Optimal Feedback in Contests*

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

September 29, 2022

Abstract

We obtain optimal dynamic contests for environments where the designer monitors effort through coarse, binary signals—Poisson successes—and aims to elicit maximum effort, ideally in the least amount of time possible, given a fixed prize. The designer has a vast set of contests to choose from, featuring termination and prize allocation rules together with real-time feedback for the contestants. Every effort-maximizing contest (which also maximizes total expected successes) has a history-dependent termination rule, a feedback policy that keeps agents fully apprised of their own success, and a prize allocation rule that grants them, in expectation, a time-invariant share of the prize if they succeed. Any contest that achieves this effort in the shortest possible time must in addition be what we call second chance: once a pre-specified number of successes arrive, the contest enters a countdown phase where contestants are given one last chance to succeed.

*We are grateful to our editor Bard Harstad, four anonymous referees, Daniel Barron, Piotr Dworczak, Willie Fuchs, Marina Halac, Navin Kartik, Igor Letina, Alessandro Pavan, Michael Powell, Agustiń Rayo, Andy Skrzypacz 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 Dayton, skhorasani1@udayton.edu; L. Rayo: Northwestern University, luis.rayo@kellogg.northwestern.edu.

1 Introduction

Contests—situations where multiple agents compete for a prize—are a common way to organize economic activity. 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 best to allocate the prize among participants, and more recently, starting with the work of Yildirim (2005), Ederer (2010), and Halac, Kartik and Liu (2017), how best to disclose information over time.

It has nevertheless proven challenging to find fully optimal contests for dynamic environments (such as innovation races, promotion tournaments, and qualification stages for athletic events) where contestants have an opportunity to work continuously for some length of time and where information about their efforts may be revealed in real-time as they do so. In such cases, characterizing an optimal contest involves finding an ideal triple—a prize allocation rule, a termination rule, and a real-time feedback policy—among a dauntingly large set.

Here we study dynamic scenarios where finding such optimal contests is possible. These scenarios involve each contestant exerting all-or-nothing effort and producing (independent) "Poisson successes" over time. As it turns out, the contest designer has multiple ways to fully convert their prize money into effort (which guarantees in expectation maximum effort, or equivalently, maximum number of successes) with these having in common a history-dependent contest deadline, a feedback policy that keeps each agent fully apprised of their own success and a prize allocation rule that grants them, in expectation, a time-invariant share of the prize if they succeed.

One such contest, which we term "cyclical egalitarian," features a cyclical structure whereby the contest is terminated at the end of each fixed-length cycle if at least one agent has succeeded by then, and it is otherwise reset. The prize is shared equally among all successful agents irrespective of *when* they succeeded, and the feedback policy keeps agents fully apprised of their own success, but only periodically informs them about their rivals' successes—at the end of each cycle—so as to not discourage further effort.

Only a much smaller set of contests, however, is capable of converting all prize money into effort *in the shortest expected length of time*—a property that would be valuable to the designer if for instance running the contest entailed a flow cost. These tournaments, which we term "2^{nd} chance," all have in common that the contest continues until some number K of successes arrive and, once that occurs, the contest enters a countdown phase where contestants are given a final (potentially random) deadline to succeed (a "second chance"), with the contest ending before that deadline if one more contestant succeeds. These contests minimize length because they guarantee that the number of agents working at a given time is as similar as possible across different histories—which in turn prevents inefficient scenarios where the contest continues with only a small number of agents are still working.

In our model a "contest," defined here as agents competing for the same prize, improves upon individual contracting, with reserved prize money for each agent, along two dimensions. First, it allows the designer to spend her budget exactly while surrendering zero rents—a feat that, save for knife-edge cases, would be impossible without having agents compete for the same prize. In addition, by inducing more agents to work at the same time, a contest shortens, potentially by a large margin, the expected time needed to extract the desired effort from the contestants.

In our baseline model, contestants are all identical and each can succeed at most once. We interpret these successes as something not necessarily of intrinsic value to the designer, but rather as a coarse measure of effort. It is perhaps surprising that despite this coarseness, the designer can attain the maximum effort possible given the prize budget. We also consider several extensions that suggest a degree of robustness to our findings: the possibility of multiple successes, heterogeneity of success rates across agents, and an increasing hazard rate that captures a notion of progress or knowledge accumulation over time. In all these cases, a 2^{nd} chance contest attains maximum effort; and in several cases, it does so in the minimum possible time. We also relax the designer's commitment power and show that an egalitarian prize structure, which is implicitly present in all our other designs, remains optimal.

Applications. Our model, while stylized, is inspired by settings where agents' efforts can only be imperfectly monitored through coarse performance measures, and where these measures are (at least approximately) discrete. A first example is an innovation contest where agents attempt a specific technological breakthrough, as defined in advance by the designer, such as inventing a device that performs a particular task. These contests have become increasingly common across a variety of sectors.[1] One

[1] Examples can be found in private crowdsourcing platforms such as Innocentive or Kaggle, federal agencies including the DoD and NASA, Challenge.Gov competitions, and in philanthropic organi-

recent example is the Netflix Prize competition that sought a 10% improvement in the prediction accuracy of one its algorithms. Similar to our *Netflix-style* design, this contest included a final 30-day countdown phase, initiated once the first success was reached, after which a winner would be declared. In these innovation settings, like in our model, the designer may care about more than a single success, as every success (and even effort absent a success) may produce new ideas.[2]

A second example is an organization, such as a professional partnership or an academic department, where large successes (such as landing a new client or publishing a home-run paper) may be rare. One could also interpret a discrete "success" as reaching a sufficiently-well-defined threshold for promotion. Here the splitting of a fixed prize—a promotion—can be interpreted as a probabilistic allocation of the full prize (e.g., on the basis of random differences in the magnitudes of different successes).

Finally, elite athletic events (for instance in track & field, cycling, or rowing) are commonly preceded by qualifying stages where athletes have a period of time to meet a performance threshold. Akin to our 2^{nd} chance contests, these qualifying stages frequently begin with a phase where athletes have an opportunity to meet or exceed a pre-specified goal, followed by a *repechage* phase where some of the athletes who initially fell short are given a final chance.

In all these scenarios, agents may in principle succeed more that once and, moreover, the principal may have the ability to refine their performance measure to more closely monitor effort. As we shall see, our lessons apply similarly to the case of multiple successes and because the principal is able to extract the maximum possible effort even with the coarse measure, the benefit from refining it may be limited.

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. Drugov and Ryvkin (2019, 2020) and Olszewski and Siegel (2020), among others, consider extensions to stochastic output, arbitrary risk-preferences and heterogeneous agents.[3]

zations like XPrize.

[2]In fact, to save on engineering effort, Netflix did not implement the prize-winning algorithm and instead opted for a combination of two lesser-performing ones; see tinyurl.com/37kdtz74.

[3]Siegel (2009, 2010) and Olszewski and Siegel (2016) provide a comprehensive equilibrium analysis

Fang, Noe and Strack (2018) show that individual 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, and find that for n contestants, a nested Tullock contest featuring $n-1$ equal prizes is optimal. This potential desirability of "turning down the heat" extends to dynamic settings: Moscarini and Smith (2011) and Ryvkin (2020) show that incentives are strongest when agents are tied and quickly collapse once one is ahead. While our work differs in that it features feedback design and more flexible termination and prize-allocation rules, the optimality of an egalitarian prize in our framework echoes their idea that reducing competitiveness motivates greater effort.

Taylor (1995) considers a dynamic contest of exogenous length where, upon paying an entry fee, agents invest in an innovation of stochastic quality. In equilibrium, 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 to allow for interim transfers and an endogenous termination date. There the optimal contest ends as soon as the highest-quality innovation exceeds a threshold, and agents invest until that time. Both models, unlike our own, restrict to a winner-takes-all prize structure, exclude feedback, and allow the principal to extract rents via entry fees. Because we rule out entry fees (e.g., due to cash constraints), the termination and prize allocation rules combined with the feedback policy play a crucial role in extracting rents.

Among the first to study endogenous feedback in contests are Lizzeri, Meyer and Persico (2005) and Yildirim (2005), who use a two-period, two-agent framework. In that same 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. Also related is Khorasani, Körpeoğlu and Krishnan (2021), who show that in a two-stage winner-takes-all contest, dynamically-adjusted rewards and probabilistic information disclosure can improve upon a fixed-reward contest.

Finally, 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

of general all-pay contests with heterogeneous players. See also Georgiadis (2022) for a review.

