Optimal Feedback in Contests*

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

December 8, 2020

Abstract

We derive optimal contests for environments where output takes the form of breakthroughs and the principal has an informational advantage over the contestants. Whether or not the principal is able to provide real-time feedback to contestants, the optimal prize allocation is *egalitarian*: all agents who have succeeded in a pre-specified time interval share the prize equally. When providing feedback is feasible, the optimal contest takes a stark cyclical form: contestants are fully apprised of their own success, and at the end of each fixed-length cycle, they are informed about peer success as well.

1 Introduction

Contests—situations where multiple agents compete for a prize—are a common way to organize economic activity: innovation races, promotions and other labor-market tournaments, all-pay auctions, athletic events, and legal battles all fall into this category.[1] Ever since the seminal work of Lazear and Rosen (1981), Green and Stokey (1983), and Nalebuff and Stiglitz (1983), researchers in economics, marketing and operations management have sought to understand how to best allocate the prize among

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

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

[1]Since 2010 federal agencies have conducted nearly 1,000 prize competitions, with the total amount of prize money growing from $247,000 in FY2011 to over $37 million in FY2018.

participants, and more recently, starting with the work of Yildirim (2005), Aoyagi (2010), Ederer (2010), Goltsman and Mukherjee (2011), and Halac, Kartik and Liu (2017), how best to disclose real-time information regarding the contestants' progress. While there is no one type of contest that is universally efficient, this literature has offered key general insights together with design ideas for specific environments of interest.

Here we are interested in scenarios, not formally considered so far, in which the contest designer (the principal) has an informational advantage over the contestants in terms of how well they are doing mid-contest. Our goal is to find the optimal contest—inclusive of prize-allocation and termination rules, as well as an information feedback policy—out of the full set of feasible designs. The principal's informational advantage may reflect, for example, the manager of a firm knowing better than their employees whether any one of them has met a given standard for promotion, or a company hosting an innovation race potentially having a clearer idea about the value of new technologies developed by the contestants.[2]

In our model, contestants exert binary flow effort, which they can interrupt/resume at any time, and their success takes the form of a Poisson breakthrough, which until the contest is over, only the principal can observe. The key challenge when searching for the optimal contest is the vast range of potential contest designs from which to choose, especially when the feedback policy is part of the design. We address this challenge by first providing a sufficient condition for a contest to be optimal—namely, that it maximizes total surplus while giving zero rents to the contestants—and then finding a contest that meets these criteria.

We first solve a baseline model where we adopt the common assumption that the agents' hazard rate of success depends only on their current effort, rather than on the full history of efforts. Whether or not the principal is able to provide real-time feedback, the optimal prize allocation is *egalitarian*; that is, all contestants who have succeeded in a pre-specified time interval share the prize equally, regardless of when they did so. This allocation smoothes incentives over time, which helps the principal maximally extend the period over which agents are willing to work.

When the principal can also design the feedback policy, the optimal contest keeps

[2]Take for instance the Netflix Prize, a competition that offered a $1 million prize for a predictive algorithm for users' film ratings that managed to improve upon Netflix' existing algorithm by at least 10%. Only Netflix was privy to the qualifying dataset needed to evaluate submissions, and hence to whether any contestant had achieved the goal.

each agent fully apprised as to whether they have succeeded, but provides them only periodic feedback on the success of their rivals. This contest takes a stark cyclical form: The principal first sets a provisional deadline, T^*. If one or more agents succeed by then, the contest ends and the prize is shared equally among those who succeeded so far. Otherwise, the contest resets and the deadline is extended until $2T^*$. The contest proceeds in this cyclical manner until the end of the first cycle in which one or more agents manage to succeed. Contestants receive feedback regarding their peers only at the end of each cycle, when they are either informed that the contest has ended, or that nobody has succeeded so far and it is time to extend the deadline. This contest is optimal because it maximally stretches out the time period over which agents are willing to work while extracting all rents from the contestants.[3]

Due to its simplicity, this cyclical egalitarian contest should not be difficult to implement in practice, as the only parameter that the designer needs to calibrate is the cycle length T^*, with a longer deadline possible the lower the cost of effort, the larger the probability of success, and the greater the prize.[4] While there is limited empirical evidence so far on the effectiveness of different contest designs, the field experiment conducted by Lim, Ahearne and Ham (2009) lends support to the effectiveness of an egalitarian prize, and the findings of Fershtman and Gneezy (2011) and Gross (2017) suggest that real-time feedback, with a flavor similar to that suggested by our model, can be effective at encouraging effort.[5]

An application that fits our simple setting especially well is the proof-of-work protocol at the core of digital currencies such as Bitcoin. The contestants (miners) exert effort in form of "hash submissions" with the goal of solving a cryptographic

[3] This contest maximizes both total expected effort and the expected number of successes.

[4] The Netflix Prize competition had a similar flavor, for example. If no team achieved the initial goal (a greater than 10% improvement in prediction accuracy) after a year, it allocated a smaller "milestone prize" ($50K) to the best-performing team so far, provided it achieved some improvement (at least 1%). This process was to be repeated each year until a team achieved >10% improvement. At that point, a 30-day countdown clock would start ticking, at the end of which a final winner would be declared. Our predicted contest is closely related, but simpler, as it does not prescribe smaller "milestone" prizes (as in our setting, success is all or nothing), ties are much more likely, and following the first success, rather than having a final 30-day period to catch up, contestants have until the end of the current cycle to do so.

