DISCUSSION PAPER SERIES

DP16308

Optimal Feedback in Contests

Jeffrey Ely, George Georgiadis, Sina Khorasani and Luis Rayo

ORGANIZATIONAL ECONOMICS

CEPR

ISSN 0265-8003

Optimal Feedback in Contests

Jeffrey Ely, George Georgiadis, Sina Khorasani and Luis Rayo

Discussion Paper DP16308
Published 29 June 2021
Submitted 25 June 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Optimal Feedback in Contests

Abstract

We derive an optimal dynamic contest for environments where the principal monitors effort through a coarse, binary performance measure and chooses prize-allocation and termination rules together with a real-time feedback policy. The optimal contest takes a stark cyclical form: contestants are kept fully apprised of their own successes, and at the end of each fixed-length cycle, if at least one agent has succeeded, the contest ends and the prize is shared equally among all successful agents regardless of when they succeeded; otherwise, the designer informs all contestants that nobody has yet succeeded and the contest resets.

JEL Classification: N/A

Keywords: N/A

Jeffrey Ely - jeff@jeffely.com
Northwestern University

George Georgiadis - g-georgiadis@kellogg.northwestern.edu
Northwestern University

Sina Khorasani - skhorasani@UCSD.EDU
University of California San Diego

Luis Rayo - luis.rayo@kellogg.northwestern.edu
Northwestern University and CEPR

Optimal Feedback in Contests*

Jeffrey Ely, George Georgiadis, Sina Khorasani, and Luis Rayo†

June 25, 2021

Abstract

We derive an optimal dynamic contest for environments where the principal monitors effort through a coarse, binary performance measure and chooses prize-allocation and termination rules together with a real-time feedback policy. The optimal contest takes a stark cyclical form: contestants are kept fully apprised of their own successes, and at the end of each fixed-length cycle, if at least one agent has succeeded, the contest ends and the prize is shared equally among all successful agents regardless of when they succeeded; otherwise, the designer informs all contestants that nobody has yet succeeded and the contest resets.

1 Introduction

Contests—situations where multiple agents compete for a prize—are a common way of organizing economic activity: innovation races, promotions and other labor-market tournaments, all-pay auctions, athletic events, and legal battles all have this property in common. Ever since the seminal work of Lazear and Rosen (1981), Green and Stokey (1983), and Nalebuff and Stiglitz (1983), researchers in economics, marketing and operations management have sought to understand how to best allocate a prize among participants, and more recently, starting with the work of Yildirim (2005),

*We are grateful to Daniel Barron, Piotr Dworczak, Marina Halac, Navin Kartik, Igor Letina, Michael Powell, Agustiń Rayo, and participants at various seminars and conferences for helpful comments. Author Jeff Ely acknowledges the support of NSF grant SES-1851883.

†J. Ely: Northwestern University, jeff@jeffely.com; G. Georgiadis: Northwestern University, g-georgiadis@kellogg.northwestern.edu; S. Khorasani: University of California–San Diego, skhorasani@ucsd.edu; L. Rayo: Northwestern University, luis.rayo@kellogg.northwestern.edu.

Ederer (2010), and Halac, Kartik and Liu (2017), how best to disclose real-time information regarding the contestants' progress. This literature has offered key general insights together with design ideas for specific environments of interest.

Here we are interested in scenarios where a designer (principal) has a fixed prize to motivate a group of agents and has an informational advantage in terms of how well the agents are doing mid-contest. Our goal is to find an optimal contest inclusive of when it ends, how the prize is allocated, and what information the designer shares with the agents throughout the event.

In our model, agents choose at each instant whether or not to exert costly effort. The principal monitors them only imperfectly through a binary measure of "success" that arrives stochastically for each agent at a time-invariant rate provided the agent is working. Agents do not observe the successes of their rivals and may or may not observe their own successes, or do so probabilistically. The principal seeks to maximize total effort, or equivalently as it turns out, the total number of successes. To this end, the principal designs prize-allocation and termination rules together with a real-time feedback policy contingent on past successes and past messages.[1]

This framework can be applied, for instance, to a professional partnership seeking to promote one of their associates to partner. Here a "success" represents an associate exceeding an exogenous threshold for promotion, and an associate's effort is arguably valuable even after they, or any of their peers, have cleared the bar for promotion. One may also consider innovation tournaments such as the Netflix Prize, where success corresponds to a pre-specified objective and the more appropriate goal for the designer is maximizing the total number of successes, rather than total effort.[2]

The key challenge when searching for the optimal contest is the vast range of potential designs from which to choose. We attack this problem by first providing a sufficient condition for a contest to be optimal—namely, that it maximizes the probability that the prize is awarded while giving zero rents to the contestants—and then displaying a contest that meets these demanding criteria. That contest, which we term *cyclical-egalitarian*, has the following features regardless of the degree to

[1] Unless the available prize is large relative to the number of agents, a contest strictly outperforms individual contracting because it pools the agents' incentive constraints: by forcing agents to share the prize in expectation with their rivals, no prize money is wasted. This pooling, as we shall see, is aided by the principal's strategic sharing of information mid-contest.

[2] The principal would care about more than a single success if each additional success generated new, distinct technology. For instance, to save on engineering effort, Netflix did not implement the prize-winning algorithm and instead opted for a combination of two lesser-performing ones. [*Link.*]

which agents directly observe their own successes:

- **An "egalitarian" prize structure** that splits the prize evenly—or, equivalently, allocates it with equal probability—among all agents who have succeeded by the time the contest ends, regardless of when they happened to succeed.

- **A "cyclical" termination rule** whereby the principal sets a provisional deadline, T^*, such that if one or more agents have succeeded by then, the contest ends; otherwise, the contest resets and the deadline is extended until $2T^*$, and so forth.

- **A feedback policy** that immediately informs agents of their own successes, but provides them only periodic feedback (at the end of each cycle) regarding the success of their rivals.

Keeping agents fully apprised of their own successes is needed for the principal to pocket all rents, as otherwise agents would be able to extract rents by strategically timing their effort.[3] The egalitarian prize smooths the agents' incentives by ensuring that their marginal return to effort remains constant over time, which is needed as well to keep them indifferent between working and not. The periodic resetting of the deadline, together with the revelation that no one has yet succeeded, replenishes the agents' incentives after they have become sufficiently worried that their peers have succeeded and diluted the prize: the length of each cycle, T^*, maximally stretches out the period over which agents are willing to work without knowing whether their peers are already ahead.[4]

The cyclical-egalitarian contest strictly outperforms the popular winner-takes-all design for a wide range of parameter values. This is because it has an additional degree of freedom relative to that contest: the designer gets to control the cycle length and, with it, the level of rents that the contestants take home. Winner-takes-all is in fact a special case of our contest with a cycle length equal to zero, and is artificially constrained by that feature.

We also consider some extensions that suggest a degree of robustness to our recommended contest design. First, if players discount time, a version of our contest with

[3]While this policy leads each agent to halt effort as soon as they succeed, it is optimal even when the principal values effort beyond that success.

[4]Because each agent exerts effort until they succeed, this contest maximizes both total expected effort and the expected number of successes.

shorter cycle length is approximately optimal when the number of agents is large, the marginal cost of effort is large, or the discount rate is small; examples suggest that this contest also performs well away from those limits. Second, if the principal lacks the credibility to provide truthful feedback, an optimal contest features no feedback and a fixed terminal date, at which time the prize is shared equally among all successful agents. Finally, when the agents' hazard rate of success is increasing in past efforts, a cyclical contest with a similar structure as before proves optimal, with the novelty that the cycles end stochastically ahead of the deadline, as this differentially rewards agents who succeed early on, before their hazard rate has grown.

Due to its simplicity, this cyclical-egalitarian contest should not be difficult to implement in practice, as the only parameter that the designer needs to calibrate is the cycle length T^*, with a longer deadline possible the lower the cost of effort, the larger the probability of success, or the greater the prize. While there is limited empirical evidence so far on the effectiveness of different contest designs, the field experiment conducted by Lim, Ahearne and Ham (2009) lends support to the effectiveness of an egalitarian prize, and the findings of Fershtman and Gneezy (2011) and Gross (2017) suggest that real-time feedback, with a flavor similar to that suggested by our model, can be effective at encouraging effort.[5]

Related Literature. Early work by Lazear and Rosen (1981), Green and Stokey (1983), and Nalebuff and Stiglitz (1983) provides conditions under which it is optimal to condition each agent's pay on the ordinal rank of their output, as opposed to its absolute value. Moldovanu and Sela (2001) show that, given a fixed prize, it is optimal to award it entirely to the best performer when the agents' cost functions are weakly concave; otherwise, some prize-sharing may be optimal. Extensions to stochastic output, arbitrary risk-preferences and heterogeneous agents are considered by Drugov and Ryvkin (2019, 2020) and Olszewski and Siegel (2020), among others.[6]

Fang, Noe and Strack (2018) find that aggregate effort in all-pay contests decreases in their competitiveness, as measured by the dispersion of prizes, contest crowding, and the number of contestants. Letina, Liu and Netzer (2020) consider a generalized version of that framework. They find that for n contestants, a nested Tullock contest featuring $n - 1$ equal prizes is optimal. While our work differs in that these papers

[5]See also Lemus and Marshall (2021) for a comparison of feedback policies in prediction contests.

[6]Siegel (2009, 2010) and Olszewski and Siegel (2016) provide a comprehensive equilibrium analysis of general all-pay contests with heterogeneous players.

study static environments with no scope for feedback design, the idea that "turning down the heat" motivates more effort echoes the optimality of an egalitarian contest in our dynamic framework.

Taylor (1995) considers a dynamic contest where players invest in an innovation of stochastic quality. In the optimal contest, players invest in a given period as long as their highest-quality innovation to date is below a threshold. Benkert and Letina (2020) extend this framework by incorporating interim transfers and an endogenous termination date. The optimal contest ends as soon as the highest-quality innovation exceeds a threshold, and agents invest until the end of the contest. In these models, the entire prize is allocated to the agent with the highest-quality innovation, and the principal extracts rents by charging agents an entry fee. In our setting, in contrast, the prize allocation rule is a choice variable, and because agents are cash-constrained, the principal requires the cyclical-egalitarian design to extract rents.