(2011), but with a principal who designs a contest to maximize the probability of a single success. Among rank-monotonic prize schemes (awarding a weakly large prize to agents who succeed earlier) and deterministic and symmetric disclosure policies, the optimal design provides no interim feedback and ends as soon as a critical number of agents have succeeded or a deadline is reached, with each successful agent winning an equal share of the prize. This result is driven by the need to balance intertemporal incentives in light of agents learning about the feasibility of the project over time, which favors a commitment to silence. What distinguishes our setting is that the principal maximizes total effort or total number of successes (which means a second chance is desirable), there is no learning (i.e., successes arrive with a constant or increasing hazard rate), and the principal is unconstrained in her choice of contest—indeed, the optimal contest in our model need not be rank-monotonic.[4] In Bimpikis, Ehsani and Mostagir (2019), an agent must succeed twice to win, with the feasibility of the first success unknown. Under certain conditions, 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, every agent 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 a promotion application, it may represent landing a client or exceeding an exogenous bar for promotion; in an innovation contest or athletic application, it may represent achieving a pre-specified target. If agent i spends effort $a_{i,t} \in \{0, 1\}$,

[4] In Halac, Kartik and Liu (2017) the principal extends the contest beyond the first success because otherwise agents would become too pessimistic over time if the contest has not yet ended; in our model she does so, in contrast, to extract additional effort/successes.

he incurs cost at rate $ca_{i,t}$, where c represents the (constant) marginal cost of effort, and conditional on not having succeeded before, achieves a success stochastically with constant instantaneous rate $\lambda a_{i,t}$ where $\lambda \in (c, nc)$; thus, agents cannot succeed while they shirk and if K agents are working, the expected time until the next success is $1/(\lambda K)$.[5] In our baseline model agents are identical and each can succeed at most once.

The principal observes successes but not efforts. Each agent 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 privately informs each agent immediately if he succeeds. 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 and messages. 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 agent 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 agent i to have succeeded, whereas an *egalitarian* contest divides the prize equally among all agents who have succeeded.

If the contest ends at time s, agent i's ex-post payoff is

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

There is no discounting and agents 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

[5]The assumption that $\lambda < nc$ implies that there are enough competitors for a contest to be desirable in the first place. When this condition fails, the principal could do as well by reserving $1/n$-th of the prize for each agent and contracting with each one individually. If $\lambda < c$, it is impossible to incentivize even a single agent to spend effort.

(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 agents' efforts even after they have succeeded. Our results would be unchanged, though, if the principal instead sought to maximize the total number of successes.[6] We shall restrict, without loss, to contests where in equilibrium each agent works continuously over some (history- and agent-dependent) time interval: agents do not pause and restart.

3 A sufficiency result

Finding an 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 history 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 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 all players' combined surplus; if the agents keep none of this surplus, it must all go to the principal. Indeed, any such contest is payoff-equivalent to the first-best outcome in which the principal chooses the agents' efforts directly subject only to the constraint that they earn non-negative payoffs.

To formally establish this result, note that for any contest and equilibrium effort profile, we can write 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

[6] Our contests 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 that case too.

one, whereas the agents' rents are bounded from below by zero; therefore, if a contest 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 exist various contests that satisfy it.

4 Effort-maximizing contests

Here we present three examples of contests meeting our sufficiency condition. Because these contests transform (in expectation) all prize money into effort, they implement total (expected) effort $1/c$. We also establish necessary conditions that every effort-maximizing contest must satisfy.

Since agents are risk neutral, we can (and henceforth will) restrict attention without loss of generality 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 t. These functions will allow us to analyze the incentives faced by each agent separately.

When describing the feedback policy of a given contest, we follow the convention that unless otherwise noted, the principal remains silent.

4.1 Cyclical-egalitarian contest

Our first example of an effort-maximizing contest, which we call **cyclical egalitarian**, has the following features. First, its prize allocation is egalitarian: it divides the prize equally among all agents who have succeeded regardless of when they happened to do so. Second, agents are kept fully apprised of their own success 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 restarts the contest, again with a provisional deadline T^* (thus informing all agents that no one has yet succeeded). The contest continues in this manner until at least one agent has succeeded by the time the next provisional deadline is reached.[7]

[7]While a cycle is ongoing, agents should not receive any feedback about their rivals' successes.

In order to extract all rents from the agents, the provisional deadline is set just long enough that 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. *The cyclical-egalitarian contest is effort-maximizing.*

This contest achieves maximum expected effort $1/c$ 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 ensure that an agent's expected reward $R_{i,t}$ is always equal to c/λ.

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 where all their continuation payoffs are zero.[8] 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

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

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^*.[9]

[8]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 infinitesimally shrinking the cycle length. This modified contest admits only the equilibrium in which all agents work until they succeed, while giving up only arbitrarily small rents.

[9]It can be shown that T^* and the expected duration of this contest, $T^*/(1 - e^{-n\lambda T^*})$, decrease in n and in c, and increase in λ.

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 his flow payoff, and hence his continuation payoff, are always held at zero regardless of his effort, and so working until he succeeds is incentive compatible. Because agents are symmetric, an equilibrium with the desired properties exists. *Q.E.D.*

A practical advantage of this cyclical-egalitarian design is that it relies on a single parameter: the cycle length. It is also a member of a larger family of cyclical-egalitarian contests that differ only in that parameter. Provided the cycle length does not exceed T^*, these contests induce all agents to work and transform 100% of the prize into a combination of effort and agent rents, and hence are on the Pareto frontier. Indeed, by varying the cycle-length, one can trace the entire Pareto frontier for contests that treat agents symmetrically and give prize money only to agents who succeed. 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 0.[10]

4.2 Beeps contest

Our second example, which we call the **beeps contest**, is inspired by the single-agent feedback policy in Ely (2017). The contest continues until at least T^* (the length of one cycle previously), and at each $t \geq T^*$, provided at least one agent has succeeded by then, the contest ends randomly: at time T^* with probability

$$\frac{c/\lambda - e^{-\lambda(n-1)T^*}}{(1 - e^{-\lambda(n-1)T^*})(c/\lambda)},$$

and during each interval dt after that with probability $(n-1)\lambda dt/(1 - c/\lambda)$. If one or more agents succeed before T^*, the prize is shared equally among those agents, regardless of when the contest happens to end. Otherwise, the entire prize is awarded to the first agent to succeed, also regardless of when the contest happens to end. The feedback policy is $\mathcal{M}^{pronto}$.

Proposition 2. *The beeps contest is effort-maximizing.*

[10]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, they are guaranteed to win the entire prize as soon as they succeed.

This more complex design also meets the sufficiency conditions in Lemma 1. It awards the entire prize by construction and extracts all rents by ensuring that each successful agent obtains expected reward c/λ regardless of when he succeeds. Indeed, at every $t > T^*$ each agent assigns probability c/λ to the event that none of his rivals have yet succeeded; in such event, his reward for succeeding is the entire prize, which together with the $\mathcal{M}^{pronto}$ feedback policy implies that the agent is just willing to work. Moreover, because agents' continuation payoffs at T^* are zero, and any agents who succeed by that time split the prize equally (irrespective of when the contest ends), the first phase of the contest is economically identical to a single cycle of the cyclical-egalitarian design.

4.3 Netflix-style contest

Our third example, which we call the **Netflix-style contest**, has two phases: the first one runs from time 0 until the first success; the second one—a countdown phase—runs for an additional T^c units of time. Letting m denote the number of agents who succeed in phase 2, the prize is split among successful agents such that the first agent to succeed receives $\alpha/(\alpha + m)$ and all others receive $1/(\alpha + m)$ each. The feedback policy is $\mathcal{M}^{pronto}$.[11]

Proposition 3. *There exist α and T^c such that the Netflix-style contest is effort-maximizing.*

This contest is a modified version of the winner-takes-all design where rival agents are allowed a grace period to catch up—an addition that ensures the first agent is not over-rewarded and any extra prize money is instead spent on prolonging the effort of rivals. A feature of this design is that $\alpha < 1$, which means that when any rival succeeds during the grace period, the first agent who succeeded is awarded the smaller share of the prize. Such feature is needed so agents are willing to work during the grace period despite at least one of their rivals having succeeded by then.

The Netflix-style contest is also an example of a broader family of contests with countdown phases that begin as soon as a pre-specified number of agents have succeeded. As we shall see in Section 5, this type of design will be of particular interest

[11] It does not matter whether agents are told what phase they are in, but to maintain their effort incentives during phase 2, they should not receive any feedback about their rivals' successes.

to us as it will allow the principal, by means of an appropriate choice of parameters, to implement the maximal effort $1/c$ in the shortest possible time.

4.4 Necessary conditions

All contests we have presented so far use $\mathcal{M}^{pronto}$, provide constant expected rewards c/λ, and rely on a history-dependent termination rule. The following proposition shows that these features are necessary for a contest to implement effort $1/c$.

Proposition 4. *Every effort-maximizing contest features:*

i. A history-dependent termination rule such that the contest does not end until at least one agent succeeds.

ii. A prize allocation rule such that the prize is awarded with probability 1 and all active agents are promised a time-invariant reward $R_{i,t} = c/\lambda$.

iii. $\mathcal{M}^{pronto}$ feedback.