[5] Gross (2017) analyzes a sample of 4294 winner-takes-all logo design competitions, and finds that feedback improves the quality of subsequent submissions, and on net increases the number of high-quality ideas. Moreover, simulations under various feedback policies indicate that private feedback allows players to improve their submissions without exposing performance differences that discourage continued participation.

puzzle, which is followed by a monetary reward. As in our model, success arrives at a constant Poisson rate—an intentional design characteristic—and marginal costs, which stem from the cost of electricity, are essentially constant. In addition, the designer can in principle reveal to each contestant as much or as little information as she wishes about their rivals' state of success (and the same goes for a contestant's own performance). Because of this close fit, we are able to carry out a calibration exercise that suggests that Bitcoin's current contest, which is winner-takes-all, underperforms the cyclical egalitarian contest suggested by our model by a considerable margin. We argue that by adopting the cyclical contest, Bitcoin would be able obtain between 83% and 126% additional effort for the same prize, depending on the price of electricity—or equivalently, it would be able to reduce the prize by 44%-54% while securing the same level of total effort.

We also consider some extensions of the model that suggest a degree of robustness to our results. When the number of contestants is unknown, the cyclical contest remains optimal, and when the contestants' hazard rate of success is increasing in both current and past effort, the optimal contest has a very similar cyclical structure as before: it features an egalitarian prize and a feedback policy that keeps agents fully apprised of their own success, but only periodically informs them of their rivals' status. The only difference is that the length of each cycle is stochastic, as this allows the principal to frontload incentives.

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, and some prize-sharing may be optimal otherwise. Extensions to stochastic output, arbitrary risk-preferences and heterogeneous agents are considered by Drugov and Ryvkin (2019, 2020) and Olszewski and Siegel (2020), among others.[6]

Fang, Noe and Strack (2018) find that aggregate effort in all-pay contests decreases in their competitiveness, as measured by the dispersion of prizes, contest crowding, and the number of contestants. Letina, Liu and Netzer (2020) consider a generalized version of that framework. They find that for n contestants, a nested Tullock contest

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

featuring $n-1$ equal prizes is optimal. While our work differs in that these papers study static environments with no scope for feedback design, the idea that "turning down the heat" motivates more effort echoes the optimality of an egalitarian contest in our dynamic framework.

Taylor (1995) considers a dynamic contest where players invest in an innovation of stochastic quality. After T periods, the player with the highest-quality innovation wins a prize. The principal chooses the prize and an entry fee, which determines the number of players, to maximize her profit. In the optimal contest, players invest in a given period as long as their highest-quality innovation to date is below a threshold, and the entry fee is chosen to extract all rents. Benkert and Letina (2020) extend this framework by incorporating interim transfers and an endogenous termination date. The optimal contest ends as soon as the highest-quality innovation exceeds a threshold, and agents invest up until the end of the contest.

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

Our paper is also related to a growing literature on contests involving experimentation, where the feasibility of success is initially unknown. Halac, Kartik and Liu (2017) consider an experimentation framework such as the one in Bonatti and Horner (2011), but with a designer who chooses a prize-sharing scheme and a feedback policy to maximize the probability of a success. Within the class of rank-monotonic prize schemes and deterministic and symmetric disclosure policies, a *cutoff-disclosure equal-sharing contest* is optimal. This contest provides no interim feedback and ends as soon as a critical number of agents have succeeded, each winning the prize with equal probability. 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" after which successes are disclosed immediately dominates contests with a constant probabilistic disclosure (including those with no or full disclosure). Our work differs in that successes are observed only by the principal, and that we manage to solve for the fully optimal contest.

2 Model

A principal designs a contest to motivate $n \geq 2$ agents to spend effort. The contest consists of a termination rule, which specifies *when* the contest will end, a rule for allocating a prize, whose value we normalize to \$1, and a feedback policy, which specifies the information transmitted to each agent at every moment in time. We formalize these objects below. The principal's objective is to maximize the expected total effort exerted throughout the contest.

At each instant t of continuous time, each player observes any message sent from the feedback policy and decides whether to spend effort. Effort is costly but potentially produces a "success." In particular, if player i spends effort for a total duration a_i he incurs cost ca_i, where c is the constant marginal cost of effort. While the agent spends effort, success arrives stochastically with constant instantaneous rate $\lambda > c$. That is, a player can succeed at most once and conditional on not having succeeded by time t, effort for an additional duration dt produces a success during the time interval $(t, t + dt)$ with probability λdt. Thus, $F(t) = 1 - e^{-\lambda t}$ is the probability that a player succeeds on or before date t if he spends effort continuously in that time, and $f(t) = F'(t) = \lambda e^{-\lambda t}$ is its "density".[7]

Each player observes his own effort, but not whether he has succeeded, or others' efforts and successes. Conversely, the principal observes successes but not efforts.

The principal's feedback policy specifies a message that she transmits to each agent at every moment as a function of her past observations. An example of a feedback policy that will be important for our results is the one we denote by $\mathcal{M}^{pronto}$, according to which the principal informs each player immediately if and when he succeeds, but otherwise keeps players uninformed. Alternative policies might inform agents about their or their rivals' successes probabilistically, about the feedback conveyed to rivals, and so forth.

[7]A constant hazard rate means there is no notion of progress over time. In Section 5, we extend our model to allow for an increasing hazard rate.

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

When the contest ends, if a_i is the total effort spent by player i, then his ex-post payoff is

$$u_i = q_i - ca_i.$$

There is no discounting and the players' objective is to maximize their expected ex-post 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 of the resulting contest is maximal among Bayesian Nash equilibria of a given set of contests. We will derive the optimal contest when there is no feedback and the fully optimal contest when the feedback policy is unconstrained.