Lizzeri, Meyer and Persico (2005) and Yildirim (2005) are among the first to study endogenous feedback in contests using a two-period, two-agent framework. For this setting, Aoyagi (2010), Ederer (2010), and Goltsman and Mukherjee (2011) characterize conditions under which a principal benefits from (publicly) revealing the outcome of the contestants' first-period efforts. Mihm and Schlapp (2019) extend this framework by considering private feedback and by allowing agents to voluntarily disclose their own progress. Khorasani (2020) considers two-stage winner-takes-all contests, and shows that the optimal design features an initial period with no disclosure and a gradually increasing prize, followed by a period of probabilistic disclosure to the laggard about the intermediate progress of the leader.

Our paper also relates to a growing literature on contests involving experimentation, where the feasibility of success is initially unknown. Halac, Kartik and Liu (2017) consider an experimentation framework such as the one in Bonatti and Horner (2011), but with a designer who chooses a prize-sharing scheme and a feedback policy to maximize the probability of a success. Within the class of rank-monotonic prize schemes and deterministic and symmetric disclosure policies, the optimal contest provides no interim feedback and ends as soon as a critical number of agents have succeeded, each winning the prize with equal probability. What distinguishes our setting from this work is that the principal maximizes total effort (or total number of successes), successes arrive with a constant or increasing hazard rate, and the principal is fully unconstrained in her choice of contest. In Bimpikis, Ehsani and Mostagir

(2019), an agent must succeed twice to win, with the feasibility of the first success unknown. The authors obtain conditions under which a contest comprising a "silent period" followed by a period where successes are immediately disclosed dominates all contests with a constant, probabilistic disclosure (including those with full disclosure or no disclosure at all).

2 Model

A principal (she) designs a contest to motivate $n \geq 2$ agents (he) to spend effort. The contest consists of a termination rule specifying *when* the contest will end, a rule for allocating a prize, whose value we normalize to \$1, and a feedback policy stipulating the information transmitted to each agent at every moment in time. We formalize these objects below.

At each instant t of continuous time, each player observes any message sent according to the feedback policy and decides whether to spend effort. Effort is costly and can only be monitored by the principal via a binary noisy signal, which we call "success." In an innovation-contest application, this signal may represent achieving a pre-specified target; in a promotion application, it may represent exceeding an exogenous bar for promotion. If player i spends effort $a_{i,t} \in \{0, 1\}$, he incurs cost at rate $ca_{i,t}$, where c represents the (constant) marginal cost of effort. While the agent spends effort, success arrives stochastically with constant instantaneous rate $\lambda > c$. That is, a player can achieve success at most once, and conditional on not having succeeded by t, effort for an additional duration dt produces a success during the interval $(t, t + dt)$ with probability λdt.[7]

The principal observes successes but not efforts. Each player observes his own effort, but not others' efforts or successes. Whether agents observe their *own* success or not, or do so probabilistically, is immaterial for our results. For concreteness, and to give the principal maximal flexibility, we assume that they do not observe them.

The principal's feedback policy specifies a message that she transmits to each agent at every moment as a function of her past observations and past messages. An example of a feedback policy that will be important for our results, and which we denote $\mathcal{M}^{pronto}$, is the one where the principal informs each player immediately if and

[7] A constant hazard rate means there is no notion of progress or knowledge accumulation over time. In Section 5.3, we extend our model to allow for an increasing hazard rate.

when he succeeds, but otherwise keeps silent. Alternative policies might publicly or privately inform agents about their or their rivals' successes, perhaps probabilistically, or inform them about the feedback conveyed to rivals, and so forth.

The principal's termination rule, τ, ends the contest possibly randomly and possibly as a function of the principal's past observations. The prize is then awarded according to the allocation rule, which specifies a share of the prize (or, equivalently, a probability of winning the prize) q_i for each player i, with $\sum_i q_i \leq 1$, as a function of the history of successes. For example, a *winner-takes-all* contest awards the entire prize ($q_i = 1$) to the first player i to have succeeded, whereas an *egalitarian* contest divides the prize equally among all players who have succeeded. Note that both these types of contest are efficient in the sense that the entire prize is awarded if and only if at least one player has succeeded.

When the contest ends, player i's ex-post payoff is

$$u_i = q_i - c \int_0^{\tau} a_{i,t} dt.$$

There is no discounting and players maximize their expected payoff.

The principal designs the termination rule, prize allocation rule, and feedback policy with the goal that the expected total effort in a Bayesian Nash equilibrium (hereafter *equilibrium*) of the resulting contest is maximal among Bayesian Nash equilibria of a given set of contests. In this formulation of the objective, the principal cares only about effort, not successes, and cares about players' effort even after they have succeeded. Our results would be unchanged, though, if the principal instead sought to maximize the total number of successes.[8]

3 A Sufficiency Result

Finding a fully optimal contest, inclusive of feedback policy, is in principle a daunting task. All of the choice variables are high-dimensional objects, as they can condition on the path of past successes and prior feedback. Thus, it is not even clear how to formulate the appropriate optimization problem.

Our first lemma offers an opportunity to overcome this challenge by means of a

[8]Our contest would also be optimal if the principal cared about the first success alone, although simpler designs, such as winner-takes-all, would be optimal in this case too.

simple sufficient condition for optimality.

Lemma 1. *A contest is guaranteed to be optimal if, in equilibrium:*
(i) the prize is awarded with probability one, and
(ii) each agent earns zero rents.

Intuitively, a contest that awards the prize with the maximum possible probability also maximizes the players' combined surplus; if the agents keep none of this surplus, it must all go to the principal.

To formally establish this result, note that for any contest and equilibrium effort profile, we can rewrite the principal's payoff as

$$\mathbb{E}\sum_{i=1}^{n}\int_{0}^{\tau} a_{i,t}dt = \frac{\sum_{i=1}^{n}\mathbb{E}[q_i] - \sum_{i=1}^{n}\mathbb{E}[u_i]}{c}.$$

The first term in the numerator represents the total prize awarded; the second term represents the agents' rents. The total prize awarded is bounded from above by one, whereas the agents' rents are bounded from below by zero. Therefore, if there exists a contest that attains these bounds (and so the principal's payoff is $1/c$), it must be optimal. *Q.E.D.*

While the condition in Lemma 1 is a stringent one, we shall see that there indeed exists a contest that satisfies it.

4 Optimal Contest

Here we establish our main result. Our goal is to characterize a contest that satisfies both criteria in Lemma 1 under the assumption that the parameters of the model satisfy $n > \lambda/c$. This assumption means that there are enough competitors for a contest to be desirable in the first place. When the assumption fails, the principal could do as well by reserving $1/n$-th of the prize for each agent and contracting with each one individually.

Because agents are risk neutral, without loss of generality, we can (and henceforth will) restrict attention to contests where an agent wins a positive share of the prize only if he succeeds. Fixing an equilibrium of a given contest, define for each agent i the *reward function*

$$R_{i,t} = \mathbb{E}\left[q_i \mid \text{agent } i \text{ succeeds at } t\right], \tag{1}$$

which represents agent i's expected share of the prize conditional on succeeding at time t.[9] These functions will allow us to analyze the incentives faced by each agent separately.

Our proposed contest, which we call *cyclical-egalitarian*, has three properties. First, its prize allocation is egalitarian; that is, it divides the prize equally among all players who have succeeded regardless of when they happened to do so. Second, agents are fully apprised of their own successes via the $\mathcal{M}^{pronto}$ feedback policy. Lastly, it has a cyclical termination rule as follows: the principal sets a provisional deadline T^*; if at least one agent has succeeded by that time, the contest ends; otherwise, the principal informs all agents that no one has yet succeeded and restarts the contest, again with a provisional deadline T^*. The contest continues in this manner until at least one agent has succeeded by the time the next provisional deadline is reached.

In order to extract all rents from the agents, the provisional deadline is set just long enough that the agents are indifferent between working and not during the entire length of the cycle (unless they have already succeeded). Formally, the termination rule is described by the stopping time

$$\tau^* = \inf\{t : t = kT^*,\ k \in \mathbb{N}, \text{ and at least one agent has succeeded}\},$$

where T^* is the unique solution to $\left(1 - e^{-n\lambda T^*}\right) / \left(n(1 - e^{-\lambda T^*})\right) = c/\lambda$.

Proposition 1. *Assume $n > \lambda/c$. The contest with an egalitarian prize, the cyclical termination rule τ^*, and the feedback policy $\mathcal{M}^{pronto}$ is optimal. In this contest, at least one agent succeeds and hence the prize is awarded with probability one. Moreover, each agent obtains 0 expected utility and the principal's profit is $1/c$.*

This cyclical-egalitarian contest is optimal because it meets both requirements of Lemma 1: since the provisional deadline keeps extending if no agent has succeeded, the contest awards the prize with probability 1; moreover, the $\mathcal{M}^{pronto}$ feedback policy, egalitarian allocation rule, and provisional deadline T^* act together to grant the agents no rents.[10]

[9] If the prize is indivisible, $R_{i,t}$ can equivalently be interpreted as the probability that agent i wins the prize conditional on succeeding at t.

[10] Because agents are indifferent between working and shirking before they succeed, the contest admits another equilibrium in which one of the agents never works. This equilibrium can be eliminated by shrinking the cycle length by an infinitesimal amount. This modified contest admits only the equilibrium in which all agents work until they succeed, while giving up only arbitrarily small rents.

To formally establish this proposition, we show that the cyclical-egalitarian contest has an equilibrium where all agents work until either they succeed or the contest ends (and hence the prize is awarded with probability 1) and where all their continuation payoffs are zero. Let $p_{i,t}$ denote agent i's belief at time t that he has succeeded, and observe that his flow payoff is $(1-p_{i,t})\lambda R_{i,t} - c$ if he works, and zero otherwise.