To achieve effort $1/c$, a contest must convert all prize money into effort. Immediately informing agents of their own success—despite the principal preferring that they keep working—is needed to keep them from earning rents. Indeed, an 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 goes up; but this would allow him to secure rents by initially withdrawing effort and working only once the expected reward has grown. Given the $\mathcal{M}^{pronto}$ policy, an agent who has not yet succeeded is willing to work so long as the expected reward is at least c/λ, and any reward greater than that would relinquish rents.[12] Finally, the contest must allow at least one agent to succeed so that the prize is always awarded.

One may interpret a "contest" as any design where an agent's expected reward sometimes falls (and never grows) when a rival succeeds. This is in contrast to individual contracting, where agents do not impose such negative externalities on each other. The present advantage of a contest—despite no aggregate productivity shocks or agent risk-aversion—is two-fold. First, it allows the principal to pool the agents' incentive constraints; that is, prize money not awarded to one agent can be awarded to another. This allows the principal to spend 100% of the prize money

[12] Because expected rewards are time-invariant, even if they could, agents would have no incentive to initially hide a success and report it at a later date.

without granting any rents. With individual contracting, in contrast, the principal would need to reserve c/λ prize money for each agent (more than that would grant rents), and hence she would only be able to reward $\lfloor \lambda/c \rfloor$ agents and spend $\lfloor \lambda/c \rfloor c/\lambda$ of the prize, which generically is less than 1.[13] Second, as we shall discuss shortly, a contest allows the principal to save time.

We conclude with some remarks on robustness:

i. Owing to its use of $\mathcal{M}^{pronto}$, any effort-maximizing contest would remain optimal if agents were able to observe their own successes directly or even probabilistically. For some contests, however, agents should not observe the successes of their rivals; thus, the principal has to make sure that when informing a successful agent, this communication occurs only informally (e.g., verbally) or using encryption so it cannot be credibly re-transmitted to other agents in a bid to discourage them.

ii. Every contest that maximizes effort also maximizes expected successes, as it never motivates already successful agents to keep working. The expected number of successes equals the success-per-cost ratio λ/c. As this ratio grows, each success becomes easier and so the designer needs to pay less for it.

iii. The principal would be unable to achieve higher effort with a more precise monitoring technology as, despite her imperfect (binary) signal, she is able to convert the entire prize into effort. However, to extract all rents, the monitoring technology must have the property that an agent cannot win any share of the prize without exerting effort.

iv. If the principal could raise her prize budget and every unit of effort was worth a dollar, she would do so until the budget reaches nc/λ (since each dollar of budget is transformed into $1/c > 1$ units of effort), at which point individual contracting becomes optimal and effort cannot be raised any further. By the same token, the principal would not want to lower the prize money below her initial (limited) budget—and provided that budget is strictly below nc/λ, a contest is optimal.

[13] 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: she could for instance promise each agent c/λ for a success and ask them all to work for a sufficiently short period of time so that she spends \$1 on average.

5 Maximum effort in the shortest possible time

While there is a variety of contests that implement effort $1/c$, those contests may differ in their expected duration. This occurs, in a nutshell, because they may differ in how many agents are working at any given moment, and the fewer agents are working on average, the longer the contest is expected to last.

Here we find contests that achieve the maximum effort in the shortest possible time. One justification for this lexicographic objective is that the principal bears a constant flow cost while the contest is in progress. Provided this cost is not too large, the principal will wish to minimize duration without sacrificing any effort in the process.

Obtaining such contests is in principle a challenging task because even computing the duration of specific contests can be hard. We are able to overcome this difficulty by finding a specific family of contests, which we call **2^{nd} chance**, that (uniquely) attains the theoretically shortest-possible duration given the target effort $1/c$.

Definition 1. *We say a contest is 2^{nd} chance if for some $K \in \{1, ..., n\}$:*

i. The contest continues until at least K agents succeed.

ii. Once K agents succeed, the contest enters a countdown phase where the remaining agents are given a (potentially random) deadline to succeed.

iii. The contest ends as soon as one more agent succeeds or the deadline elapses, whichever occurs first.

The following proposition shows that all minimum-duration contests must be 2^{nd} chance contests with a specific value of K and a deadline that meets a stringent condition. Such contests must of course also meet the necessary conditions of Proposition 4; that is $\mathcal{M}^{pronto}$ feedback and expected rewards equal to c/λ.

Proposition 5. *A contest achieves expected effort $1/c$ in the shortest expected time possible if and only if it is 2^{nd} chance. Any such contest must have $K = K^* := \lfloor \lambda/c \rfloor$ and a (potentially random) deadline such that the ex-ante probability of $K+1$ successes is $\lambda/c - K^*$.*

Intuitively, the way to minimize expected duration is to find a termination rule that maximizes the expected number of agents that work per unit of time. A 2^{nd} chance contest achieves this goal by ensuring that the number of successes at the moment of termination is as similar as possible across different histories so as to avoid histories where only a small number of agents work for some length of time. Moreover, by setting a specific K and an appropriate deadline, the contest ensures that agents keep zero rents.

A simple heuristic derivation provides further insight. For each j, let T_j denote the (expected) amount of time during which exactly j agents are working. To obtain a lower bound on a contest's (expected) duration, imagine that we can directly select the values of T_j without worrying about whether such values can actually be implemented by a contest. That is, solve

$$\min_{T_1,\ldots,T_n} \sum_{j=1}^{n} T_j \quad \text{s.t.} \quad \sum_{j=1}^{n} jT_j = \frac{1}{c} \quad \text{and} \quad 0 \leq T_j \leq \frac{1}{\lambda j},$$

where the first constraint ensures that the contest implements effort $1/c$ (each term jT_j represents the total expected effort obtained during the time j agents are working) and the second one follows from the fact that when j agents are working, the next success occurs with Poisson rate λj, and hence the expected time that j agents work cannot exceed $1/(\lambda j)$.

Because the greater the j the greater the weight on T_j in the first constraint, the unique solution features a cutoff (specifically, $n - K^*$), such that T_j takes its upper bound for all j greater than the cutoff, its lower bound for all j smaller than the cutoff, and its value *at* the cutoff is chosen—zero or interior—such that total effort is $1/c$. It follows that the ideal contest must always continue whenever less than K^* agents have succeeded and must never do so after $K^* + 1$ agents have succeeded. This is precisely what a 2^{nd} chance contest achieves. The expected duration of the countdown phase—and hence the probability of success during that phase—is pinned down by the requirement that agents receive zero rents (and hence total effort is $1/c$).

This heuristic also clarifies why other contests fail to minimize duration. Take, for instance, the cyclical-egalitarian design that resets every T^* units of time if no one has succeeded. Because there are histories where the contest continues even when a large fraction of agents have already succeeded—and therefore only a small number of agents are still working—there must also exist histories where the contest ends after only a small number of agents have succeeded (per the constraint on total effort), which is an inefficient way to allocate the T_j's. The Beeps and Netflix-style designs have this drawback as well.

A specific type of contest that meets all requirements in Proposition 5, which we call **egalitarian** 2^{nd} **chance**, is a 2^{nd} chance contest with $K = K^*$, $\mathcal{M}^{pronto}$ feedback, and the following additional features:

i. The countdown phase has a deterministic deadline

$$T^{sc} = \frac{\log(1 + K^* - \lambda/c)}{-\lambda(n - K^*)}.$$

ii. If an agent succeeds during the countdown, he earns c/λ and the original K^* successful agents split the remaining prize equally; otherwise, those K^* agents split the entire prize equally.

Different egalitarian 2^{nd} chance contests differ only in their feedback about rivals' successes. Conveniently for the designer, this feedback is of no consequence.

Corollary 1. *Every egalitarian* 2^{nd} *chance contest implements expected effort* $1/c$*, and among contests that do so, has minimum expected duration.*

These contests extract all rents because the length of the countdown together with the value of K ensure that the first K successful agents receive in expectation c/λ, and any one who succeeds during the countdown receives the same reward. That they achieve minimum expected duration follows from the fact that they are 2^{nd} chance.[14]

We conclude with some remarks. First, there exist other 2^{nd} chance contests with minimum expected duration. These differ from an egalitarian 2^{nd} chance contest, at most, in that the deadline for the countdown phase and the prize allocation could each be random. However, the ex-ante probability of a success during the countdown phase (together with the expected duration of this phase) and the expected prize for each successful agent must be the same.

Second, in any optimal 2^{nd} chance contest, if an agent succeeds during the countdown, that agent will take prize money away from the first K^* successful agents (a negative externality) and will by construction earn more than them despite succeeding later. Thus, the prize structure is not "rank monotonic" in the Halac, Kartik and Liu (2017) sense.

[14]Because 2^{nd} chance contests limit the total number of successes (and hence the number agents who can share the prize) they are robust to contestants learning about their rivals' outcomes.