Allowing the principal to be better informed than the agents concerning their success allows us to capture scenarios such as a manager being better informed about their employees' performance than the employees themselves, or the designer of an innovation contest knowing best what constitutes progress. Our main result would be unchanged if contestants were also privy to their own success.

3 No-Feedback Contests

In this section, as a benchmark, we characterize the optimal no-feedback contest, that is, the effort-maximizing contest within the class of contests that provide no feedback, terminate at a pre-specified deterministic time T, and allocate the \$1 prize according to some rule $\{q_i\}$.[8] Formally, we solve

$$\max_{T,\{q_i\},\{a_i\}} \sum_{i=1}^{n} a_i \tag{1}$$
$$\text{s.t. } \{a_i\} \text{ is an equilibrium profile.}$$

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

$$R_{i,t} = \mathbb{E}\left[q_i \mid \text{agent } i \text{ succeeds at } t\right], \tag{2}$$

which represents agent i's expected share of the prize conditional on succeeding at t. This object will allow us to decompose the problem and design optimal incentives for each agent separately.

To illustrate, consider the egalitarian contest that divides the prize equally among all players who succeed prior to the terminal date. Conditional on player i succeeding, his expected share of the prize is $R_{i,t} = \mathbb{E}[1/(1+M)]$ where M is the number of rivals who eventually succeed. An important observation is that this benefit is independent of the time at which i succeeds, and is the same for all players. In other words, the reward functions for the egalitarian contest are constant and symmetric across players. Any other contest with this property is for all intents and purposes identical to the egalitarian one, and so we shall treat it as being identical to the egalitarian design.[9]

For another example, consider the winner-takes-all contest. Conditional on player i succeeding at time t, he earns the entire prize if and only if he is the first to succeed. Thus, assuming all of his rivals have worked throughout the $[0, t]$ interval, $R_{i,t} = [1 - F(t)]^{n-1}$ since this is the conditional probability that he is the first to

[8] Observe that a contest can provide feedback indirectly via the termination rule. For example, if the contest ends as soon as the first success occurs, then at every moment the contest is still ongoing, agents will know that no one has yet succeeded, which can (and will) affect their incentives. To rule out such indirect feedback, we constrain the principal to choose a deterministic deadline.

[9] One example is a two-player contest that randomizes equally between awarding the prize to the first agent to succeed and awarding it to the last one to succeed, while not informing the agents of the outcome of this draw until after the contest is over.

succeed. Notice that this reward function is strictly decreasing in t.

Using the reward function, we can express agent i's expected utility from working continuously during an interval $[0, T_i]$ as

$$\int_0^{T_i} f(s) R_{i,s}\, ds - cT_i. \tag{3}$$

When searching for an optimal contest, one can restrict attention to contests where each agent indeed works continuously over some interval $[0, T_i]$. This is because neither the principal nor any agent gains from delays in effort, and hence any contest can be replaced, without loss, with one that frontloads all effort.

The following result provides a necessary condition for a reward function to incentivize continuous effort.

Lemma 1. *Consider a no-feedback contest that gives agent i the reward function $R_{i,t}$. Working continuously throughout $[0, T_i]$ is incentive compatible for this agent only if, for all $t \in [0, T_i]$,*

$$f(t) R_{i,t} + \int_t^{T_i} f'(s) R_{i,s} ds \geq c. \tag{IC}$$

This incentive constraint states that the marginal benefit of effort at time t, which is captured by the left-hand side, should be no smaller than the marginal cost. To understand the expression for the marginal benefit, note that the first term is the instantaneous marginal benefit of effort at time t. The second term captures a forward-looking incentive effect: success today precludes success in the future. In particular, $f'(s)$, which is negative, is the amount by which the success probability at some future date s is reduced when the agent spends effort at date t. The second term thus aggregates the reduction in future instantaneous benefits that results from spending effort in the current date.

Egalitarian Contest. When there is no feedback, the egalitarian contest that divides the prize equally among all successful agents admits a simple symmetric pure-strategy equilibrium.

Proposition 1. *Let T^{EGA} uniquely satisfy*

$$\frac{1 - \left(1 - F(T^{\text{EGA}})\right)^n}{nF(T^{\text{EGA}})} = \frac{c}{f(T^{\text{EGA}})}.$$

So long as the deadline is no smaller than T^{EGA}, the egalitarian contest has a symmetric pure-strategy equilibrium where each player spends total effort T^{EGA}.

The condition defining T^{EGA} has a simple intuition. As shown in the proof, located in Appendix A, the left-hand side equals the expected share of the prize earned by player i were he to succeed. Multiplying by $f(T^{\text{EGA}})$, the incremental probability of success from working beyond T^{EGA}, gives the marginal benefit from increasing total effort. The condition equates this marginal benefit to the marginal cost c.

We now use the incentive constraint in Lemma 1 to show that, absent feedback, the egalitarian contest is optimal.

Proposition 2. *The egalitarian contest with deadline T^{EGA} is optimal among no-feedback contests.*

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

Here is a more detailed heuristic argument that highlights the crucial role of Lemma 1. For brevity, restrict attention to symmetric contests with symmetric equilibria. Normalizing $\lambda = 1$, and substituting into (IC) the expressions $f(s) = e^{-s}$ and $f'(s) = -e^{-s}$ for the Poisson arrival process, yields the following necessary condition for the (symmetric) reward functions of such a contest:

$$\underbrace{e^{-t} R_{i,t}}_{\text{instantaneous marg. benefit at } t} - \int_t^T \underbrace{e^{-s} R_{i,s}}_{\text{impact of effort on future reward}} \geq c \tag{4}$$

for all t in $[0, T]$. Recall that the left-hand side is the full marginal benefit of effort at time t.