Now suppose that all of agent i's rivals work until they succeed. Because the allocation rule is egalitarian and the contest ends at the next provisional deadline if any agent has succeeded, agent i's expected reward conditional on success is

$$R_{i,t} = \mathbb{E}\left[\frac{1}{1+M}\right] = \frac{1-e^{-\lambda n T^*}}{n(1-e^{-\lambda T^*})} = \frac{c}{\lambda}, \tag{2}$$

where $M \sim Binom(n-1, 1-e^{-\lambda T^*})$ is the number of rivals who succeed by the next provisional deadline, the second equality follows from writing the binomial sum and rearranging terms, and the third equality follows from the definition of T^*.[11]

The feedback policy $\mathcal{M}^{pronto}$ ensures that $p_{i,t} = 0$ until this agent succeeds, at which moment his belief jumps to one. This implies that each agent's flow payoff, and hence his continuation payoff, is always held at zero irrespective of his effort, and so working until he succeeds is incentive compatible. Because agents are symmetric, there is indeed an equilibrium with the desired properties. *Q.E.D.*

As it turns out, there are other optimal contests as well. However, all these other contests have in common with the cyclical-egalitarian contest that: (i) they keep agents fully apprised of their own successes, and so each agent quits as soon as he succeeds; (ii) they have an egalitarian prize structure; that is, they prescribe reward functions $R_{i,t}$ that are time-invariant and symmetric across players; and (iii) they do not end until at least one agent has succeeded. As a result, they differ only in the details of the termination rule.[12]

[11] To be specific, letting $p = 1-e^{-\lambda T^*}$, we have

$$\mathbb{E}\left[\frac{1}{1+M}\right] = \frac{1}{np}\sum_{k=0}^{n-1}\binom{n}{k+1}p^{k+1}(1-p)^{n-1-k} = \frac{1}{np}\sum_{j=1}^{n}\binom{n}{j}p^j(1-p)^{n-j} = \frac{1-(1-p)^n}{np},$$

where the first equality follows from expanding the binomial coefficient and manipulating the expression, the second one results from the change of variables $j = k+1$, and the last one follows from collecting terms.

[12] One example is a modified version of the contest given in Proposition 1 with an arbitrary provisional deadline $T > T^*$ and where agents are asked to work only a fraction T^*/T of each cycle.

The reason it is necessary to immediately inform agents of their own successes—despite the principal preferring that they keep working—is that they would otherwise be able to obtain rents from the principal by strategically withdrawing effort. Intuitively, a hard-working agent who is not fully informed will gradually come to believe that he has already succeeded, and hence will only continue working if his expected reward increases. But this would allow him to secure rents by initially withdrawing effort and working only once the expected reward has grown.[13] The egalitarian rule is necessary, in turn, because given the $\mathcal{M}^{pronto}$ feedback policy, for a contest to extract all rents, each agent's expected reward conditional on succeeding must be c/λ regardless of when he happens to succeed. Non-egalitarian contests are unable to offer such time-invariant rewards.

Because in our model agents are risk neutral and there are no aggregate productivity shocks, one may wonder why a contest—rather than contracting individually with each agent—is necessary in the first place. The advantage of the contest is that it allows the principal to pool the agents' incentive constraints in the face of a fixed prize budget; that is, prize money not awarded to one agent can be awarded to another. In pooling these constraints, the principal benefits from strategically controlling what each agent learns about his peers: by (temporarily) keeping agents in the dark, she can motivate them to continue working even after their peers have already succeeded. The benefit of the contest can be seen from the fact that the sum of the agents' marginal benefits of effort, $\sum_i R_{i,t} = nc/\lambda$ where $R_{i,t}$ is defined in (2), exceeds the \$1 prize money (since $nc > \lambda$), whereas if the principal contracted with each agent individually, the marginal benefit of effort for any given agent would be equal to the prize money reserved for him.[14]

We conclude this section with some remarks:

i. Because the cyclical-egalitarian contest keeps agents fully apprised of their own

[13] To see this more formally, consider a contest that is intended to grant zero rents and suppose that there are times where an agent is expected to exert effort and yet $p_{i,t} > 0$. Then there must be a time interval in which the agent is supposed to work and yet his belief $p_{i,t}$ strictly increases. So that he is willing to work meanwhile earning 0 rents, $(1 - p_{i,t})\lambda R_{i,t}$ must equal c. But then he can pause effort during the first half of this interval so that his private belief diverges from, and is strictly smaller than, the equilibrium belief (as he knows that he cannot have possibly succeeded while shirking), which in turn allows him to extract rents during the second half.

[14] Another way to reap the benefits of pooling would be for the principal to meet her prize budget constraint only in ex-ante terms (i.e., spend \$1 only in expectation). If she were able to do so, individual contracting would suffice.

successes, that contest would remain optimal if agents were able to observe these successes directly, or observe them probabilistically.[15]

ii. While designed to maximize total effort, the cyclical contest also maximizes the expected total number of successes, as it transforms 100% of the prize into effort (rather than rents) and never motivates already successful agents to keep working.

iii. To implement the optimal contest, the principal needs to commit to only one cycle at a time: if no agent has succeeded by the end of a cycle, it is in her interest to commit to another identical one.

iv. The principal would be no better off with a more precise monitoring technology as, despite her imperfect (binary) signal, she is able to convert the entire prize into effort.

v. If agents were uncertain about the number of rivals they face, the cyclical-egalitarian contest would remain optimal provided agents shared a common prior over that number, with support on $(\lambda/c, \infty)$. The provisional deadline T^* would now satisfy $\mathbb{E}[(1 - e^{-n\lambda T^*})/(n(1 - e^{-\lambda T^*}))] = c/\lambda$.[16]

4.1 A family of cyclical-egalitarian contests

The contest in Proposition 1 is a member of a larger family of cyclical-egalitarian contests that differ only in the length of their cycle. Such contests transform 100% of the prize into any desired combination of effort and agent rents, and hence are on the Pareto frontier.

Corollary 1. *For any $T \leq T^*$, the cyclical-egalitarian contest with $\mathcal{M}^{pronto}$ feedback policy and cycle length T induces all agents to work until either they succeed or the contest ends. Such contest delivers total expected effort*

$$\mathbb{E}\sum_{i=1}^{n}\int_0^{\tau} a_{i,t}dt = \frac{n}{\lambda} \times \frac{1 - e^{-\lambda T}}{1 - e^{-\lambda n T}},$$

which is increasing in T, and delivers rents $\sum_i u_i = 1 - c\,\mathbb{E}\sum_i \int_0^{\tau} a_{i,t}dt$.

[15] What is important is that agents do not observe the successes of their rivals. Note that although each agent has an incentive to tell his rivals that he has succeeded to discourage them, they have no reason to believe him.

[16] Every time the deadline is extended, seeing that nobody has yet succeeded, agents will update their beliefs about n, and so each subsequent deadline must account for those updated beliefs.

The commonly used *winner-takes-all* contest—which ends as soon as the first agent succeeds and awards the entire prize to that agent—is a special case of this family with cycle length equal to 0.[17] The cyclical-egalitarian design thus adds a degree of flexibility—cycle length—to the winner-takes-all design. By controlling that parameter, the principal is able to extract as much rent as she desires while transforming it into effort at rate $1/c$; the winner-takes-all contest instead selects a particular point along the Pareto frontier (see Figure 1).

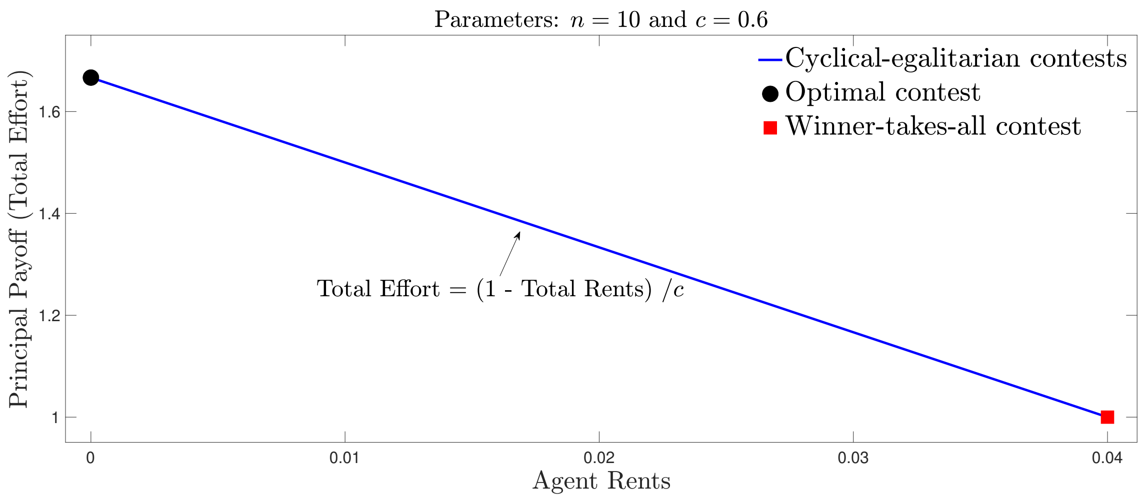

Figure 1: The payoff frontier of the cyclical-egalitarian contest as the cycle-length varies from zero to T^*.

5 Extensions

Here we discuss three extensions that suggest a degree of robustness to our findings.