Third, because K^* is the maximum number of agents that can be motivated using individual contracts by reserving c/λ prize money for each, every optimal contest outperforms that form of contracting both in expected effort (as it uses the entire budget) and in expected duration per unit of effort (as more agents work at the same time).

Fourth, the expected contest duration falls with n because more agents work to generate the same expected effort, falls with c because total effort falls, and grows with λ because even though each success occurs quicker, the designer needs to pay less for each success and hence waits for more of them to occur.[15]

Finally, we have assumed that the designer has a lexicographic objective whereby she first maximizes effort and then minimizes contest length. If her objective was instead to maximize expected effort minus a flow cost γ times length, then as long as $\gamma \leq n - K^*$, the same 2^{nd} chance contests as before would be optimal; otherwise, every optimal contest would end as soon as $n - \lfloor \gamma \rfloor$ agents succeed.[16]

6 Extensions

Here we consider four extensions that suggest a degree of robustness to our findings. The first three involve, respectively, multiple successes, agent heterogeneity, and a hazard rate of success that grows in an agent's past effort. With multiple successes, a 2^{nd} chance contest is fully optimal. With heterogeneous agents, a generalized version of such contest maximizes effort but may fail to minimize duration—although with either two agents or sufficiently similar ones it minimizes duration as well. Similarly, with a growing hazard rate, a generalized 2^{nd} chance contest maximizes effort—and minimizes duration provided the hazard rate does not grow too quickly over time. We then reduce the commitment power of the principal and show that an egalitarian prize allocation is optimal, albeit one embedded in a simpler contest with a fixed deadline.

[15]Per the proof of Proposition 5, every minimum-duration contest with expected effort $1/c$ lasts in expectation $\underline{D} = (\lambda/c - K^*)/(\lambda n - \lambda K^*) + \sum_{k=n+1-K^*}^{n} (\lambda k)^{-1}$. The desired comparative statics follow from the facts that $\underline{D}$ decreases with c and n, and increases with λ and K^*, which itself increases with λ and decreases with c.

[16]Since the designer's flow payoff equals the number of agents still working minus γ, she wishes to end the contest immediately whenever $\lfloor \gamma \rfloor$ or fewer agents still work. Because agents inevitably earn rents in this case, the split of the prize need not be exactly egalitarian, although an equal split is guaranteed to be optimal.

6.1 Multiple successes

Recall that even when each agent can succeed only once, the designer is able to transform all her prize money into effort. The upshot is that if agents were able to succeed multiple times, the designer would be unable to raise effort further. Multiple successes, however, may help her reduce contest length as agents need not stop working after their first success.

Proposition 6. *Suppose agents can succeed multiple times, with each success arriving at rate λ provided an agent is working. The following version of a 2^{nd} chance contest implements effort $1/c$ in the shortest possible time:*

i. The contest continues until K^* successes occur. At that moment, it ends with probability $1-\lambda/c+K^*$; otherwise, it continues until one more success arrives.[17]

ii. If the contest ends after the first K^* successes, the entire prize is split evenly across the first K^* successes; otherwise, the K^*+1^{st} success is awarded c/λ and the remaining prize is spit evenly across the first K^* successes.

This design is very similar to a 2^{nd} chance contest for the case of a single success. The key differences are that all agents work throughout, which ensures minimum contest length, and the prize is now divided on the basis of individual successes rather than successful agents. The probability with which the contest continues after the first K^* successes is chosen so that each success receives in expectation c/λ.[18]

Observe that the countdown phase is either indefinite or non-existent. This guarantees a fixed number of successes in each phase, and therefore ensures that agents have no fear of diluting their own prize when continuing to work after succeeding. This feature was not needed when agents could succeed only once.

Another case of potential practical interest is where each agent can succeed multiple times, but his cost of effort grows after every success (e.g., because the agent works on the most enjoyable opportunities first). In this case, to maximize effort, the designer must focus on each agent's first success alone as these are the cheapest to incentivize; hence, provided the primary objective is total effort (or equivalently total successes) the model is equivalent to the baseline model where by assumption there is a single success.

[17] The feedback policy is immaterial.

[18] This contest with the addition of the $\mathcal{M}^{pronto}$ feedback policy is also optimal for the baseline model with a single success.

6.2 Heterogeneous agents

Here we return to the baseline case of a single success per agent but allow agents to vary in their success hazard rates, which we assume are common knowledge. We show that a generalized version of a 2^{nd} chance contest that rewards successful agents in inverse proportion to their hazard rates is effort maximizing. This contest is also duration minimizing when agents are sufficiently similar (or when there are two agents only).

Proposition 7. *Suppose agent i succeeds with rate $\lambda_i > c$ such that $\sum_{i=1}^{n} c/\lambda_i > 1$, and let $m_i := c/\lambda_i$ denote the minimum prize required for that agent to work. Any contest with $\mathcal{M}^{pronto}$ feedback and the following termination and prize allocation rules implements maximum effort:*

i. If the sum of the m_i's across all successful agents by time t, denoted M_t, weakly exceeds 1, the contest ends instantly. The last agent to succeed, say j, is awarded m_j, and every other successful agent i is awarded $(1 - m_j) \times m_i/(M_t - m_j)$.

ii. At any t at which a success occurs and $M_t < 1 \leq M_t + m_j$ for some yet-unsuccessful agent j, the contest ends with some positive probability (defined in the proof) and if so, each successful agent i is awarded prize m_i/M_t.

iii. Otherwise, the contest continues until at least the next success.

In the special case where all agents are identical, this contest reduces to a 2^{nd} chance one where the countdown phase has either a zero or infinite deadline. The more complex design is needed so that successful agents each receive prize m_i in expectation regardless of their individual success rates. Akin to the simpler 2^{nd} chance contest, it achieves this by sometimes over-rewarding successful agents, which occurs when the contest ends with the principal "under budget" (i.e., $M_t < 1$), and sometimes under-rewarding them, which occurs when the principal is instead "over budget." Feedback about rivals is immaterial.[19]

[19] As in the baseline model, maximizing total effort is equivalent to maximizing total number of successes because the rate at which prize money is transformed into effort remains equal across agents (i.e., $1/c$). The case of heterogeneous costs is more challenging (and the two objectives are no longer equivalent) because the rate now differs across agents; hence Lemma 1 no longer applies. However, if costs are similar across agents, a 2^{nd} chance contest with agent-specific prizes remains approximately optimal as even though it grants rents and over-spends on motivating less efficient agents, these losses are small.

This contest need not be duration minimizing because the heterogeneity across m_i's means that the number of successful agents upon termination may vary significantly across different histories. However, if there are either two agents only or agents are sufficiently similar, the contest will have minimum length. The reason is that in such cases there exists a K such that the principal is under budget after any K agents succeed and over budget after any $K+1$ of them do so, and so the number of successful agents upon termination is either K or $K+1$, regardless of the history, which is the smallest variation possible.

6.3 Increasing hazard rate

In some settings, an agent's probability of success may grow over time as he accumulates knowledge or otherwise progresses towards a solution. For example, contestants might be sampling among a finite set of possible solutions or may need to accumulate a number of intermediate Poisson successes before they solve the overall problem (as for example in Doraszelski, 2003). In this case, a generalized 2^{nd} chance design similar to that of Proposition 7 (for heterogeneous agents) maximizes effort. As before, the designer must adjust prizes on the basis of the agents' hazard rates, but rather than this adjustment varying across agents, it must vary over time as the hazard rate grows.[20]

Formally, we assume that the success rate for an agent who has worked for t units of time, denoted λ_t, is weakly increasing. Let $m_i := c/\lambda_{\tau(i)}$ where $\tau(i)$ denotes the success time of agent i, and let M_t denote the sum of the m_i's across all successful agents by time t.

Proposition 8. *Assume $\lambda_t \in (c, nc)$ is weakly increasing. Any contest with $\mathcal{M}^{pronto}$ feedback and the following termination and prize allocation rules implements maximum effort:*

i. If $M_t \geq 1$, the contest ends instantly. The last agent to succeed, say j, is awarded m_j, and every other successful agent i is awarded $(1-m_j) \times m_i/(M_t - m_j)$.

ii. At any t at which a success occurs and $M_t \in (1 - c/\lambda_t, 1)$, the contest ends instantly with some probability (defined in the proof), and each successful agent i is awarded m_i/M_t.

[20]When the hazard rate is decreasing, finding an optimal contest is challenging because agents inevitably earn rents.

iii. Otherwise, the contest continues until at least the next success.

Akin to the case of heterogeneous agents, this design minimizes length so long as the hazard rate does not grow too quickly over time. This ensures that there is a K such that the principal is under budget when K agents have succeeded and over budget when $K+1$ of them have done so, regardless of *when* these successes occurred.