The constant reward function $R_{i,t} = e^{T^{\text{EGA}}} c$, which corresponds to the egalitarian contest, satisfies constraint (4) with equality at all $t \leq T^{\text{EGA}}$. Figure 1 plots the corresponding instantaneous marginal benefit schedule $e^{T^{\text{EGA}} - t} c$, together with the agent's marginal cost. Notice that at every $t' \leq T^{\text{EGA}}$, the instantaneous marginal benefit exceeds c by exactly area 1, which corresponds to the integral in the left-hand side of (4).

Consider now a non-egalitarian contest (with a non-constant reward schedule) that attempts to implement the same total effort as the egalitarian one. As illustrated in the figure, constraint (4) 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 (4) 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.

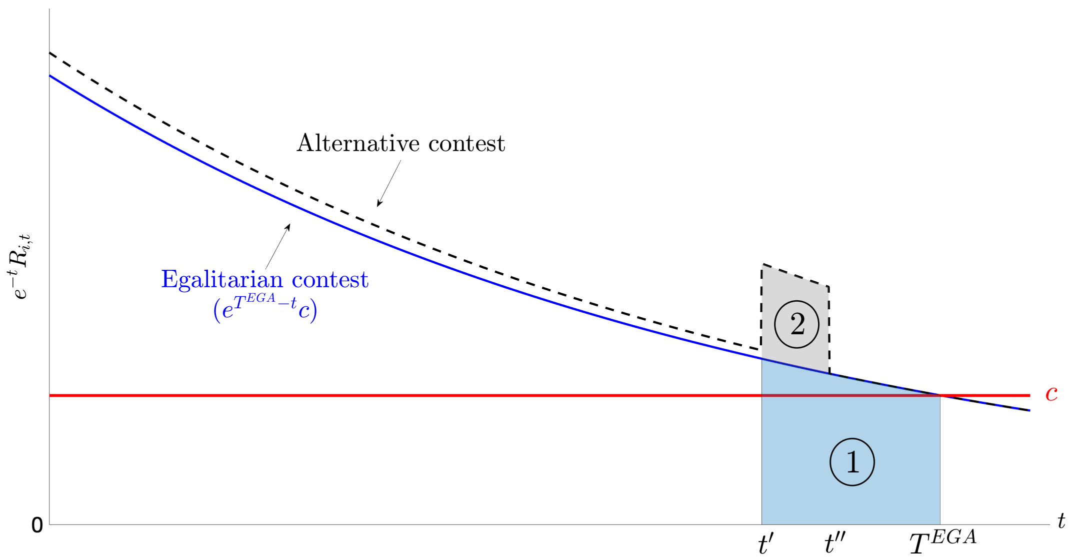

Figure 1: Meeting the incentive constraint.

Making use of this heuristic, it is easy to show why other seemingly-appealing contests are not optimal. In the winner-takes-all contest, for instance, the reward schedule is strictly decreasing over time. Therefore, per the above argument, in order to incentivize the same amount of effort it would need to lie strictly above the egalitarian one (that is, it over-rewards agents early on), which is not feasible given that the principal only has \$1 of prize money.

Consider now the polar opposite "last-takes-all" contest, where it is the last (rather than the first) to succeed who is guaranteed the prize. This contest in principle seems appealing because it maximally backloads incentives and therefore might appear to relax the incentive constraint. However, because the associated reward schedule is strictly increasing over time (i.e., a later success gives rivals a smaller chance to win), agents are tempted to delay their effort. The only way to prevent this is again to offer higher overall rewards, which is not feasible.

4 Optimal Contest With Feedback

In this section we characterize the optimal contest once feedback is allowed. We begin by proving a result that will serve as a key stepping stone for our analysis.

Lemma 2. *A contest is guaranteed to be optimal if, in equilibrium:*
(i) the prize is awarded with probability one, and
(ii) each agent earns zero rents.

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

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

$$\mathbb{E}\sum_{i=1}^{n} a_i = \frac{\sum_{i=1}^{n} \mathbb{E}q_i - \sum_{i=1}^{n} u_i}{c}.$$

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

The no-feedback contest we derived in Section 3 fails both conditions in this lemma and hence can in principle be improved upon once feedback is allowed. That it fails to award the entire prize follows from the fact that there is a chance that no agent has succeeded by the deadline T^{EGA}. That it leaves rents can be seen from Figure 1: the overall rents for each agent are equal to the area between the instantaneous marginal benefit and the marginal cost curves.

Our main result, Proposition 3, characterizes a contest that satisfies both criteria in the lemma under the assumption that the parameters of the model satisfy $n > \lambda/c$. This assumption means that there are enough competitors for a contest to be desirable in the first place. When the assumption fails, the principal could do at least as well by reserving $1/n$-th of the prize for each agent and contracting with each one individually.