5.1 Discounting

Assume players are impatient and discount the future at a common rate $r > 0$. Thus, if the contest ends at τ, agent i's payoff u_i is his expected discounted prize $e^{-r\tau}q_i$ minus his discounted cost of effort $c\int_0^\tau e^{-rt}a_{i,t}dt$, and the principal's payoff is the

[17] Because agents learn about their peers' successes at the end of each cycle, a zero cycle length means that agents are always fully informed, and barring zero-measure events, are guaranteed to win 100% of the prize as soon as they succeed.

total discounted effort

$$\sum_{i=1}^{n} \int_0^{\tau} e^{-rt} a_{i,t} dt = \frac{e^{-r\tau} \sum_{i=1}^{n} q_i - \sum_i u_i}{c}.$$

A contest is guaranteed to be optimal if it simultaneously maximizes the first term in the numerator and minimizes the second. This would require awarding the prize as soon as the first agent succeeds (so as to minimize discounting) while also surrendering zero rents to the agents.[18] Were it possible to achieve both goals at once, the principal's payoff would be

$$\overline{\Pi} := \frac{\lambda n}{c(r + \lambda n)}.$$

Unfortunately, meeting both goals simultaneously is impossible: granting the prize immediately would allow the agents to earn rents; conversely, extracting all rents would require delaying the prize.

Because of that impossibility, it is challenging to find a fully optimal contest. We can show nonetheless that a cyclical contest akin to that in Proposition 1, but with a shorter provisional deadline, is approximately optimal in the sense of delivering a payoff close to $\overline{\Pi}$, provided r is close to 0, n is large, or λ/c is close to 1.

Proposition 2. *Consider the contest with an egalitarian prize, the $\mathcal{M}^{pronto}$ feedback policy, and a cyclical termination rule with provisional deadline T^{**} chosen such that the agents' incentive constraints bind at the start of each cycle. This contest is approximately optimal when r is close to 0, n is large, or λ/c is close to 1.*

The reason the cycle length must be shortened is that, due to discounting, if agents were to be indifferent between working and not at a given time in the cycle, they would strictly prefer not to work before then. To minimize rents, the contest in Proposition 2 uses a cycle length just long enough that agents are barely willing to work at the start of each cycle, and earn rents after that.[19]

That this contest is approximately optimal when r converges to zero can be seen from the fact that the deadline converges to the original one, T^*. This means agents

[18] In this case, $\mathbb{E}[e^{-r\tau} \sum_{i=1}^{n} q_i] = \int_0^{\infty} e^{-rt} \lambda n e^{-\lambda n t} dt = \lambda n/(r + \lambda n)$.

[19] Another possibility is to use a stochastic-length cycle, as a random early termination is a way to counteract the agents' preference to delay effort. This more complex design allows the principal to extract all rents, but performs very similarly to the simpler deterministic design.

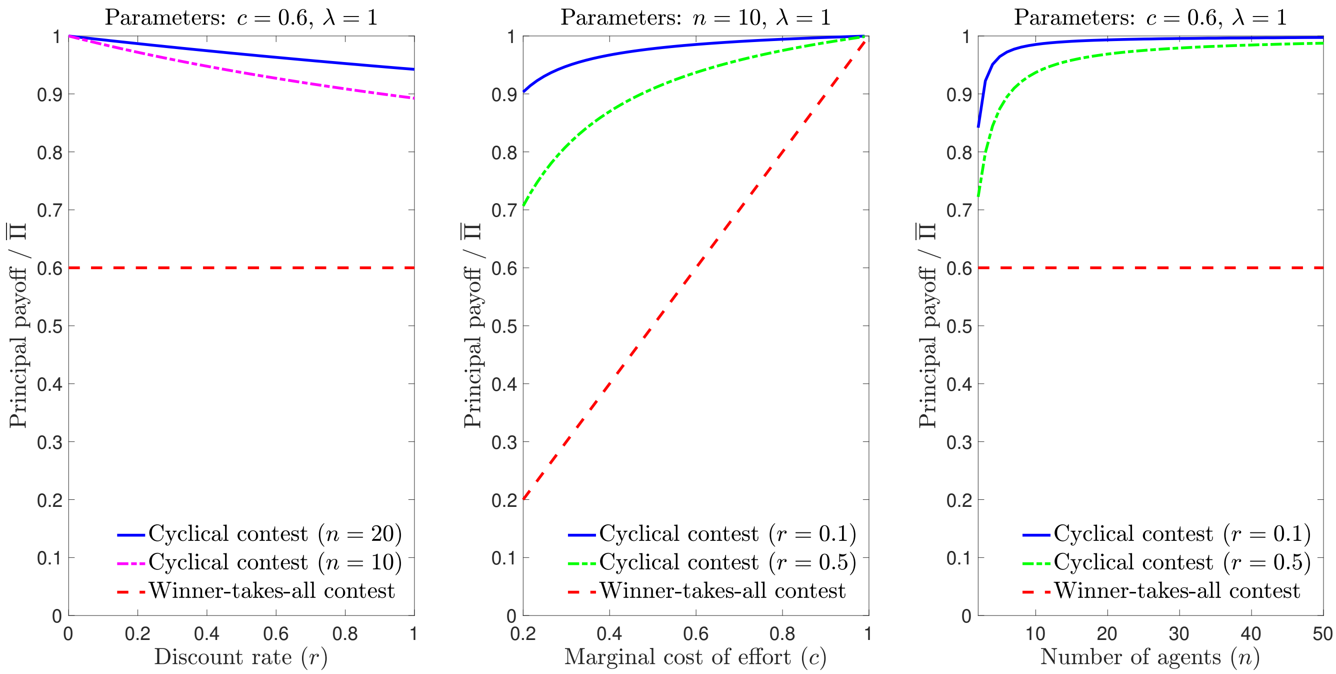

Figure 2: Principal's payoff as a fraction of the (unattainable) upper bound $\overline{\Pi}$ for the cyclical-egalitarian and the winner-takes-all contest.

earn close-to-zero rents, and, moreover, there is minimal loss in overall surplus due to the prize being delayed. When the number of contestants grows to infinity or the ratio of effort productivity λ/c falls to one, the deadline must converge to zero so that agents are willing to work. This means that regardless of r they earn minimal rents, and because there is also minimal delay in awarding the prize once the first success has arrived, the principal's payoff must be close to $\overline{\Pi}$.

Figure 2 illustrates the contest's performance relative to the upper bound $\overline{\Pi}$ when the parameters are away from these limits. The lower dashed curve in each panel depicts, for comparison, the outcome of the commonly used winner-takes-all contest, which gives the principal a payoff of $n/(r + \lambda n)$.

5.2 Limited Commitment

Crucial for our results is the assumption that the principal is able to commit to a real-time feedback policy and a history-dependent termination rule—on top of a prize-allocation rule. If the principal was instead unable to credibly communicate with the agents mid-contest—e.g., due to her interest in having agents work for as long as possible—and could only commit to ending the contest at a pre-specified deadline, it would no longer be possible to meet the sufficiency conditions of Lemma 1. An optimal (no-feedback, fixed-deadline) contest can be obtained nonetheless.

To begin, restrict attention to contests where each agent works continuously over some interval $[0, T_i]$.[20] Now express agent i's utility from working as prescribed as $\int_0^{T_i} \lambda e^{-\lambda t} R_{i,t} dt - cT_i$, where the reward function $R_{i,t}$ is as defined in (1).

The following lemma provides a necessary condition for incentive compatibility.

Lemma 2. *Consider a no-feedback contest that gives agent i the reward function $R_{i,t}$. Working continuously throughout $[0, T_i]$ is incentive compatible for this agent only if*

$$\lambda e^{-\lambda t} R_{i,t} - \int_t^{T_i} \lambda^2 e^{-\lambda s} R_{i,s} ds \geq c \text{ for all } t \in [0, T_i]. \tag{IC}$$

This incentive constraint states that the marginal benefit of effort at time t, which is captured by the left-hand side, should be no smaller than the marginal cost. To understand the expression for the marginal benefit, note that the first term is the instantaneous marginal benefit of effort at time t. The second term captures a forward-looking incentive effect: success today precludes success in the future. In particular, $\lambda^2 e^{-\lambda s}$ is the amount by which the success probability at some future date s is reduced when the agent spends effort at date t. The second term thus aggregates the reduction in future instantaneous benefits that results from spending effort in the current date.

To find an optimal contest, we can solve

$$\max_{T, \{T_i\}, \{q_i\}} \sum_{i=1}^{n} T_i \quad \text{subject to (IC) and } T_i \leq T \text{ for all } i, \tag{3}$$

and then verify that the contest indeed has an equilibrium in which each agent i works continuously until T_i. To this end, define T^{EGA} to uniquely solve $(1 - e^{-\lambda n T})/(n(1 - e^{-\lambda T})) = ce^{\lambda T}/\lambda$.

Proposition 3. *The egalitarian contest with deadline T^{EGA} is optimal among no-feedback, fixed-deadline contests. This contest admits an equilibrium where each agent works continuously until the deadline.*

[20] This restriction is without loss of generality because neither the principal nor any agent gains from delays in effort.

The simple intuition for this result is that non-egalitarian contests, unlike the egalitarian one, create unequal effort incentives over time, leading to potential gaming by the agents in how they time their effort. The only way to prevent this gaming is to spend additional money on the prize, which the principal does not have.

Here is a more detailed heuristic argument that highlights the role of Lemma 2. For brevity, let $\lambda = 1$ and restrict attention to symmetric contests with symmetric equilibria. The constant reward function $R_{i,t} = e^{T^{\text{EGA}}}c$, which corresponds to the egalitarian contest, satisfies constraint (IC) with equality at all $t \leq T^{\text{EGA}}$. Figure 3 plots the corresponding instantaneous marginal benefit schedule $e^{T^{\text{EGA}}-t}c$, together with the agent's marginal cost. Notice that at every $t' \leq T^{\text{EGA}}$, the instantaneous marginal benefit exceeds c by exactly area 1, which corresponds to the integral on the left-hand side of (IC).