6.4 Limited commitment

Here we assume that the principal is unable to credibly communicate with agents mid-contest—e.g., due to her interest in having agents work for as long as possible—and can commit (e.g., via a court-enforced contract) only to ending the contest at a pre-specified date and allocating the prize at that time. While it is no longer possible to meet the sufficiency conditions of Lemma 1 in this case, an optimal (no-feedback, fixed-deadline) contest can be obtained nonetheless.[21]

We begin with a necessary condition for a contest to be incentive compatible.

Lemma 2. *Consider a no-feedback contest that promises agent i a 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 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. The first term on the left 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. Specifically, $\lambda^2 e^{-\lambda s}$ is the amount by which the success probability at some future date s is reduced when the agent spends effort at date t; thus, the second term aggregates all future reductions in instantaneous benefits that result from spending effort now.

To find a contest that maximizes effort, we can solve

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

[21] In practice the principal may also find it easier to commit to a contest that always pays out the full prize to someone who succeeded than to individual contracts; though a contest dominates even when individual contracts are possible.

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 9. *The contest with deadline T^{EGA} and an egalitarian prize allocation maximizes effort among no-feedback, fixed-deadline contests. Because all agents work continuously until either they succeed or the fixed deadline is reached, it achieves this effort in the shortest possible time.*

This contest splits the prize equally among all successful agents irrespective of *when* they happened to succeed. A simple intuition 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.

For further detail, consider a simple heuristic. Set $\lambda = 1$ 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 (IC) with equality at all $t \leq T^{\text{EGA}}$. Figure 1 plots the corresponding instantaneous marginal benefit $e^{T^{\text{EGA}}-t}c$, together with the agent's marginal cost. Notice that at every $t' \leq T^{\text{EGA}}$, the marginal benefit exceeds c by exactly area 1, which corresponds to the integral on the left-hand side of (IC).

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.

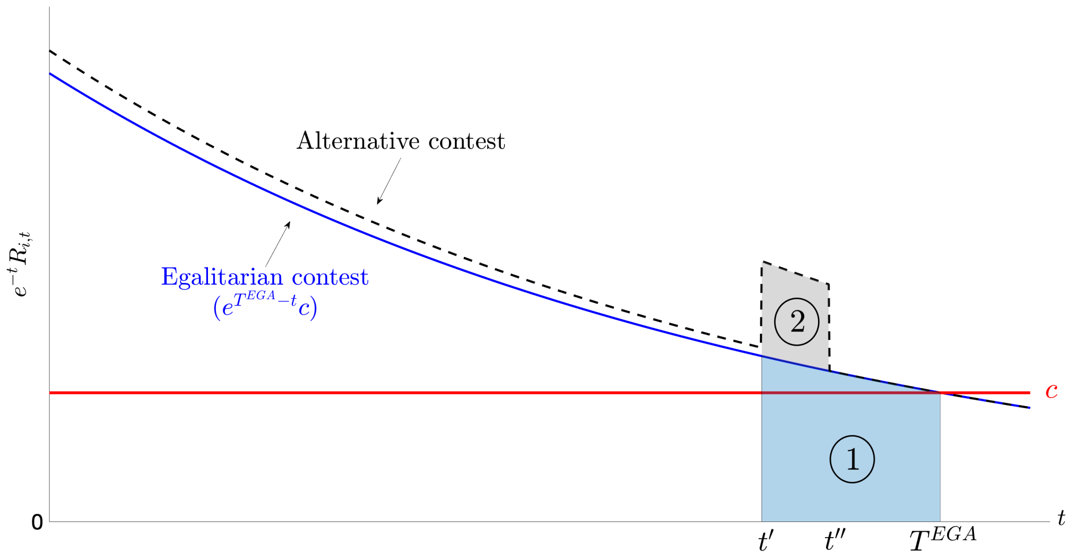

Figure 1: Meeting the incentive constraint.

7 Conclusion

We have obtained optimal contests for dynamic environments where the contest designer monitors effort through coarse, binary signals and has a fixed prize to award. The designer's primary goal—to maximize total effort or, equivalently, the number of successes—is achieved by any contest that spends all of the prize while granting the agents no rents. As it turns out, there is a variety of contests that do so. Her secondary goal—to obtain such effort in the shortest time possible—is achieved, in contrast, by only one family of contests, which we term 2^{nd} chance. These contests wait for some pre-specified number of successes to arrive, after which contestants are given one last chance to succeed. All optimal contests rely on $\mathcal{M}^{pronto}$ feedback, which keeps agents fully apprised of their own success; 2^{nd} chance contests achieve their goal regardless of what agents learn about their peers.

To find these contests amongst the vast set of possible contests, we first established an upper bound for total effort and a lower bound for expected duration given that effort, and then showed how the designer can achieve them. A challenge left for future work is finding fully-optimal contests for dynamic environments where attaining those bounds is impossible, such as when agents' successes are correlated or when they have heterogeneous effort costs. An alternative direction is letting the designer adjust the value of her prize. In those cases, as in the present setting, an adjustable deadline and a judicious choice of real-time feedback may help mitigate agency frictions.

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.

Doraszelski, Ulrich. 2003. "An R&D race with knowledge accumulation." *RAND Journal of Economics*, 20–42.

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.

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

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

Georgiadis, George. 2022. "Contracting with Moral Hazard: A Review of Theory & Empirics."

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.

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

Khorasani, Sina, Ersin Körpeoğlu, and Vish Krishnan. 2021. "Dynamic Development Contests." *Available at SSRN 3804454*.

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

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

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

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.

Moscarini, Giuseppe, and Lones Smith. 2011. "Optimal dynamic contests."

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.

Ryvkin, Dmitry. 2020. "To fight or to give up? Dynamic contests with a deadline." Working paper.

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 Proposition 2

Because the contest does not end until at least one agent succeeds, the first condition in Lemma 1 is satisfied. Therefore, given that the feedback policy is $\mathcal{M}^{pronto}$, it suffices to show that each agent's expected reward conditional on succeeding is time-invariant and equal to c/λ. This will imply that each unsuccessful agent is just willing to work and earns no rents.

Fix a time $t > T^*$ such that the contest is in progress, pick an agent who has yet to succeed, and suppose that he assigns probability $1 - c/\lambda$ to the event that at least one of his rivals has succeeded. Let $q_{rate}dt := (n-1)\lambda dt/(1 - c/\lambda)$ denote the probability that the contest ends during the interval $(t, t + dt)$ conditional on at least one success having occurred. By Bayes' rule, this agent's updated belief at $t + dt$ that at least one of his rivals has succeeded is equal to

$$\frac{(1 - c/\lambda)(1 - q_{rate}dt) + (c/\lambda)(n-1)\lambda dt}{(1 - c/\lambda)(1 - q_{rate}dt) + c/\lambda} = 1 - c/\lambda$$

as desired, where we substituted the value of q_{rate}.[22] Therefore, if unsuccessful agents assign probability $1 - c/\lambda$ to the event that at least one of their rivals has succeeded by any $t > T^*$, they will also do so at every $t' > t$. In that case, their expected reward conditional on succeeding is $(c/\lambda) \times 1 + (1 - c/\lambda) \times 0 = c/\lambda$ as desired. (Recall that an agent who succeeds after T^* wins the entire prize if he is the first to succeed, and none of it otherwise.)

Next, we show that at T^*, each unsuccessful agent indeed assigns probability c/λ to the event that none of their rivals have succeeded. First, we claim that the probability that no other agent has succeeded by T^*, $e^{-\lambda(n-1)T^*} < c/\lambda$.[23] Given that

[22] The first term in the numerator captures the probability that at least one agent has succeeded by t times the conditional probability that the contest did not end during $(t, t+dt)$. The second term captures the probability that the first success arrived during $(t, t + dt)$. The denominator captures the probability that the contest is not terminated by $t + dt$.

[23] To see why, observe that the expected reward of success during the first phase is $\sum_{k=0}^{n-1} \Pr\{k \text{ others succeed by } T^*\}/(1 + k) = c/\lambda$ by the choice of T^*. If the probability that an agent assigns to the event that nobody succeeds by T^* is weakly greater than c/λ, then the first term in the above sum is by itself weakly larger than c/λ, which is a contradiction.

the contest is terminated at T^* with probability

$$q_{T^*} := \frac{c/\lambda - e^{-\lambda(n-1)T^*}}{(1 - e^{-\lambda(n-1)T^*})(c/\lambda)}$$

if at least one agent has succeeded, by Bayes' rule each unsuccessful agent's posterior belief that none of their rivals have succeeded is

$$\frac{e^{-\lambda(n-1)T^*}}{e^{-\lambda(n-1)T^*} + (1 - q_{T^*})[1 - e^{-\lambda(n-1)T^*}]} = \frac{c}{\lambda}$$

as desired, where we have substituted the expression for q_{T^*}. We have therefore shown that from T^* onward, each unsuccessful agent assigns probability c/λ to the event that nobody has succeeded yet, and so his expected reward from success is c/λ.