The optimal contest has three crucial properties. First, its prize allocation is egalitarian, just as in the no-feedback case. Second, agents are fully apprised of their own success, which is achieved 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^* (which indirectly informs 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. This termination rule is formally 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 3. *Assume $n > \lambda/c$. The contest with egalitarian prize, the cyclical termination rule τ^*, and the feedback policy $\mathcal{M}^{pronto}$ is optimal. In this contest at least one agent succeeds, and hence the prize is awarded with probability one. Moreover, each agent obtains 0 expected utility and the principal's profit is $1/c$.*

This cyclical contest is optimal because it meets the requirements of Lemma 2: it awards the prize with probability 1 because the provisional deadline keeps getting extended if no agent has succeeded; and it grants agents zero rents owing to the $\mathcal{M}^{pronto}$ feedback policy combined with a provisional deadline T^* just long enough so that agents who have not yet succeeded are at each moment just barely willing to work.

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

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

$$R_{i,t} = \mathbb{E}\left[\frac{1}{1+M}\right] = \frac{1-e^{-\lambda n T^*}}{n(1-e^{-\lambda T^*})} = \frac{c}{\lambda}, \tag{5}$$

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 is shown in the proof of Proposition 1, and the third equality follows from the definition of T^*.

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

As it turns out, there are other optimal contests as well. All these other contests, however, have in common with our contest that they keep agents fully apprised of their own success and they have an egalitarian prize structure, and hence differ only in the details of the termination rule.[10]

The reason it is necessary to keep agents fully informed of their own success is that, otherwise, they would be able to obtain rents from the principal by strategically withdrawing effort.[11] The egalitarian rule is necessary, in turn, because given the $\mathcal{M}^{pronto}$ feedback policy, for a contest to extract all rents, each agent's expected reward conditional on succeeding must be c/λ regardless of when he happens to succeed. Non-egalitarian contests are unable to offer such time-invariant rewards.

We conclude this section with some remarks:

i. Because the cyclical contest keeps agents apprised of their own success, it would remain optimal if agents were able to observe this success directly. It is key, however, that agents do not learn about their rivals' successes until each cycle ends. The principal must therefore make sure that when informing a success-

[10] One example is a modified version of the contest given in Proposition 3 with an arbitrary provisional deadline $T > T^*$ and where agents work only a fraction T^*/T of the time.

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

ful agent, this communication cannot itself be credibly re-transmitted to other agents in a bid to discourage them. One of various ways to do so is via an anonymous message sent to an address specified by the receiver, whose legitimacy can therefore only be ascertained by that receiver.

ii. The cyclical contest maximizes the total expected number of successes, and not just total expected effort, as it guarantees that agents work until either they succeed or the contests ends.

iii. To implement the optimal contest, it suffices that the principal is able to commit to one cycle at a time, as after a cycle ends it is in her interest to commit to another identical one if no agent has yet succeeded.

5 Extensions

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

Unknown number of contestants. In practice, there are contests where participants do not know how many rivals they face, and perhaps even the principal ignores this as well. Fortunately, Proposition 3 immediately generalizes to the case where n is random and unknown, provided all agents share a common prior over n. The only thing that changes is that the provisional deadline T^* must now satisfy $\mathbb{E}\left[\left(1-e^{-n\lambda T^*}\right) / \left(n(1-e^{-\lambda T^*})\right)\right] = c/\lambda$, where the expectation is taken with respect to n. As before, this provisional deadline ensures that agents are just barely willing to work, thus allowing the principal to extract all rents, and the cyclical structure ensures that the prize is awarded with probability one.

Increasing hazard rate. In some settings, the agents' instantaneous probability of success might grow as they work and make progress on the problem.[12] To capture this possibility, let $F(t)$ denote the probability that an agent succeeds at or before date t if he works continuously until that time, and suppose the hazard rate $\lambda_t = F'(t)/[1-F(t)]$ exists and is weakly increasing.[13]

Under the assumption that $\lambda_t \in (c, nc)$ for all t and is differentiable almost everywhere, the optimal contest is similar to the one with constant hazard rate character-

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

[13] This implies that if an agent has spent s units of effort by some date, his hazard rate at t is λ_s.

ized in Proposition 3 except that the length of each provisional deadline is stochastic. In particular, it comprises the egalitarian prize, the feedback policy $\mathcal{M}^{pronto}$, and a cyclical termination rule which operates as follows: At time 0, the principal privately draws the duration of the first provisional deadline T_1 from some commonly-known distribution $G^0(\cdot)$. If at least one agent succeeds by this deadline, the contest ends and the prize is awarded according to the egalitarian rule. Otherwise, a new provisional deadline T_2 is drawn from a distribution $G^{T_1}(\cdot)$. The contest proceeds in this cyclical manner until the first deadline by which at least one agent has succeeded. Proposition 4, which is presented and proven in Appendix B, shows that this contest meets the conditions of Lemma 2, and is therefore optimal.

6 Proof of Work

Here we calibrate the parameters of our model in order to match Bitcoin's proof-of-work protocol.[14] This exercise will allow us to compare within the structure of the model, the performance of Bitcoin's current winner-takes-all design against the more efficient cyclical egalitarian design.

Contestants and effort. We equate a contestant to an ASIC miner (a machine specialized in competing for cryptocurrency), whose effort entails submitting candidate solutions to a cryptographic puzzle (i.e., calculating "hashes"). For concreteness, we consider the Bitmain Antminer S19 Pro Miner device, with an output of 110 trillion attempts (hashes) per second.[15]

Success rate. By design, the likelihood of success is time-invariant (i.e., progress does not accumulate) and is hence described by a Poisson process. One success currently takes in expectation 17.597×2^{32} trillion attempts, and so for the device in question, the success rate $\lambda = 110/(19.157 \times 2^{32}) = 1.337 \times 10^{-9}$ successes per second.[16]

Marginal costs. This cost comes primarily from the use of electricity, which each of our contestants consumes at a rate of 3.25 kW. We take the price of electricity to

[14]See Appendix C for a summary of Bitcoin's institutional details.