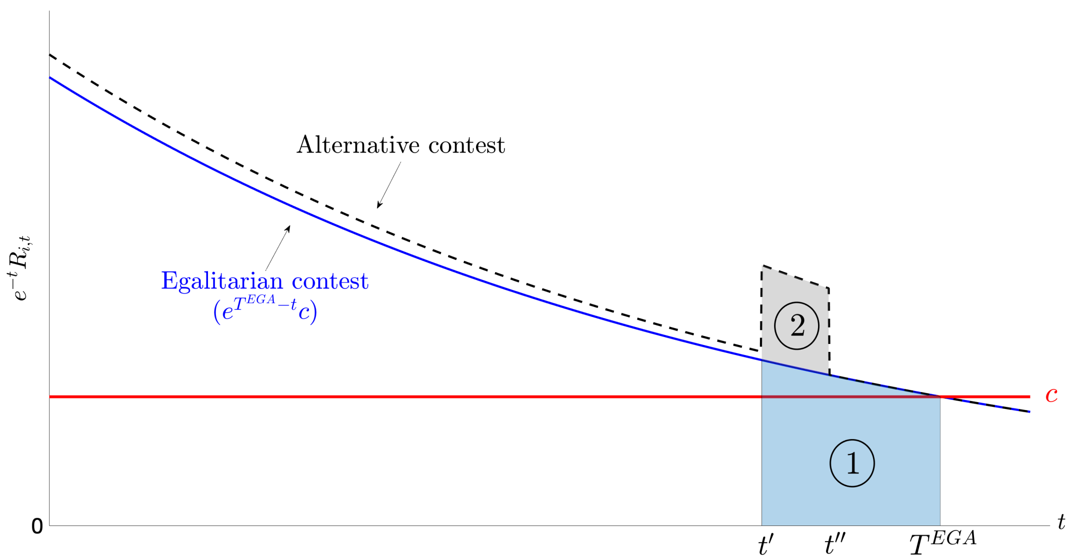

Figure 3: Meeting the incentive constraint.

Consider now a non-egalitarian contest (i.e., one with a non-constant reward schedule) that attempts to implement the same total effort as the egalitarian one. As illustrated in the figure, (IC) implies that if there is a time interval $[t', t''] \leq T^{\text{EGA}}$ where this alternative schedule exceeds the egalitarian one, it must also exceed the egalitarian schedule at all times prior to t', since the integral in (IC) grows from area 1 to area 1 + 2. In other words, a higher reward at any future date forces a higher reward today, as otherwise the agent would prefer to pause his effort today and gain access to this higher future gain. Thus, in order to implement the same effort as the egalitarian contest, the reward schedule would need to be uniformly higher, which is only possible with a prize greater than \$1.

5.3 Increasing hazard rate

In some settings, the agents' instantaneous probability of success might grow as they work and make progress on the problem and accumulate knowledge.[21] To capture this possibility, let $F(t)$ denote the probability that an agent succeeds at or before date t if he works continuously until that time, and suppose the hazard rate $\lambda_t = F'(t)/[1-F(t)]$ exists and is weakly increasing. This implies that if an agent has spent s units of effort by some date t, his hazard rate is λ_s.

Under the assumption that $\lambda_t \in (c, nc)$ for all t, and is differentiable almost everywhere, the optimal contest is similar to the cyclical-egalitarian contest with constant hazard rate characterized in Proposition 1, except that the length of each provisional deadline is stochastic. The termination rule now operates as follows: At date 0, the principal privately draws $T_1 \sim H^0(\cdot)$, where the distribution function $H^t(T) := 1 - e^{-\int_t^T \gamma_s^t ds}$ and

$$\gamma_s^t := \frac{c\dot{\lambda}_s}{\lambda_s^2}\left[\frac{1-e^{-n\int_t^s \lambda_v dv}}{n\left(1-e^{-\int_t^s \lambda_v dv}\right)} - \frac{c}{\lambda_s}\right]^{-1}. \tag{4}$$

If at least one agent has succeeded by T_1, the contest ends and the prize is awarded according to the egalitarian rule. Otherwise, a new cycle begins, and ends at the random date $T_2 \sim H^{T_1}(\cdot)$, and so forth.[22] Let τ^{***} denote the termination rule thus defined.

Proposition 4. *Assume $\lambda_t \in (c, nc)$ for all t, is weakly increasing, and its derivative exists almost everywhere. The contest with egalitarian prize, the cyclical termination rule τ^{***}, and the feedback policy $\mathcal{M}^{pronto}$ is optimal.*

To explain the logic of this design, notice first that Lemma 1 remains valid: a contest which awards the entire prize and concedes zero rents to the agents gives the principal profit $1/c$, and is guaranteed to be optimal. Given that the feedback

[21]For instance, contestants might be sampling among a finite set of possible solutions, or they might need to accumulate a number of intermediate Poisson successes before they solve the problem.

[22]It is shown that γ_s^t is non-negative and $H^t(\cdot)$ has finite support; that is, for every t, there is a finite cutoff date such that, as it is approached, $\gamma_s^t \to \infty$ and the cycle ends arbitrarily quickly.

policy is $\mathcal{M}^{pronto}$, it suffices to show that each agent's expected reward conditional on success, $R_{i,t}$, is equal to c/λ_t until he succeeds. Then, by the same argument as in Section 4, there exists an equilibrium in which agents work continuously until they succeed and earn zero rents.[23]

What differs from the baseline model is that an earlier success must be rewarded more dearly than a later one, once the hazard rate has had a chance to grow. Proposition 4 shows that this can be achieved using an egalitarian prize with a stochastic cyclical termination rule: since an agent who succeeds early on will expect to share the prize with fewer of his rivals, he secures a greater reward. By choosing the distribution of each cycle's length, it is possible to fine-tune $R_{i,t}$ so that it exactly equals c/λ_t.

6 Conclusion

We have proposed a contest with an egalitarian prize, a cyclical structure involving a periodic resetting of the contest, and a partial type of feedback: contestants are kept fully apprised of their own successes, but are only periodically informed about their rivals' successes so as to not discourage further effort. In our setting, this contest manages to convert 100% of the prize money into effort (i.e., is maximally efficient) as it manages to extract all rents from the contestants. For this reason, it is able to deliver large gains relative to commonly-used contests, including the winner-takes-all design. Because of its relative simplicity, it may also be attractive from an applied point of view.

The proposed contest is a member of a larger family of "cyclical-egalitarian contests" that manage to transform 100% of the prize money into any desired combination of effort and total rents for the agents, with an efficient transformation rate between the two equal to the inverse marginal cost of effort. The winner-takes-all contest is one member of this family.

Our model has abstracted from features that may be relevant for specific applications, including technological asymmetries across players and a decreasing hazard rate of success. When these features are present, the optimal control of information

[23] A crucial observation for this argument is that if an agent ever shirked prior to date t, then his hazard rate at t would be strictly smaller than λ_t, and so he would strictly prefer to shirk at every subsequent date. Therefore, agents cannot extract positive rents by strategically withdrawing effort.

is likely to be more complex. We leave these possibilities for future work.

References

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

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

Bimpikis, Kostas, Shayan Ehsani, and Mohamed Mostagir. 2019. "Designing dynamic contests." *Operations Research*, 67(2): 339–356.

Bonatti, Alessandro, and Johannes Horner. 2011. "Collaborating." *American Economic Review*, 101(2): 632–663.

Drugov, Mikhail, and Dmitry Ryvkin. 2019. "Optimal prizes in tournaments with risk-averse agents." Working paper.

Drugov, Mikhail, and Dmitry Ryvkin. 2020. "Tournament rewards and heavy tails." *Journal of Economic Theory*, 190: 105116.

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

Fang, Dawei, Thomas Noe, and Philipp Strack. 2018. "Turning up the heat: The demoralizing effect of competition in contests." *Journal of Political Economy*, 71.

Fershtman, Chaim, and Uri Gneezy. 2011. "The tradeoff between performance and quitting in high power tournaments." *Journal of the European Economic Association*, 9(2): 318–336.

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

Gross, Daniel P. 2017. "Performance feedback in competitive product development." *The RAND Journal of Economics*, 48(2): 438–466.

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

Khorasani, Sina. 2020. "Dynamic Innovation Contests and Information Design." Working paper.

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

Lemus, Jorge, and Guillermo Marshall. 2021. "Dynamic Tournament Design: Evidence from prediction contests." *Journal of Political Economy*, 129(2): 383–420.

Letina, Igor, Shuo Liu, and Nick Netzer. 2020. "Optimal Contest Design: A General Approach." Working Paper.

Lim, Noah, Michael J Ahearne, and Sung H Ham. 2009. "Designing sales contests: Does the prize structure matter?" *Journal of Marketing Research*, 46(3): 356–371.

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

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.

Olszewski, Wojciech, and Ron Siegel. 2016. "Large contests." *Econometrica*, 84(2): 835–854.

Olszewski, Wojciech, and Ron Siegel. 2020. "Performance-maximizing large contests." *Theoretical Economics*, 15(1): 57–88.

Siegel, Ron. 2009. "All-Pay Contests." *Econometrica*, 77(1): 71–92.

Siegel, Ron. 2010. "Asymmetric Contests with Conditional Investments." *American Economic Review*, 100(5): 2230–2260.

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 Proof of Corollary 1

Consider the cyclical-egalitarian contest with the $\mathcal{M}^{pronto}$ feedback policy and cycle-length T, where $T \leq T^*$. Let $u_{i,t}$ denote agent i's continuation payoff at time t conditional on not having yet succeeded. By the derivation in (2), this agent's reward function $R_{i,t} = (1-e^{-\lambda n T})/[n(1-e^{-\lambda T})]$, and note that it is time-invariant, it decreases in T, and equals c/λ if $T = T^*$.

An agent finds it optimal to work at t if and only if $\lambda(R_{i,t} - u_{i,t}) \geq c$. That is because if he works at t, he succeeds with rate λ, in which case he earns reward $R_{i,t}$ meanwhile forgoing his continuation payoff $u_{i,t}$. Naturally, for working to be incentive compatible, this must exceed the marginal cost of effort, c.