It remains to show that each agent's expected reward from success equals c/λ at every $t < T^*$ as well. Recall that T^* has the property that when the prize is shared equally among the agents who succeed during $[0, T^*]$, each agent is kept fully apprised of his own success, then his expected reward from success is c/λ, as desired.

A.2 Proof of Proposition 3

We establish the proof assuming the countdown is triggered publicly, and then argue that it does not matter whether agents know what phase they are in.

For given α and T^c, the expected rewards from success are

$$R^1(\alpha, T^c) = \mathbb{E}\left[\frac{\alpha}{\alpha + m^1(T^c)}\right] \quad \text{and} \quad R^2(\alpha, T^c) = \mathbb{E}\left[\frac{1}{\alpha + 1 + m^2(T^c)}\right]$$

for the first phase and the countdown phase, respectively, where we have omitted the time subscript due to rewards being time-invariant, $m^1(T^c) \sim binom(n-1, 1-e^{-\lambda T^c})$ and $m^2(T^c) \sim binom(n-2, 1-e^{-\lambda T^c})$. That is, if an agent succeeds during the next instant and is the first to do so (which will trigger the countdown), then he shares the prize (not equally) with $m^1(T^c)$ of his rivals. On the other hand, any agent who succeeds during the countdown phase shares the prize with $1 + m^2(T^c)$ of his rivals. Because the feedback policy is $\mathcal{M}^{pronto}$ and the contest does not end until at least one agent has succeeded, it suffices to show that there exist α and T^c such that $R^1(\alpha, T^c) = R^2(\alpha, T^c) = c/\lambda$.

Fix any $\alpha \in (0, 1]$, and observe that $\lim_{T^c \to 0} R^1(\alpha, T^c) = 1$, $R^1(\alpha, T^c)$ continuously decreases in T^c, and it converges to $\alpha/(\alpha+n-1)$ as $T^c \to \infty$ (assuming each agent works until he succeeds). Since $\alpha/(\alpha+n-1) \le 1/n < c/\lambda < 1$, by the intermediate value theorem there is a $T^1(\alpha)$ such that $R^1(\alpha, T^1(\alpha)) = c/\lambda$. By the implicit function theorem, $T^1(\alpha)$ increases in α.

Similarly, for any $\alpha > 0$, $\lim_{T^c \to 0} R^2(\alpha, T^c) = 1/(\alpha+1)$, $R^2(\alpha, T^c)$ continuously decreases in T^c, and it converges to $1/(\alpha+n-1)$ as $T^c \to \infty$ (again assuming each agent works until he succeeds). Observe that $1/(\alpha+n-1) < c/\lambda \le 1/(\alpha+1)$ for all $\alpha \in (\underline{\alpha}, \overline{\alpha}]$ where $\underline{\alpha} = \max\{0, \lambda/c-(n-1)\}$ and $\overline{\alpha} = \min\{\lambda/c-1, 1\}$. Therefore, for any such α, by the intermediate value theorem there exists a $T^2(\alpha)$ such that $R^2(\alpha, T^2(\alpha)) = c/\lambda$. By the implicit function theorem, $T^2(\alpha)$ decreases in α.

If $\overline{\alpha} < 1$ (or equivalently $\lambda/c < 2$), because $\lim_{T^c \to 0} R^2(\overline{\alpha}, T^c) = 1/(\overline{\alpha}+1) = c/\lambda$ by definition, we have $\lim_{\alpha \to \overline{\alpha}} T^2(\alpha) = 0$. If instead $\overline{\alpha} = 1$, because $R^2(1, T^c)$ decreases in T^c, $R^2(1, T^c) < R^1(1, T^c)$, and $R^1(1, T^*) = c/\lambda$ by the definition of T^* (given in Section 4.1), it must be the case that $T^2(1) < T^*$.

Recall that $T^1(\alpha)$ increases in α and notice that $T^1(1) = T^*$. On the other hand $\lim_{\alpha \to \underline{\alpha}} T^2(\alpha) > T^*$, $T^2(\alpha)$ decreases in α, and $T^2(\overline{\alpha}) < T^*$.[24] Therefore there exists an α^* such that $T^1(\alpha^*) = T^2(\alpha^*)$. Letting $T^{c*} = T^1(\alpha^*)$, we have $R^1(\alpha^*, T^{c*}) = R^2(\alpha^*, T^{c*}) = c/\lambda$ as desired.

We have shown that each (unsuccessful) agent's expected reward from success is time-invariant and equal to c/λ. This implies that it is immaterial whether agents are told what phase they are in.

A.3 Proof of Proposition 4

To implement effort $1/c$, a contest must maximize total value and leave agents with zero rents. To satisfy the first criterion, it must award the entire prize with probability 1, and hence must not end before at least one agent has succeeded.

Next we turn to full rent-extraction. Let $p_{i,t}$ denote the probability that agent i assigns at time t to the event that he has already succeeded. Then, agent i's expected payoff can be expressed as $\mathbb{E} \int_0^\tau [\lambda(1-p_{i,t})R_{i,t} - c] a_{i,t} dt$. Therefore, for a mechanism to leave agents with no rents, it must be the case that in equilibrium for each agent i,

[24] If $\underline{\alpha} = 0$, because $\lim_{\alpha \to \underline{\alpha}} R^2(\alpha, T^c) = \mathbb{E}[1/(1+m^2(T^c))] > \mathbb{E}[1/(1+m^1(T^c))]$, $R^2(\alpha, T^c)$ decreases in T^c, and $\mathbb{E}[1/(1+m^1(T^*))] = c/\lambda$ by definition, it must be the case that $\lim_{\alpha \to \underline{\alpha}} T^2(\alpha) > T^*$. If instead $\underline{\alpha} > 0$, then by definition $1/(\underline{\alpha}+n-1) = c/\lambda$, and so $\lim_{\alpha \to \underline{\alpha}} T^2(\alpha) = \infty$. To see why $T^2(\overline{\alpha}) < T^*$, observe that either (i) $\overline{\alpha} < 1$ and $\lim_{\alpha \to \overline{\alpha}} T^2(\alpha) = 0$, or (ii) $\overline{\alpha} = 1$ and $T^2(1) < T^*$.

$\lambda(1-p_{i,t})R_{i,t} = c$ for all t such that $a_{i,t} = 1$. We claim that this condition holds if and only if in equilibrium $p_{i,t} = 0$ and $R_{i,t} = c/\lambda$ for all i and t such that $a_{i,t} = 1$. Towards a contradiction, suppose there is an interval $(t', t'+dt)$ during which $\mathbb{E}[p_{i,t}|p_{i,t} < 1] > 0$ for some agent i. For the mechanism to incentivize effort meanwhile extracting all rents, it must be the case that $\mathbb{E}[R_{i,t}] > c/\lambda$ during that interval. Suppose this agent deviates from the equilibrium and exerts effort during that interval only (while shirking at all other times). This agent privately knows that he has not succeeded by t', and the fact that $\mathbb{E}[R_{i,t}] > c/\lambda$ during that interval means that he can earn strictly positive rents—a contradiction. Therefore, any mechanism that leaves agents with no rents must have $p_{i,t} = 0$ (i.e., it must feature the $\mathcal{M}^{pronto}$ feedback policy which keeps agents fully apprised of their own successes) and $R_{i,t} = c/\lambda$ for all i and t such that $a_{i,t} = 1$.

A.4 Proof of Proposition 5 and Corollary 1

We begin by establishing a lower bound for the expected duration of a contest that implements total expected effort $1/c$. Fix a contest that implements that effort level, and for each $j \in \{1, \dots, n\}$, let T_j denote the expected amount of time that exactly j agents are working. Note that total effort $\sum_{j=1}^{n} jT_j = 1/c$ with $T_j \leq 1/(\lambda j)$ owing to the fact that when j agents are working, the next success occurs with Poisson rate λj. Let $D := \sum_{j=1}^{n} T_j$ denote the expected duration of the contest.