[15]https://www.asicminervalue.com/miners/bitmain/antminer-s19-pro-110th.

[16]The 19.157 trillion figure (retrieved on Dec. 4, 2020 from https://btc.com/stats/diff) is known as the "bitcoin difficulty." This parameter varies over time as the computing power on the network changes with the intention to maintain the average time between successes at approximately 10 minutes.

range between 7.71 and 9.56 cents per kWh, which corresponds to the average retail price of electricity in Louisiana and Illinois.[17] Thus, c ranges between $\$6.96 \times 10^{-5}$ and $\$8.63 \times 10^{-5}$ per second.

Prize. The current prize for a success is 6.25 Bitcoins, worth around $18,860 each (as of Dec. 4, 2020).

Principal's payoff. In the current winner-takes-all design, where the contest ends with the first contestant to succeed, total effort across all contestants combined is $1/\lambda = 7.48 \times 10^8$ seconds per success.[18] Under the cyclical egalitarian design, in contrast, the principal induces total effort $1/c$ per dollar of prize money, and hence obtains a payoff $\Pi^* = 6.25 \times \$18,860/c$, which for our range of electricity prices ranges between 13.66×10^8 and 16.94×10^8 seconds of activity per success, representing a 83-126% improvement.[19] Equivalently, the principal would be able to implement the same total effort using only 44-54% of the prize.

This large improvement can be traced to the fact that the current design leaves large rents to the contestants. For example, if we assume $n = 1$ million, which matches current estimates, each contestant earns in expectation $7.6 per day.[20] The cyclical egalitarian contest manages to extract any such rents by keeping contestants going during a cycle even if some of their rivals might have already succeeded.

Our model also suggests that a mining protocol with a constant rather than an increasing success rate is justified, as the former design induces the same total effort (i.e., $1/c$) and in addition it prevents any single miner from gaining too much power—which is an additional goal of Bitcoin's.

7 Conclusion

We have proposed a contest with an egalitarian prize, a cyclical structure involving a periodic resetting of the contest, and a partial type of feedback: leaders are im-

[17]Source: https://www.eia.gov/electricity/state. Louisiana prices are the lowest in the U.S.; Illinois prices are close to the median.

[18]In particular, the first success arrives in expectation after $(n\lambda)^{-1}$ units of time, and because all contestants exert effort until then, total effort is equal to $1/\lambda$. Note also that effort is incentive compatible as long as λ times the prize is greater than c.

[19]If for instance $n = 1$ million, the optimal contest would feature a cycle length between 16 and 24 minutes.

[20]Because there is free entry of miners, these rents are presumably dissipated in the form of investments in mining equipment.

mediately informed of their success and laggards are kept in the dark so as to not discourage further effort. In our setting, this contest manages to convert 100% of the prize money into effort (i.e., is maximally efficient) as it manages to extract all rents from the contestants.

This contest is attractive from an applied perspective because of its relative simplicity. It is also capable of delivering large efficiency gains relative to commonly-used contests such as the winner-takes-all design. In the proof-of-work application, for instance, our calibration exercise suggests that holding the prize fixed, the cyclical contest generates roughly twice the effort relative to the design currently in place.

Our model has abstracted from features—such as technological asymmetries across players, decreasing hazard rates of success, and the possibility that success is more continuous than all or nothing—that may be relevant for specific applications. When these features are present, the optimal control of information is likely to be more complex. We leave these possibilities for future work.

References

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

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

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

Bohme, Rainer, Nicolas Christin, Benjamin Edelman, and Tyler Moore. 2015. "Bitcoin: Economics, technology, and governance." *Journal of Economic Perspectives*, 29(2): 213–238.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Nakamoto, Satoshi. 2008. "Bitcoin: A peer-to-peer electronic cash system." *The Cryptography Mailing List*.

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

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

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

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

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

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

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

A Proofs

A.1 Proof of Lemma 1.

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

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

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

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

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

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

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

$$R_{i,t} f(t) + \int_t^{T_i} R_{i,s} f'(s) \geq c\,,$$

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

A.2 Proof of Proposition 1.

Consider any symmetric pure strategy profile in which each player works for a duration T. Let R denote the expected share of the prize enjoyed by player i should he succeed:

$$R = \mathbb{E}\left[\frac{1}{1+M}\right],$$

where $M\sim\ Binom(n-1, F(T))$ is the random variable equal to the number of players other than i who also succeed. Letting $p = F(T)$, we have

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

where the second line follows by expanding the binomial coefficient and manipulating the expression, the third line follows by changing variables $j = k + 1$, and the last line follows by collecting terms. Therefore,

$$R = \frac{1-(1-F(T))^n}{nF(T)}.$$

Now, taking as given the strategy profile of the other players, the net expected payoff of player i from spending effort for duration T is given by

$$F(T)R - Tc.$$

Note that because F is concave, this is a concave objective and therefore, the best-response for player i is the duration T' given by

$$f(T')R = c;$$

in other words

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

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

$$\frac{1-\left(1-F(T^{\mathrm{EGA}})\right)^n}{nF(T^{\mathrm{EGA}})} = \frac{c}{f(T^{\mathrm{EGA}})}.$$

$\square$

A.3 Proof of Proposition 2.

