stop is permanent, p and r are non-increasing, and $r(t)[1 - F(t)] \leq p(t) \leq 1$.

We first note that, for any recommendation policy $\big(q(\cdot \mid \cdot), r(\cdot)\big)$, Proposition 1, the definition of Φ, integration by parts and Fubini's theorem yield

$$\int_0^T r(t)\, w^{\min}(t)\, f(t)\, dt \;=\; c\int_0^T \Big([2 + \Phi(t)]\, p(t) - r(t)\Big)dt. \tag{8}$$

We next argue that Proposition 1 applies to $\mathcal{M}$ even though its recommendations may be correlated across players. Recommendations are work/stop instructions and stopping is permanent, so at each date the only history at which a player is told to work is "not yet stopped." Her payoff from any effort strategy therefore depends on $\mathcal{M}$ only through the joint law of her success date, her stopping date and her prize share, and through how this law responds to her effort. Before she succeeds, her stopping date is a function of X, which is independent of her effort, so for any effort path with cumulative effort e_t at date t she is unsuccessful and working at t with probability $r(t)[1 - F(e_t)]$. Once she succeeds, at date s say, the designer's information about her consists of s alone, so her continuation probabilities and expected prize share are $q(\cdot \mid s)$ and $w(s)$ whatever effort path produced s. Hence her gain from pausing at any date coincides with her gain under the recommendation policy (q, r) with reward w, and Proposition 1, applied to (q, r), implies $r(t)w(t) \geq r(t)w^{\min}(t)$ for all t. Together with (8), incentive compatibility of $\mathcal{M}$

therefore requires

$$\int_0^T r(t)\, w(t)\, f(t)\, dt \;\geq\; c \int_0^T \Big([2 + \Phi(t)]\, p(t) - r(t) \Big) dt. \tag{9}$$

Finally, an obedient player never accumulates more than T units of effort, so she succeeds only if $\theta_i \leq T$. Since the θ_i are independent,

$$\mathcal{S}(\mathcal{M}) \;\leq\; 1 - \big[1 - F(T)\big]^N = \overline{\mathcal{S}}, \tag{10}$$

with equality whenever every unsuccessful player works through T. Combining (9), the budget constraint (3) and (10), the reduced form of every incentive compatible mechanism that satisfies the budget constraint obeys

$$c \int_0^T \Big([2 + \Phi(t)]\, p(t) - r(t) \Big) dt \;\leq\; \frac{\overline{\mathcal{S}}}{N}, \tag{BC}$$

and, by (2) and (10), the designer's payoff under such a mechanism is at most $N \int_0^T p(t)\, dt + \nu\, \overline{\mathcal{S}}$. Notice that both (BC) and this payoff bound involve the mechanism only through $p(\cdot)$ and $r(\cdot)$. Consider therefore the program

$$\begin{aligned} \max_{0 \leq r(\cdot) \leq 1,\, p(\cdot)\ \text{decreasing}} \quad & \int_0^T p(t)\, dt + \frac{\nu}{N}\, \overline{\mathcal{S}} \\ \text{s.t.} \quad & (BC) \ \text{ and } \ 0 \leq r(t)[1 - F(t)] \leq p(t) \leq 1. \end{aligned} \tag{P}$$

The reduced form (p, r) of every feasible mechanism is feasible in (P), and the designer's payoff under that mechanism is at most N times the objective of (P) evaluated at (p, r). Hence N times the value of (P) is an upper bound on the designer's payoff under every mechanism, and it remains to show that the weighted egalitarian contest attains it.

Step 2: Dualizing (BC), define for $\eta \geq 0$

$$L(\eta) := \max_{0 \leq r \leq 1,\, p \text{ decr.}} \int_0^T p(t)\Big(1 - \eta c\,[2 + \Phi(t)]\Big)\,dt + \eta c \int_0^T r(t)\,dt + \frac{\eta + \nu}{N}\,\overline{\mathcal{S}},$$

where the maximization is over $0 \leq r(t)[1 - F(t)] \leq p(t) \leq 1$ and $0 \leq r(t) \leq 1$. For any (p, r) feasible in (P), the objective of (P) is at most the expression maximized in $L(\eta)$, because the two differ by η times the slack in (BC), which is non-negative. Hence the value of (P) is at most $L(\eta)$ for every $\eta \geq 0$.