To obtain our lower bound, we introduce an auxiliary program that represents the problem of minimizing the duration of a contest by directly selecting the values of T_j, without regard for whether these values can be implemented in an actual contest as follows

$$\min_{T_1,\dots,T_n} \sum_{j=1}^{n} T_j \quad \text{s.t.} \quad \sum_{j=1}^{n} jT_j = \frac{1}{c} \quad \text{and} \quad 0 \leq T_j \leq \frac{1}{\lambda j}. \tag{4}$$

We claim that this program has the following unique solution:

$$\underline{T}_j = \begin{cases} 1/(\lambda j) & \text{if } j > n - K^* \\ (\lambda/c - K^*)/[\lambda(n - K^*)] & \text{if } j = n - K^* \\ 0 & \text{if } j < n - K^*, \end{cases}$$

where $K^* = \lfloor \lambda/c \rfloor$. To see why, observe that the auxiliary program is linear and in the first constraint T_j has a greater weight the larger the j. As a result, there is a cutoff j^* such that it is optimal to set T_j to its upper bound for all $j > j^*$, to its lower bound for all $j < j^*$, and to a possibly interior value for $j = j^*$.[25] To solve for j^* and T_{j^*}, we substitute these values of T_j into the constraint, which yields

$$\sum_{j=1}^{n} j\underline{T}_j = \sum_{j=1+j^*}^{n} \left(j \times \frac{1}{\lambda j} \right) + j^* \underline{T}_{j^*} = \frac{n - j^*}{\lambda} + j^* \underline{T}_{j*} = \frac{1}{c}.$$

Because $j^* T_{j^*} \leq 1/\lambda$ (from the last constraint), it follows that $j^* = n - K^*$ and $\underline{T}_{j^*} = (\lambda/c - \lfloor \lambda/c \rfloor)/(\lambda j^*)$, as desired. We have therefore shown that every contest that implements total expected effort $1/c$ has expected duration $D \geq \sum_{j=1}^{n} \underline{T}_j$.

Next, we show there exits a contest that achieves this lower bound, and hence this bound corresponds to the minimum contest duration among effort-maximizing contests. Consider the egalitarian 2^{nd} chance contest described immediately following the proposition.

Observe that each agent who succeeds during the first phase (i.e., before K^* agents succeed) receives prize $1/K^*$ if there is no success in phase 2, which occurs with probability $e^{-\lambda(n-K^*)T^{sc}}$, and otherwise receives $(1 - c/\lambda)/K^*$. Therefore, his expected reward from success is equal to

$$\frac{1}{K^*} \times e^{-\lambda(n-K^*)T^{sc}} + \frac{1 - c/\lambda}{K^*} \times [1 - e^{-\lambda(n-K^*)T^{sc}}] = \frac{c/\lambda}{K^*} \times (1 + K^* - \lambda/c) + \frac{1 - c/\lambda}{K^*} = \frac{c}{\lambda},$$

where we used the expression for T^{sc}. This implies that he is just willing to work until he succeeds, and earns no rents. During the second phase, by construction, if an agent succeeds, his prize is c/λ, so again he is just willing to work while earning no rents. Since the entire prize is awarded with probability 1, the contest satisfies the criteria of Lemma 1, and therefore implements maximal total effort $1/c$.

Notice that with probability 1, the contest ends no earlier than K^* and no later than $K^* + 1$ agents have succeeded. Moreover, the expected amount of time that

[25] T_{j^*} takes an interior value whenever λ/c is not an integer; otherwise it is zero.

$n - K^*$ agents work,

$$\begin{aligned} T_{n-K^*} &= \int_0^{T^{sc}} \lambda(n-K^*) t e^{-\lambda(n-K^*)t} dt + T^{sc} e^{-\lambda(n-K^*)T^{sc}} \\ &= \frac{1 - \lambda(n-K^*)T^{sc} e^{-\lambda(n-K^*)T^{sc}} - e^{-\lambda(n-K^*)T^{sc}}}{\lambda(n-K^*)} + T^{sc} e^{-\lambda(n-K^*)T^{sc}} \\ &= \frac{\lambda/c - K^*}{\lambda(n-K^*)} = \underline{T}_{n-K^*}, \end{aligned}$$

and so the expected duration of the contest meets the lower bound established above.

We are now ready to establish the proposition. We begin with necessity ($\Longrightarrow$). Consider a contest that implements maximal effort. Because the auxiliary problem (4) has a unique solution, every contest that achieves the lower bound on duration must be such that $T_j = \underline{T}_j$ for all j. As a result, the contest must never end before the K^* agents have succeeded, it must never continue after $K^* + 1$ agents have succeeded, and because $T_j^* \leq 1/(\lambda j^*)$, it must end with positive probability before $K^* + 1$ agents succeed. To be specific, because each successful agent must earn c/λ in expectation and the prize is unit-sized, letting q denote the probability that a success occurs during the countdown phase, it must be the case that

$$(1-q)\frac{1}{K^*} + q\frac{1 - c/\lambda}{K^*} = \frac{c}{\lambda} \Leftrightarrow q = \frac{\lambda}{c} - K^*.$$

Thus, the contest must be 2^{nd} chance with $K = \lfloor \lambda/c \rfloor$ and a deadline such that a success occurs during the countdown phase with probability $\lambda/c - \lfloor \lambda/c \rfloor$.

Finally, we turn to sufficiency ($\Longleftarrow$). Consider a second chance contest that implements expected effort $1/c$. Because the contest is 2^{nd} chance, there is a cutoff $\widehat{j}$ such that $T_j = 1/(\lambda j)$ for all $j > \widehat{j}$, $T_j = 0$ for all $j < \widehat{j}$, and because total effort is $1/c$,

$$\widehat{j} T_{\widehat{j}} = \frac{1}{c} - \sum_{j=\widehat{j}+1}^{n} j \times \frac{1}{\lambda j} = \frac{1}{c} - \frac{n - \widehat{j}}{\lambda}.$$

Because $T_{\widehat{j}} \leq 1/(\lambda \widehat{j})$, it must be the case that $\widehat{j} = j^*$ and hence $T_j = \underline{T}_j$ for all j. Hence, this contest achieves the lower bound on expected duration.

A.5 Proof of Proposition 6

It suffices to show that each agent's expected reward from an additional success is always equal to c/λ. Towards this goal, suppose the contest continues after the K^{*th} success. Then a K^*+1^{st} success will be obtained with certainty, and since this success is rewarded with c/λ, each agent irrespective of how many times he has already succeeded, is just willing to work during this second phase. During the first phase, each agent's expected reward from an additional success is equal to

$$(1-\lambda/c+K^*)\times\frac{1}{K^*}+(\lambda/c-K^*)\times\frac{1-c/\lambda}{K^*}=\frac{c}{\lambda},$$

as desired. That is, with probability $1-\lambda/c+K^*$ he earns $1/K^*$ per success and otherwise he earns $(1-c/\lambda)/K^*$.[26]

It follows that all agents work throughout the contest, they earn no rents, and the entire prize is awarded with probability 1. Therefore, this contest implements expected effort $1/c$ in the shortest duration possible.

A.6 Proof of Proposition 7

For a given contest and time t, let I_t denote the set of agents who have already succeeded (recall that $M_t=\sum_{i\in I_t}m_i$), let J_t denote the set of unsuccessful agents such that $m_j+\sum_{i\in I_t}m_i\geq 1$ for any given agent $j\in J_t$, and let K_t denote the remaining unsuccessful agents.

Per Lemma 1 and because the prize is awarded with probability 1 and the feedback policy is $\mathcal{M}^{pronto}$, it suffices to show that agent i's expected reward from succeeding is always m_i. Assume for now that the sets I_t, J_t, and K_t are common knowledge, though we will argue below that this knowledge is immaterial.

Consider the expected reward of agent i if he succeeds at the next instant. If $i\in J_t$, M_t weakly exceeds 1 upon the arrival of his success, he receives m_i, and the contest ends by Proposition 7(i). Otherwise, letting q_t (a function of I_t,J_t,K_t) denote the probability that the contest ends if he succeeds and noting that $M_t<1$,

[26] The feedback policy is immaterial because during the first phase agents know that the contest won't end until at least K^* successes occur, so whether they observe their own or others' successes does not affect their incentives. During the countdown–provided there is one—agents know that the K^*+1^{st} success has not occurred yet regardless of the feedback policy. Moreover, the agents' beliefs about which phase the contest is in are irrelevant for their incentives because their expected reward from succeeding is the same in both phases.

his expected reward is

$$q_t \times \frac{m_i}{M_t} + (1-q_t) \times \frac{\sum_{k \in J_t} \lambda_k}{\sum_{k \in J_t \cup K_t} \lambda_k} \times \left(1 - \frac{|J_t|}{\sum_{k \in J_t} \lambda_k / c}\right) \times \frac{m_i}{M_t}$$
$$+(1-q_t) \times \frac{\sum_{k \in K_t} \lambda_k}{\sum_{k \in J_t \cup K_t} \lambda_k} \times v_{i,t} \tag{5}$$

Let us unpack this expression. With probability q_t, the contest ends and i's prize is that in Proposition 7(ii). Otherwise, the contest continues and with probability $\sum_{k \in J_t} \lambda_k / \sum_{k \in J_t \cup K_t} \lambda_k$, the next success is achieved by an agent $j \in J_t$, in which case the contest ends and i's prize is the expected value of the expression given in Proposition 7(i).[27] With the remaining probability the contest continues past the next success and i's expected prize, which we denote $v_{i,t}$, remains to be determined.