Next, we compute $u_{i,t}$. We guess and verify later that this agent finds it optimal to work until he succeeds or the contest ends. Notice that because the contest resets at the beginning of each cycle, $u_{i,t} = u_{i,t+kT}$ for all t and $k \in \mathbb{N}$, and so it suffices to focus on $t \in [0, T]$. We have

$$u_{i,t} = [1 - e^{-\lambda(T-t)}](R_{i,t} - c/\lambda) + e^{-\lambda n T + \lambda t} u_{i,0},$$

where the first term captures this agent's rents during the current cycle, and the second term represents the probability that the contest continues for another cycle conditional on him not having succeeded by t times his continuation payoff. Using that $u_{i,0} = (1 - e^{-\lambda T})(R_{i,t} - c/\lambda)/(1 - e^{-\lambda n T})$, we can compute

$$u_{i,t} = (R_{i,t} - c/\lambda)\left[1 - e^{\lambda(t-T)}\left(1 - \frac{e^{\lambda T} - 1}{e^{\lambda n T} - 1}\right)\right],$$

and it is straightforward to verify that it decreases in t. This observation implies that if effort is incentive compatible at the beginning of a cycle, it is also incentive compatible at every moment thereafter. Therefore, agent i finds it optimal to work at every moment until he succeeds or the contest ends if and only if $R_{i,t} - c/\lambda \geq u_{i,0}$, which is satisfied for any $T \leq T^*$.

Using the expression for $u_{i,0}$, the principal's objective

$$\mathbb{E}\sum_{i=1}^{n}\int_0^{\tau} a_{i,t}dt = \frac{1-\sum_{i=1}^{n}u_{i,0}}{c} = \frac{n}{\lambda}\times\frac{1-e^{-\lambda T}}{1-e^{-\lambda n T}},$$

while the agents' rents can be written as $\sum_i u_{i,0} = 1 - c\,\mathbb{E}\sum_i\int_0^{\tau} a_{i,t}dt$. □

A.2 Proof of Proposition 2

To prove the proposition, we proceed in two steps. First, we consider a contest with a cyclical termination rule, where each cycle has length T, the egalitarian prize allocation rule, and the $\mathcal{M}^{pronto}$ feedback policy. We derive each agent's payoff function, and we show that there exists a T such that, anticipating all other contestants to work until they succeed, each agent's incentive compatibility constraint binds at the beginning of each cycle and is slack thereafter. This implies that there exists an equilibrium in which all agents work continuously until they succeed or the contest ends. In the second step, we compute the principal's payoff, and show that it converges to the upper bound, $\overline{\Pi}$, as $r\to 0$, $\lambda/c\to 1$, or $n\to\infty$.

Step 1. Because the contest resets every cycle, without loss of generality, we can focus on the cycle that starts at time zero. Recall that each agent's expected (undiscounted) reward conditional on succeeding at $t\in[0,T]$ when he anticipates all of his rivals to work until they succeed is equal to $R_{i,t} = (1-e^{-n\lambda T})/[n(1-e^{-\lambda T})]$.

Letting $u_{i,t}$ denote agent i's continuation payoff at t, this agent finds it optimal to work at t if and only if $\lambda(R_{i,t}-u_{i,t})\geq c$. Intuitively, if an agent works at t, he succeeds with rate λ, in which case he earns reward $R_{i,t}$ but forgoes his continuation payoff $u_{i,t}$. Naturally, for working to be incentive compatible, this net marginal benefit must exceed the marginal cost c. Next, we derive an explicit expression for $u_{i,t}$. Suppose that effort is indeed incentive compatible for each agent until he succeeds. Then

$$u_{i,t} = \int_t^T e^{-(r+\lambda)(s-t)}\left[\lambda e^{-r(T-s)}R_{i,s}-c\right]ds + e^{-(r+\lambda)(T-t)-(n-1)\lambda T}u_{i,0}, \tag{5}$$

where the expression in the integrand represents the agent's flow payoff within the current cycle, while the second term is the probability that no one succeeds during the current cycle times the agent's continuation payoff, exploiting the fact that the contest resets at the end of each cycle.

Using the last observation and the expression for $R_{i,t}$, each agent's payoff at the beginning of each cycle

$$u_{i,0} = \frac{\frac{e^{-rT}\left[1-e^{-n\lambda T}\right]}{n} - \frac{c}{r+\lambda}\left[1-e^{-(r+\lambda)T}\right]}{1-e^{-(r+n\lambda)T}}. \tag{6}$$

Substituting (5) into the incentive compatibility condition, $R_{i,t} \geq c/\lambda + u_{i,t}$ if and only if

$$\frac{1-e^{-n\lambda T}}{n\left(1-e^{-\lambda T}\right)} \geq \frac{rc}{\lambda(r+\lambda)}\left[e^{(r+\lambda)(T-t)}-1\right] + \frac{c}{\lambda} + e^{-(n-1)\lambda T}u_{i,0}. \tag{7}$$

Define T^{**} to be the largest T such that the above inequality binds at $t=0$. To see why such a T^{**} exists, first, note that the left-hand side is strictly decreasing in T, it converges to 1 as $T \to 0$, and to c/λ as $T \to T^*$, where T^* is the provisional deadline in Proposition 1. On the other hand, using the expression for $u_{i,0}$, one can (tediously) show that as $T \to 0$, the right-hand side converges to

$$\frac{c}{\lambda} + \frac{\lambda - c}{r+n\lambda} \in \left(\frac{c}{\lambda}, 1\right),$$

and it is strictly greater than c/λ for any $T > 0$. Since both sides are continuous in T, by the intermediate value theorem, there exists a $T^{**} \in (0, T^*)$ such that (7) binds at $t=0$. Finally, observe that the right-hand side of (7) is strictly decreasing in t, whereas the left-hand side is constant. Together with the fact that (7) binds at $t=0$, this implies that it is satisfied for every $t \in (0, T^{**}]$.

Since T^{**} is chosen such that it is optimal for agent i to work continuously throughout the cycle until he succeeds, by symmetry, there exists an equilibrium in which all agents work continuously until they succeed or the contest ends.

Step 2. We now compute the principal's profit when she implements the cyclical contest with the egalitarian prize allocation rule and the $\mathcal{M}^{pronto}$ feedback policy, where each cycle has length T^{**}. The principal's payoff

$$\Pi = \mathbb{E}\left[\sum_{i=1}^{n}\int_0^{\tau} e^{-rt}a_{i,t}dt\right] = \frac{\mathbb{E}[e^{-r\tau}\sum_{i=1}^{n} q_i] - \mathbb{E}[\sum_{i=1}^{n} u_{i,0}]}{c} = \frac{n}{r+\lambda} \times \frac{1-e^{-(r+\lambda)T^{**}}}{1-e^{-(r+n\lambda)T^{**}}},$$

where we have used (6) and the fact that the total surplus forms a geometric sequence and can thus be computed as

$$\mathbb{E}\left[e^{-r\tau}\sum_{i=1}^{n} q_i\right] = \left(1-e^{-n\lambda T^{**}}\right)e^{-rT^{**}} + e^{-n\lambda T^{**}}\left(1-e^{-n\lambda T^{**}}\right)e^{-2rT^{**}} + ...$$
$$= \frac{\left(1-e^{-n\lambda T^{**}}\right)e^{-rT^{**}}}{1-e^{-(r+n\lambda)T^{**}}}.$$

Define the ratio

$$\rho(T^{**}) := \frac{\Pi}{\overline{\overline{\Pi}}} = \frac{c(r+n\lambda)}{\lambda(r+\lambda)} \times \frac{1-e^{-(r+\lambda)T^{**}}}{1-e^{-(r+n\lambda)T^{**}}},$$

which is bounded from above by one. A contest is guaranteed to be optimal if $\rho(T^{**}) = 1$.

As the discount rate $r \to 0$, one can verify from (7) that $T^{**} \to T^*$, and using the definition of T^*, it follows that $\rho \to 1$. By continuity, whenever r is close to zero, the contest is approximately optimal.

Next, fix a λ and consider the limit as $c \to \lambda$. It is easy to verify that $\lim_{c\to\lambda} T^* = 0$, and since $T^{**} \in (0, T^*)$, it follows that $T^{**} \to 0$ in this limit as well. Therefore, by L'Hôpital's rule we have

$$\lim_{c\to\lambda} \rho = \lim_{c\to\lambda} \frac{c(r+n\lambda)}{\lambda(r+\lambda)} \times \frac{1-e^{-(r+\lambda)T^{**}}}{1-e^{-(r+n\lambda)T^{**}}} = 1.$$

Hence by continuity, our proposed design is approximately optimal as $\lambda/c \to 1$.

Finally, we consider the limit as $1/n \to 0$ (which is equivalent to, but more convenient as it turns out, than taking the limit as $n \to \infty$). Because $\lim_{1/n\to 0} T^* = 0$ and $T^{**} \in (0, T^*)$, it follows that $\lim_{1/n\to 0} T^{**} = 0$. Towards showing that $\lim_{1/n\to 0} \rho = 1$, define $K(T)$ to be equal to the left-hand side minus the right-hand side of (7) after multiplying both sides by $(1 - \exp(-\lambda T^{**}))$, and note that $K(T^{**}) = 0$ by the definition of T^{**}. By the implicit function theorem, we have $\partial T^{**}/\partial(1/n) = -[\partial K(T^{**})/\partial(1/n)]/[\partial K(T^{**})/\partial T^{**}]$. Defining $\eta := \lim_{1/n\to 0} \partial T^{**}/\partial(1/n)$, and assuming that this limit exists (which we verify below), it can be shown using the algebraic limit theorem and L'Hôpital's rule that

$$\lim_{1/n\to 0} -\frac{\partial K(T^{**})}{\partial(1/n)} = -1 + e^{-\lambda\eta} + \lambda\eta e^{-\lambda\eta} \text{ and } \lim_{1/n\to 0} \frac{\partial K(T^{**})}{\partial T^{**}} = \lambda e^{-\lambda\eta} - c,$$

which implies that

$$\eta = \lim_{1/n\to 0} \frac{\partial T^{**}}{\partial(1/n)} = \frac{\lim_{1/n\to 0} -\partial K(T^{**})/\partial(1/n)}{\lim_{1/n\to 0} \partial K(T^{**})/\partial T^{**}} = \frac{-1 + e^{-\lambda\eta} + \lambda\eta e^{-\lambda\eta}}{\lambda e^{-\lambda\eta} - c}, \tag{8}$$