Define $\eta^* := 1/(c\,[2 + \Phi(\tau^*)])$ and consider $L(\eta^*)$. We maximize in two steps. First, for fixed r, the coefficient on $p(t)$ equals $[\Phi(\tau^*) - \Phi(t)]/[2 + \Phi(\tau^*)]$, which is non-negative for $t \leq \tau^*$ and non-positive thereafter because Φ is increasing. Hence it is optimal to set $p(t) = 1$ for $t \leq \tau^*$ and $p(t) = r(t)[1 - F(t)]$ otherwise; this p is decreasing whenever r is, so the monotonicity constraint on p does not bind. Substituting, the objective becomes

$$\int_0^{\tau^*} \Big(1 - \eta^* c\,\big[2 + \Phi(t) - r(t)\big]\Big)\,dt + \int_{\tau^*}^T r(t)\,G(t)\,dt + \frac{\eta^* + \nu}{N}\,\overline{\mathcal{S}}, \tag{11}$$

where $G(t) := [1 - F(t)]\Big(1 - \eta^* c\,[2 + \Phi(t)]\Big) + \eta^* c$. The coefficient on $r(t)$ is $\eta^* c > 0$ for $t \leq \tau^*$ and $G(t)$ for $t > \tau^*$. The following lemma shows that the latter is non-negative, so that (11) is maximized by $r(\cdot) \equiv 1$, which is decreasing.

Lemma 2. *Under Assumption 1, $G(t) \geq 0$ for all $t \in [\tau^*, T]$.*

Proof. Since $\eta^* c = 1/[2 + \Phi(\tau^*)]$ and $\Phi(\tau^*) \geq 0$,

$$[2 + \Phi(\tau^*)]\,G(t) = [1 - F(t)]\big[\Phi(\tau^*) - \Phi(t)\big] + 1 \;\geq\; 1 - [1 - F(t)]\,\Phi(t).$$

Writing $f = \lambda\,[1 - F]$ and differentiating gives $f' = [1 - F]\,[\lambda' - \lambda^2]$, hence $[1 - F]\,\Phi = -F[1 - F]f'/f^2 = F\,[1 - \lambda'/\lambda^2]$. Thus $[1 - F]\Phi \leq 1$ is equivalent to $-F\lambda'/\lambda^2 \leq 1 - F$, i.e. to $-\lambda'/\lambda \leq f/F$, which is Assumption 1(i). $\square$

Intuitively, $G(t)$ is the value, at the shadow price η^*, of keeping an unsuccessful player at work at a date $t > \tau^*$: she supplies effort $1 - F(t)$ at incentive cost $c[2 + \Phi(t)][1 - F(t)]$, while her continued eligibility relaxes the incentive constraints at all earlier dates by c (the future-work credit of Proposition 1). Assumption 1(i) ensures that the credit covers the cost net of effort, so that no unsuccessful player is asked to stop before T.

It follows that $L(\eta^*)$ is attained at

$$p^*(t) = \begin{cases} 1 & \text{if } t \leq \tau^* \\ 1 - F(t) & \text{if } \tau^* < t \leq T \end{cases} \qquad \text{and} \qquad r^*(t) = 1 \ \text{ for all } t \leq T,$$

which is the effort path induced by the recommendation policy (6) with $\mathcal{T} = T$.

Step 3: By (8) applied to (6) with $\mathcal{T} = T$, and by (7),

$$c \int_0^T \Big([2 + \Phi(t)]\, p^*(t) - r^*(t) \Big) dt = \int_0^T w^{\min}(t; \tau^*, T)\, f(t)\, dt = \frac{\overline{\mathcal{S}}}{N},$$

so (p^*, r^*) satisfies (BC) with equality and is feasible in (P). The term multiplying η^* in $L(\eta^*)$ therefore vanishes at (p^*, r^*), and $L(\eta^*)$ equals the objective of (P) at (p^*, r^*). Since the value of (P) is at least this objective and at most $L(\eta^*)$ by Step 2, (p^*, r^*) solves (P).

Now consider the weighted egalitarian contest with parameters (α^*, τ^*, T). It induces (p^*, r^*); it is incentive compatible, by Lemma 1 for instantaneous pauses and by Step 4 for all other deviations; and it satisfies (3), indeed with equality, since it awards the entire prize whenever at least one player succeeds, which by (10) occurs with probability $\overline{\mathcal{S}}$. Its payoff is therefore $N \int_0^T p^*(t)\, dt + \nu\, \overline{\mathcal{S}}$, which is N times the value of (P), the upper bound established in Step 2 for every mechanism. Hence the contest is optimal, whether or not recommendations are allowed to be correlated across players.

Step 4: We conclude the proof by ruling out global deviations. As in **?**, 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. $\square$