Our aim is to pick q_t such that $v_{i,t}$ and (5) are equal to m_i for all i and t. Specifically, we claim that

$$q_t = \frac{\frac{\sum_{k \in J_t} \lambda_k}{\sum_{k \in J_t \cup K_t} \lambda_k} \times \left(M_t - 1 + \frac{|J_t|}{\sum_{k \in J_t} \lambda_k / c}\right)}{1 - M_t + \frac{\sum_{k \in J_t} \lambda_k}{\sum_{k \in J_t \cup K_t} \lambda_k} \times \left(M_t - 1 + \frac{|J_t|}{\sum_{k \in J_t} \lambda_k / c}\right)}$$

achieves this goal. To see why, observe that because the set K_t is always finite and $\sum_{i=1}^{n} c/\lambda_i > 1$, there is a subgame that is reached with positive probability where this set is empty. In that subgame, $v_{i,t}$ is irrelevant (as the contest will end for sure no later than upon the next success); thus, after substituting the expression for q_t, (5) equals m_i. Working backward in time, it follows that if $v_{i,t} = m_i$ and q_t satisfies the above expression, then (5) is equal to m_i, as desired. Notice that if $|J_t| > 0$, because $0 < 1 - M_t < c|J_t| / \sum_{k \in J_t} \lambda_k$, q_t is positive and strictly smaller than one, and if $|J_t| = 0$, $q_t = 0$ as stated in Proposition 7(iii).

Finally, because i's expected reward is m_i for any I_t, J_t, and K_t, and all agents work until they succeed, it is immaterial whether these sets are observed.

A.7 Proof of Proposition 8

Since the prize is awarded with probability 1 and the feedback policy is $\mathcal{M}^{pronto}$, it suffices to show that each agent's expected reward conditional on succeeding at t

[27] Agent j is awarded $[\sum_{k \in J_t} \lambda_k \times (c/\lambda_k)]/[\sum_{k \in J_t} \lambda_k] = c|J_t| / \sum_{k \in J_t} \lambda_k$ in expectation, and the remaining prize is split among the remaining successful agents in proportion to their values of m_i.

equals c/λ_t for all t. A key thing to note is that if an agent ever shirked prior to date t, 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, he cannot extract positive rents by strategically withdrawing effort.

Let I_t denote the set of agents who have succeeded by t, $k_t := n-|I_t|$ the number of agents who have not succeeded by t, $M_t := \sum_{i\in I_t} c/\lambda_{\tau(i)}$, and Δ_t the unique solution to $c/\lambda_{t+\Delta_t} + M_t = 1$ for any given $M_t \in [1-c/\lambda_t, 1)$.[28]

If an agent succeeds at t such that $M_t + c/\lambda_t \geq 1$, per Proposition 8(i) his reward is c/λ_t, as desired. Suppose instead he succeeds at a t such that $M_t + c/\lambda_t < 1$. Here the contest ends instantly with some probability q_t and otherwise continues until at least the next success. If that next success arrives at some $s \in (t, t+\Delta_t]$, which occurs with probability $1-e^{-k_t\int_t^{t+\Delta_t}\lambda_r dr}$, the contest ends at that moment, and agent i earns in expectation

$$\left(1-\int_t^{t+\Delta_t} k_t\lambda_s\frac{c}{\lambda_s}e^{-k_t\int_t^s\lambda_r dr}ds\right)\times\frac{c/\lambda_{\tau_i}}{\sum_{i\in I_t}c/\lambda_{\tau_i}} = \left(1-ck_t\int_t^{t+\Delta_t} e^{-k_t\int_t^s\lambda_r dr}ds\right)\times\frac{m_i}{M_t},$$

where the term in the parenthesis is the expected value of the prize money that remains after awarding c/λ_s to the last agent to succeed. If instead the next success (after t) arrives at some $s > t+\Delta_t$, then $M_s < 1$ and the contest continues. Denote by $v_{i,t}$ the agent's expected reward in that contingency.

We wish to show that the agent's expected reward from succeeding at t is

$$q_t\frac{m_i}{M_t} + (1-q_t)(1-e^{-k_t\int_t^{t+\Delta_t}\lambda_r dr})\left(1-ck_t\int_t^{t+\Delta_t}e^{-k_t\int_t^s\lambda_r dr}ds\right)\frac{m_i}{M_t}$$
$$+(1-q_t)e^{-k_t\int_t^{t+\Delta_t}\lambda_r dr}\times v_{i,t} = m_i.$$

To this end, let us guess (and later verify) that $v_{i,t} = m_i$ for all t. Under this guess,

[28] Per Proposition 8(ii), if a success occurs at t, $M_t \in (1-c/\lambda_t, 1)$, and the contest does not end instantly, then it continues at least until the next success. If the next success occurs at say $s \leq t+\Delta_t$, the contest ends instantly at s. Otherwise, it ends at s with some positive probability strictly smaller than 1.

the above equation can be expressed as

$$\frac{q_t}{M_t} + (1-q_t)\left(1 - ck_t \int_t^{t+\Delta_t} e^{-k_t \int_t^s \lambda_r dr} ds\right)\left(\frac{1 - e^{-k_t \int_t^{t+\Delta_t} \lambda_r dr}}{M_t}\right) + (1-q_t) e^{-k_t \int_t^{t+\Delta_t} \lambda_r dr} = 1$$

$$\Leftrightarrow q_t = \frac{M_t - \left(1 - ck_t \int_t^{t+\Delta_t} e^{-k_t \int_t^s \lambda_r dr} ds\right)\left(1 - e^{-k_t \int_t^{t+\Delta_t} \lambda_r dr}\right) - M_t e^{-k_t \int_t^{t+\Delta_t} \lambda_r dr}}{1 - \left(1 - ck_t \int_t^{t+\Delta_t} e^{-k_t \int_t^s \lambda_r dr} ds\right)\left(1 - e^{-k_t \int_t^{t+\Delta_t} \lambda_r dr}\right) - M_t e^{-k_t \int_t^{t+\Delta_t} \lambda_r dr}}. \tag{6}$$

Since $1 - ck_t \int_t^{t+\Delta_t} e^{-k_t \int_t^s \lambda_r dr} < M_t < 1$, the numerator in (6) is positive and strictly smaller than the denominator. Therefore, $0 < q_t < 1$.

It remains to show that $v_{i,t} = c/\lambda_{\tau(i)}$ for all i and t. Recall that we have shown that if $v_{i,t} = c/\lambda_{\tau(i)}$, agent i's expected reward from success is $c/\lambda_{\tau(i)}$, as desired. Now consider a time t where all but one agent have succeeded (i.e., $k_t = 1$). Because $c/\lambda_t > 1/n$ for all t, the contest ends with certainty upon the next success. In that case, $\Delta_t = \infty$, $v_{i,t}$ is immaterial, and from (6), agent i's expected reward is c/λ_{τ_i}; therefore, working backward in time, it follows that each i's expected reward is c/λ_{τ_i}, as desired.

A.8 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 $T_i \leq T$, 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 $t + \Delta t$ 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.9 Proof of Proposition 9

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 agent works throughout the interval $[0, T^{\mathrm{EGA}}]$.*

Proof of Lemma 3. Consider any symmetric pure-strategy profile in which all agents work for a duration T. Then agent 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 agents other than i who also succeed, and the second equality is established by writing the binomial sum and rearranging terms.

Now, taking as given the strategy profile of the other agents, the net expected payoff of agent 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 agent 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 agents 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 agent 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 agent succeeds; i.e., the expression on the right-hand side of (BC). □

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

$$\max_{\{T_i\},\{R_{i,t}\}} \sum_{i=1}^{n} T_i \text{ subject to (IC) and (BC).} \tag{7}$$

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 (7) 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 (7). 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

$$\begin{aligned} Z_i(t) = \lim_{k\to\infty} Z_i^k(t) = \lim_{k\to\infty} \frac{1}{f(t)}\left[c - \int_t^{T_i} f'(s)Z_i^{k-1}(s)ds\right] \\ = \frac{1}{f(t)}\left[c - \int_t^{T_i} f'(s)Z_i(s)ds\right] \quad (8) \end{aligned}$$

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:

$$f(t)\tilde{R}_{i,t} + \int_t^{T_i} f'(s)\tilde{R}_{i,s}ds - c = 0. \tag{9}$$

Differentiating both sides of (9) 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 agent i equals $\int_0^{T_i} f(t)\tilde{R}_{i,t}dt = cF(T_i)/f(T_i)$, we have

$$c\sum_{i=1}^{n} \frac{F(T_i)}{f(T_i)} - \left[1 - \prod_{i=1}^{n}(1 - F(T_i))\right] \leq 0. \tag{10}$$

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 (10) 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 (10) 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

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

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 (10) 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 (7) 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 (7) can be rewritten as

$$\max_{T} \left\{ nT \text{ s.t. } cn\frac{F(T)}{f(T)} \leq 1 - [1 - F(T)]^n \right\}, \tag{11}$$

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 (11).

First notice that the constraint in (11) 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 (11) 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 (7), 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 (7). Therefore, this egalitarian contest is an optimal no-feedback contest.