or equivalently, $1 - e^{-\lambda\eta} = c\eta$. It is straightforward to verify that for any c and λ, there exists a unique $\eta > 0$ such that the last equality is satisfied.[24] We therefore have

$$\begin{aligned}
\lim_{1/n\to 0} \rho &= \lim_{1/n\to 0} \frac{cr}{\lambda(r+\lambda)} \times \frac{1 - e^{-(r+\lambda)T^{**}}}{1 - e^{-(r+n\lambda)T^{**}}} + \lim_{1/n\to 0} \frac{cn}{r+\lambda} \times \frac{1 - e^{-(r+\lambda)T^{**}}}{1 - e^{-(r+n\lambda)T^{**}}} \\
&= 0 + \frac{c}{r+\lambda} \times \lim_{1/n\to 0} \frac{1}{1 - e^{-(r+n\lambda)T^{**}}} \times \lim_{1/n\to 0} \frac{1 - e^{-(r+\lambda)T^{**}}}{1/n} \\
&= \frac{c}{r+\lambda} \times \lim_{1/n\to 0} \frac{1}{1 - e^{-(r+n\lambda)T^{**}}} \times \lim_{1/n\to 0} (r+\lambda) e^{-(r+\lambda)T^{**}} \frac{\partial T^{**}}{\partial(1/n)} \\
&= c \times \frac{1}{1 - e^{-\lambda\eta}} \times \eta = 1,
\end{aligned}$$

where the second, third, and fourth equalities follow from the algebraic limit theorem, L'Hôpital's rule, and the facts that $\lim_{1/n\to 0} T^{**} = 0$ and $\lim_{1/n\to 0} nT^{**} = \eta$, while the last equality follows from (8). Again by continuity, it follows that our design is approximately optimal as n grows large. □

A.3 Proof of Lemma 2

Faced with a reward function $R_{i,t}$ defined on $[0, T]$, agent i chooses his effort by solving

$$\max_{a_{i,t}} \int_0^T (\lambda R_{i,t} e^{-\lambda \int_0^t a_{i,s} ds} - c a_{i,t})\, dt.$$

Suppose that for some $T_i \le T$, this agent finds it optimal to choose $a_{i,t} = 1$ for all $t \in [0, T_i]$. Consider a deviation in which he pauses effort between times t and $t + \Delta t$

[24] Note that $\eta = 0$ is not a feasible solution: in that case, $\lim_{1/n\to 0} -\partial K(T^{**})/\partial(1/n) = 0^-$ and $\lim_{1/n\to 0} \partial K(T^{**})/\partial(T^{**}) = \lambda - c > 0$ which together would imply that $\eta = 0^-$. Then T^{**} would decrease in $1/n$, contradicting the facts that $\partial T^*/\partial(1/n) > 0$ and $T^{**} \in (0, T^*)$.

for $\Delta t > 0$. He gains

$$c\Delta t - \int_t^{t+\Delta t} \lambda R_{i,s} e^{-\lambda s} ds + \int_{t+\Delta t}^{T_i} \lambda R_{i,s} \left[e^{-\lambda(s-\Delta t)} - e^{-\lambda s} \right] ds.$$

If working continuously throughout $[0, T_i]$ is incentive compatible, this gain must be non-positive. Dividing through by Δt we have

$$c - \frac{1}{\Delta t} \int_t^{t+\Delta t} \lambda R_{i,s} e^{-\lambda s} ds + \int_{t+\Delta t}^{T_i} \lambda R_{i,s} \frac{e^{-\lambda(s-\Delta t)} - e^{-\lambda s}}{\Delta t} ds \leq 0.$$

In the limit as $\Delta t \to 0$ we have

$$\lambda R_{i,t} e^{-\lambda t} - \int_t^{T_i} \lambda^2 R_{i,s} e^{-\lambda s} \geq c\,,$$

where the first term is obtained by L'Hôpital's rule, and the second term is obtained via bounded convergence. □

A.4 Proof of Proposition 3

It will be convenient to write $F(t) = 1 - e^{-\lambda t}$ to denote the probability that an agent succeeds by date t if he works continuously until that time, and by $f(t)$ the corresponding probability density function.

We begin by establishing two lemmas. The first shows that the egalitarian contest admits a simple symmetric pure-strategy equilibrium.

Lemma 3. *The egalitarian contest with deadline T^{EGA} has a symmetric pure-strategy equilibrium where each player works throughout the interval $[0, T^{\mathrm{EGA}}]$.*

Proof of Lemma 3. Consider any symmetric pure strategy profile in which all players work for a duration T. Then player i's expected reward conditional on succeeding is

$$R = \mathbb{E}\left[\frac{1}{1+M}\right] = \frac{1 - (1 - F(T))^n}{nF(T)},$$

where $M\sim$ $Binom(n-1, F(T))$ is the random variable equal to the number of players other than i who also succeed, and the second equality is established in footnote 11.

Now, taking as given the strategy profile of the other players, the net expected payoff of player i from spending effort for duration T is given by $F(T)R - Tc$. Note that because F is concave, this is a concave objective and therefore, the best-response for player i is the duration T' given by $f(T')R = c$; in other words

$$\frac{1-(1-F(T))^n}{nF(T)} = \frac{c}{f(T')}.$$

Finally, in a symmetric equilibrium, all players choose best-responses. Therefore, they work for a duration T^{EGA} given by

$$\frac{1-\left(1-F(T^{\text{EGA}})\right)^n}{nF(T^{\text{EGA}})} = \frac{c}{f(T^{\text{EGA}})} \Leftrightarrow \frac{1-e^{-\lambda n T^{EGA}}}{n(1-e^{-\lambda T^{EGA}})} = \frac{ce^{\lambda T^{EGA}}}{\lambda},$$

after substituting the expressions for $f(\cdot)$ and $F(\cdot)$. □

The second lemma shows that in any contest, the reward functions must satisfy a certain "budget constraint," which stems from the fact that the prize's value is \$1.

Lemma 4. *In an equilibrium of a contest in which each player i spends effort continuously through an interval $[0, T_i]$, the reward functions $R_{i,t}$ must satisfy the following "budget constraint"*

$$\sum_{i=1}^{n} \int_0^{T_i} f(t)R_{i,t}dt \leq 1 - \prod_{i=1}^{n}(1-F(T_i)). \tag{BC}$$

Proof of Lemma 4. Note that

$$\int_0^{T_i} f(t)R_{i,t}dt$$

is the expected share of the prize earned by agent i. Thus, the left-hand side of (BC) is the total expected share of the prize promised to the agents. In a feasible contest in which an agent can earn a share of the prize only if he succeeds, this total expected share cannot exceed the total probability that at least one player succeeds; i.e., the expression on the right-hand side of (BC). □

Using Lemmas 2 and 4, we consider the following relaxation of (3):

$$\max_{\{T_i\},\{R_{i,t}\}} \sum_{i=1}^{n} T_i \text{ subject to (IC) and (BC).} \tag{9}$$

In this problem, the principal chooses for each agent, a time cutoff T_i and a reward function $R_{i,t}$ such that the necessary condition for incentive compatibility (IC) and the budget constraint (BC) is satisfied.

Notice that the egalitarian contest characterized in Lemma 3 has $T_i = T^{\text{EGA}}$ and $R_{i,t} = \left[1 - \left(1 - F(T^{\text{EGA}})\right)^n\right] / \left[nF(T^{\text{EGA}})\right] = c/f(T^{\text{EGA}})$ for all i and t, and it satisfies the constraints in (9) with equality at all times.

Pick an arbitrary set of time cutoff and reward function pairs $\{T_i, R_{i,t}\}$ (one for each agent) that are feasible for (9). We will show that this solution achieves a smaller objective than the egalitarian contest characterized in Lemma 3, that is, $\sum_i T_i < nT^{\text{EGA}}$. Because the egalitarian contest is feasible for the original problem (3), it will immediately follow that this contest must be optimal.

Define the function Z_i^1 for each i as follows

$$Z_i^1(t) = \frac{1}{f(t)} \left[c - \int_t^{T_i} f'(s) R_{i,s} ds\right].$$

Because F is concave and hence $f'(s) \leq 0$, we have

$$0 \leq Z_i^1(t) \leq R_{i,t}$$

for all $t \in [0, T_i]$. The second inequality follows because $R_{i,t}$ is incentive compatible. Continuing in this manner, define for all $k \geq 2$, the function Z_i^k by

$$Z_i^k(t) = \frac{1}{f(t)} \left[c - \int_t^{T_i} f'(s) Z_i^{k-1}(s) ds\right].$$

Since F is concave and $Z_i^1(s) \leq R_{i,s}$ for all s, we have $Z_i^2(t) \leq Z_i^1(t)$. By induction we have that $0 \leq Z_i^k(t) \leq Z_i^{k-1}(t)$ for all $t \in [0, T_i]$. We have thus constructed a pointwise decreasing sequence of non-negative-valued functions on the domain $[0, T_i]$. Let Z_i be the pointwise limit. For each i we have

$$\begin{aligned} Z_i(t) = \lim_{k\to\infty} Z_i^k(t) = \lim_{k\to\infty} \frac{1}{f(t)} \left[c - \int_t^{T_i} f'(s) Z_i^{k-1}(s) ds\right] \\ = \frac{1}{f(t)} \left[c - \int_t^{T_i} f'(s) Z_i(s) ds\right] \end{aligned} \tag{10}$$

by dominated convergence.