Towards proving Proposition 2, we first establish a lemma showing 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 3. *In an equilibrium of a contest in which each player i spends effort continuously through an interval $[0, T_i]$, the reward functions $R_{i,t}$ must satisfy the following "budget constraint"*

$$\sum_i \int_0^{T_i} f(t) R_{i,t} dt \le 1 - \prod_i (1 - F(T_i)). \tag{BC}$$

Proof of Lemma 3. Note that

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

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

Using Lemmas 1 and 3, we consider the following relaxation of (1):

$$\max_{\{T_i\}, \{R_{i,t}\}} \sum_{i=1}^{n} T_i \tag{6}$$
$$\text{s.t. (IC) and (BC).}$$

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 Proposition 1 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 (6) 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 (6). We will show that this solution achieves a smaller objective than the egalitarian contest characterized in Proposition 1, that is, $\sum_i T_i < nT^{\text{EGA}}$. Because the egalitarian contest is feasible for the original problem (1), it will immediately follow that this contest must be optimal.

Define the function Z_i^1 for each i as follows

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

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

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

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

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

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

$$\begin{aligned} Z_i(t) = \lim_{k\to\infty} Z_i^k(t) &= \lim_{k\to\infty} \frac{1}{f(t)} \left[c - \int_t^{T_i} f'(s) Z_i^{k-1}(s) ds \right] \\ &= \frac{1}{f(t)} \left[c - \int_t^{T_i} f'(s) Z_i(s) ds \right] \end{aligned} \tag{7}$$

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:

$$f(t)\tilde{R}_{i,t} + \int_t^{T_i} f'(s) \tilde{R}_{i,s} ds - c = 0. \tag{8}$$

Differentiating both sides of (8) reveals that $\tilde{R}_{i,t}$ is a constant function $\tilde{R}_{i,t} \equiv c/f(T_i)$.

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

$$c\sum_i \frac{F(T_i)}{f(T_i)} - \left[1 - \prod_i (1 - F(T_i))\right] \leq 0. \tag{9}$$

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 (9) 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 (9) 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

$$c\sum_i e^{\lambda T_i} + \lambda e^{-\lambda \sum_i T_i}.$$

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

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

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

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

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 (9) 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 (6) 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 (6) can be rewritten as

$$\max_{T} \left\{ nT \text{ s.t. } cn\frac{F(T)}{f(T)} \leq 1 - [1 - F(T)]^n \right\}, \tag{10}$$

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 (10).

First notice that the constraint in (10) 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 (10) 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 (6), and its objective equals nT^{EGA}. Since this is a relaxation of the original problem, (1), the objective of the original problem is bounded above by nT^{EGA}. By Proposition 1, the egalitarian contest with deadline $T \geq T^{\text{EGA}}$ 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 (6). Therefore, this egalitarian contest is an optimal no-feedback contest. □

B Increasing Hazard Rate

In this section, we assume that each agent's hazard rate of success increases in both past and current efforts. To be specific, let $F(t)$ denote the probability that an agent succeeds on or before date t if he works continuously until that date, and suppose that the hazard rate $\lambda_t = F'(t)/[1 - F(t)]$ exists and is weakly increasing. Thus, if an agent has expended s units of effort by date t, his hazard rate at t is λ_s.

The following proposition characterizes the optimal contest under the assumptions that the hazard rate $\lambda_t \in (c, nc)$ for all t and is differentiable almost everywhere.[21] This contest is similar to the one with constant hazard rate characterized in Proposition 3 except that the length of each cycle is stochastic. In particular, it comprises the egalitarian prize allocation rule, the feedback policy $\mathcal{M}^{pronto}$, and a cyclical termination rule which operates as follows: A cycle which *starts* at t, *ends* at a random date $T > t$ distributed according to the CDF $G^t(T) = 1 - e^{-\int_t^T \gamma_s^t ds}$, where

$$\gamma_s^t := \frac{c\dot{\lambda}_s}{\lambda_s^2}\left[\frac{1 - e^{-n\int_t^s \lambda_v dv}}{n\left(1 - e^{-\int_t^s \lambda_v dv}\right)} - \frac{c}{\lambda_s}\right]^{-1}. \tag{11}$$

[22] To elaborate, at date 0, the principal privately draws $T_1 \sim G^0(\cdot)$. The first cycle ends at T_1 and agents are informed so. During each cycle, agents know only the distribution of the cycle's duration and that the current cycle is still ongoing. If at least one agent has succeeded by that date, the contest ends and the prize is awarded according to the egalitarian rule. Otherwise, a new cycle begins, which ends at random date $T_2 \sim G^{T_1}(\cdot)$. The contest proceeds in this cyclical manner until the end of a cycle in which at least one agent has succeeded. Let τ^{**} denote the termination rule defined by this algorithm.

Proposition 4. *Assume $\lambda_t \in (c, nc)$ for all t and its derivative exists almost everywhere. The contest with egalitarian prize, the cyclical termination rule τ^{**}, and the feedback policy $\mathcal{M}^{pronto}$ is optimal. In this contest, at least one agent succeeds, and hence the prize is awarded with probability one. Moreover, in equilibrium, each agent obtains 0 expected utility and the principal's profit is* $1/c$.

To explain the logic of this design, notice first that Lemma 2 remains valid: a contest which awards the entire prize and concedes zero rents to the agents is guaranteed to be optimal. Given that the feedback policy is $\mathcal{M}^{pronto}$, it suffices to show that each agent's expected reward conditional on success, $R_{i,t} = c/\lambda_t$ until he succeeds. Then by the same argument as in Section 4, there exists an equilibrium in which agents

[21]The first assumption is analogous to the one imposed in Proposition 3, while the second is a technical one.