Define a new reward function $\tilde{R}_{i,t} = Z_i(t)$. Then $\tilde{R}_{i,t}$ satisfies the incentive constraint with equality at all times:

$$f(t)\tilde{R}_{i,t} + \int_t^{T_i} f'(s)\tilde{R}_{i,s}ds - c = 0. \tag{11}$$

Differentiating both sides of (11) reveals that $\tilde{R}_{i,t}$ is the constant function $\tilde{R}_{i,t} \equiv c/f(T_i)$. This reward function satisfies the budget constraint (BC) because $0 \leq Z_i(t) \leq R_{i,t}$ for all t and $R_{i,t}$ is feasible by assumption. In particular, since the expected share of the prize earned by player i equals $\int_0^{T_i} f(t)\tilde{R}_{i,t}dt = cF(T_i)/f(T_i)$, we have

$$c\sum_{i=1}^{n} \frac{F(T_i)}{f(T_i)} - \left[1 - \prod_{i=1}^{n}(1 - F(T_i))\right] \leq 0. \tag{12}$$

Note for further reference that if any of the $R_{i,t}$ were non-constant, then the $\tilde{R}_{i,t}$ satisfy the budget constraint with a strict inequality.

We will conclude the proof by showing that the expression on the left-hand side of (12) is jointly strictly convex in $(T_1, \ldots, T_n)$. For this will imply that the following symmetric reward function profile also satisfies the budget constraint:

$$R_{i,t}\star = \frac{c}{f(\bar{T})},$$

where $\bar{T}$ is the average effort duration; i.e., $\bar{T} = \sum_i T_i/n$. Indeed the budget constraint will be satisfied with a strict inequality as long as not all the T_i were equal.

To prove that the left-hand side of (12) is strictly convex, substitute the expressions $F(T_i) = 1 - e^{-\lambda T_i}$ and $f(T_i) = \lambda e^{-\lambda T_i}$, and after some simplification and eliminating constants, the left-hand side equals

$$c\sum_{i=1}^{n} e^{\lambda T_i} + \lambda e^{-\lambda \sum_{i=1}^{n} T_i}.$$

Its Hessian, $\mathbf{H} \in \mathbb{R}^{n \times n}$, has entries

$$H_{ii} = c\lambda^2 e^{\lambda T_i} + \lambda^3 e^{-\lambda \sum_{i=1}^{n} T_i} \quad \text{for each } i \text{, and}$$
$$H_{ij} = \lambda^3 e^{-\lambda \sum_{i=1}^{n} T_i} \quad \text{for all } i \neq j.$$

For any vector $\mathbf{z} \in \mathbb{R}^n_+$, we have

$$\mathbf{z}^T \mathbf{H} \mathbf{z} = c\lambda^2 \sum_{i=1}^{n} e^{\lambda T_i} z_i^2 + \lambda^3 e^{-\lambda \sum_{i=1}^{n} T_i} \left(\sum_{i=1}^{n} z_i \right)^2 \geq 0 \,,$$

and this inequality is strict if $\mathbf{z}$ has at least one strictly positive entry, implying that the Hessian is positive semidefinite, and hence the left-hand side of (12) is strictly convex.

We have shown that the set of time cutoff and reward function pairs $\{\bar{T}, R_{i,t}\star\}$ are feasible for (9) and achieve a bigger objective than $\{T_i, R_{i,t}\}$; i.e., $n\,\bar{T} \geq \sum_i T_i$, where the inequality is strict if not all the T_i were equal. Therefore, the relaxed problem given in (9) can be rewritten as

$$\max_{T} \left\{ nT \text{ s.t. } cn \frac{F(T)}{f(T)} \leq 1 - [1 - F(T)]^n \right\}, \tag{13}$$

where we have substituted $R_{i,t} = c/f(T)$, which satisfies (IC) with equality for all $t \in [0, T]$. We will show that $T = T^{\text{EGA}}$ solves (13).

First notice that the constraint in (13) binds when $T = T^{\text{EGA}}$. Using the expressions $F(T) = 1 - e^{-\lambda T}$ and $f(T) = \lambda e^{-\lambda T}$, this constraint can be rewritten as $cn(e^{\lambda T} - 1)/\lambda \leq 1 - e^{-n\lambda T}$. We claim that this inequality is satisfied if and only if $T \leq T^{\text{EGA}}$. To see why, define $\varphi(T) = 1 - e^{-n\lambda T} - cn(e^{\lambda T} - 1)/\lambda$ and observe that

$$\varphi(0) = 0,\ \varphi'(0) = n(\lambda - c) > 0, \text{ and } \varphi \text{ is strictly concave.}$$

Therefore, $\varphi(T)$ single-crosses zero from above at $T = T^{\text{EGA}}$, and so T^{EGA} is the largest deadline for which the constraint in (13) is satisfied. Since the objective is to maximize T, $T = T^{\text{EGA}}$ solves this problem.

We have therefore shown that $T = T^{\text{EGA}}$ and $R_{i,t} = c/f(T^{\text{EGA}})$ for each i solves (9), and its objective equals nT^{EGA}. Since this is a relaxation of the original problem, (3), the objective of the original problem is bounded above by nT^{EGA}. By Lemma 3, the egalitarian contest with deadline T^{EGA} has an equilibrium in which each agent spends total effort T^{EGA}, and so the principal's objective is equal to nT^{EGA}, that is, it achieves the upper bound obtained from the solution of (9). Therefore, this egalitarian contest is an optimal no-feedback contest. $\square$

A.5 Proof of Proposition 4

The proof is organized as follows. First, we show that each agent's expected reward conditional on succeeding at date t, $R_{i,t} = c/\lambda_t$. Then we argue that there exists an equilibrium in which all agents work continuously until they succeed or the contest ends. Finally, we argue that both conditions of Lemma 1 are met, and hence this contest is optimal.

For each t, define $\overline{T}_t$ to be the smallest T which solves

$$\frac{1 - e^{-n\int_t^T \lambda_v dv}}{n\left(1 - e^{-\int_t^T \lambda_v dv}\right)} = \frac{c}{\lambda_T}. \tag{14}$$

The left-hand side is strictly decreasing in T, it converges to 1 as $T \to t$, and to $1/n$ as $T \to \infty$. Meanwhile, $\lambda_T \in (c, nc)$ by assumption, and so the right-hand side takes values strictly between $1/n$ and 1. Since both sides are continuous in T, by the intermediate value theorem, there exists a smallest T such that (14) is satisfied. Moreover, because the left-hand side of (14) is strictly larger for $T \simeq t$, this is also true for all $T < \overline{T}_t$. Therefore, for every t, we have $\gamma_s^t \geq 0$ for all $s \in [0, \overline{T}_t)$, and $\lim_{s \to \overline{T}_t} \gamma_s^t = \infty$; i.e., a cycle which starts at t ends with certainty by $\overline{T}_t$.

Consider a cycle that started at t. Fix a date $s > t$, and suppose that agent i has worked continuously until this date. Then his expected reward conditional on succeeding at s is

$$R_{i,s} = \int_s^{\overline{T}_t} \frac{\left(1 - e^{-n\int_t^z \lambda_v dv}\right)}{n\left(1 - e^{-\int_t^z \lambda_v dv}\right)} \gamma_z^t e^{-\int_s^z \gamma_r^t dr} dz + \frac{\left(1 - e^{-n\int_t^{\overline{T}_t} \lambda_v dv}\right)}{n\left(1 - e^{-\int_t^{\overline{T}_t} \lambda_v dv}\right)} e^{-\int_s^{\overline{T}_t} \gamma_r^t dr}.$$

To interpret this expression, suppose this agent succeeds at s. During every interval $(z, z + dz) \subseteq (s, \overline{T}_t)$, the current cycle ends with probability $\gamma_z^t e^{-\int_s^z \gamma_r^t dr} dz$. In this case, his expected share of the prize is

$$\mathbb{E}\left[\frac{1}{1 + M_{t,z}}\right] = \frac{1 - e^{-n\int_t^z \lambda_v dv}}{n\left(1 - e^{-\int_t^z \lambda_v dv}\right)},$$

where $M_{t,z} \sim Binom(n-1, 1-e^{-\int_t^z \lambda_v dv})$ represents the number of rivals who succeed between the date that the current cycle started and z, and we have assumed that in equilibrium, they work continuously until they succeed. Integrating over the interval $(s, \overline{T}_t)$ yields the first term of $R_{i,s}$. With probability $e^{-\int_s^{\overline{T}_t} \gamma_r^t dr}$, the cycle survives until (and ends at) $\overline{T}_t$, in which case agent i's expected share of the prize is $[1 - e^{-n\int_t^{\overline{T}_t} \lambda_v dv}]/[n(1 - e^{-\int_t^{\overline{T}_t} \lambda_v dv})]$.

We now show that $\lambda_s R_{i,s}$ is constant and equal to c. By the definition of $\overline{T}_t$, we have $\lambda_{\overline{T}_t} R_{i,\overline{T}_t} = c$; i.e., the desired equality is satisfied for $s = \overline{T}_t$. By differentiating $\lambda_s R_{i,s}$ with respect to s, we have

$$\frac{d}{ds}\lambda_s R_{i,s} = \dot{\lambda}_s R_{i,s} + \lambda_s \gamma_s^t \left[R_{i,s} - \frac{\left(1 - e^{-n\int_t^s \lambda_v dv}\right)}{n\left(1 - e^{-\int_t^s \lambda_v dv}\right)} \right] = 0$$

whenever $R_{i,s} = c/\lambda_s$. The first equality follows from the Leibniz integral rule, and the second equality follows by substituting γ_s^t defined in (4) and $R_{i,s} = c/\lambda_s$. Therefore, $\lambda_s R_{i,s} = c$ at $s = \overline{T}_t$, and moving backwards in time, $d\left(\lambda_s R_{i,s}\right)/ds = 0$, implying that $\lambda_s R_{i,s} = c$ for all $s \in [t, \overline{T}_t]$.

Because the hazard rate of F is increasing, if an agent has worked continuously until date t, then his hazard rate will be equal to λ_t; otherwise, it will be smaller. So an agent who has worked continuously until t without success, taking as given that his rivals work until they succeed, weakly prefers to work at t. By symmetry, it follows that there exists an equilibrium in which all agents work continuously until they succeed, meanwhile earning zero rents.

Finally, because the contest does not end until at least one agent succeeds, the prize is awarded with probability one, which implies that this contest satisfies both conditions of Lemma 1, and is therefore optimal. $\square$