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

work continuously until they succeed and earn zero rents.[23]

When the hazard rate is increasing, to extract all rents, incentives must be front-loaded so that an earlier success is rewarded more dearly than one that occurs later. Proposition 4 shows that this can be achieved using an egalitarian prize and a stochastic cyclical termination rule. Intuitively, when the length of each cycle is random, an agent's expected reward from succeeding early in a cycle is larger than if he succeeds later, because probabilistically, he will have to share the prize with fewer of his rivals. By choosing the distribution of each cycle's length, it is possible to fine-tune $R_{i,t}$ so that it is always equal to c/λ_t as desired.

Proof of Proposition 4.

The proof is organized as follows. First, we show that each agent's expected reward conditional on succeeding at date t, $R_{i,t} = c/\lambda_t$. Then we will argue that there exists an equilibrium in which all agents work continuously until they succeed or the contest ends. Finally, we will argue that both conditions of Lemma 2 are met, and hence this contest is optimal.

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

$$\frac{1 - e^{-n\int_t^T \lambda_v dv}}{n\left(1 - e^{-\int_t^T \lambda_v dv}\right)} = \frac{c}{\lambda_T}. \tag{12}$$

The left-hand side is strictly decreasing in T, it converges to 1 as $T \to 0$, and to $1/n$ as $T \to \infty$. Meanwhile, $\lambda_T \in (c, nc)$ by assumption, and so the right-hand side takes values strictly between $1/n$ and 1. Since both sides are continuous in T, by the intermediate value theorem, there exists a smallest T such that (12) is satisfied. Moreover, because the left-hand side of (12) is strictly larger for $T \simeq 0$, this is also true for all $T < \overline{T}_t$. Therefore, for every t, we have $\gamma_s^t \geq 0$ for all $s \in [0, \overline{T}_t)$, and $\lim_{s \to \overline{T}_t} \gamma_s^t = \infty$; i.e., a cycle which starts at t ends with certainty by $\overline{T}_t$.

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

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

succeeding at s is

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

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

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

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

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

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

whenever $R_{i,s}=c/\lambda_s$. The first equality follows from the Leibniz integral rule, and the second equality follows by substituting γ_s^t defined in (11) and $R_{i,s}=c/\lambda_s$. Therefore, $\lambda_s R_{i,s}=c$ at $s=\overline{T}_t$, and moving backwards in time, $d\left(\lambda_s R_{i,s}\right)/ds=0$, implying that $\lambda_s R_{i,s}=c$ for all $s\in[t,\overline{T}_t]$.

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

Finally, because the contest does not end until at least one agent succeeds, the prize is awarded with probability one, which implies that this contest satisfies both conditions of Lemma 2, and is therefore optimal. □

C Bitcoin Institutional Details

Here we briefly describe some institutional details surrounding Bitcoin, and in particular, its proof-of-work protocol. For a more detailed description, see for instance Nakamoto (2008) and Bohme et al. (2015).

A key challenge faced by digital currencies designed to operate without a trusted intermediary (such as a central bank), including Bitcoin, is preventing users from spending the same token more than once. To overcome it, Bitcoin and other cryptocurrencies rely on a timestamp server and a proof-of-work system, which operate as follows:

i. When a transaction is requested by a Bitcoin holder, a timestamped record is added to a list of transactions waiting to be included in a public ledger (called blockchain). This timestamped record includes the entire ownership history of the specific Bitcoins that the owner intends to use in their current transaction, thus forming a chain. A subset of this list in waiting is called a block.

ii. Blocks are added to the blockchain by so-called miners, who are parties all over the world acting on their own accord. Before a miner is allowed to add a new block to the blockchain, they must first solve a cryptographic puzzle called "proof of work," whose intention is to make modifications to the blockchain sufficiently costly while also making it essentially impossible for the same person to make back-to-back changes to it.

   Proof of work involves finding a string of characters that solves a puzzle. This occurs when a candidate string, once appended to a string that identifies all transactions in the block plus a random string, and imputed to a hash function, leads this function to deliver an output that begins with a predetermined number of zeros. The only (known) way to solve this puzzle is by trial and error, and by design, the probability that any one trial succeeds is constant.

iii. Once a miner solves the puzzle, the block is transmitted to all users on the network. Provided none of the transactions in the block are double-spent, users proceed to implicitly accept it by working to add the next block to the chain. Adding to a given chain makes that chain grow through what is called a leaf (a branch of the chain). In principle, there can be multiple parallel leafs, but the longest chain is implicitly considered to be the valid one.

The miner who solved the puzzle is awarded newly issued Bitcoin. This "block reward," currently 6.25 Bitcoins, is halved approximately every four years as a way to limit the inflation of the currency.

How exactly does this decentralized system prevent double spending? By design, an attacker cannot simply spend or take Bitcoin that never belonged to them. Suppose, however, that they try to spend a given Bitcoin more than once. To do so, they would need to go back in the chain and take back Bitcoin that they have already spent, thus undoing that earlier transaction. Specifically, they would need to go back in the chain to the point just before the block that contains the earlier transaction, and then add blocks starting from that point onward. In doing so, the attacker would attempt to create a new chain parallel to the legitimate one.

By convention, however, this new chain would only be accepted by other users on the network, thus becoming the new legitimate one, if it is the longest one, which means that the attacker, who starts at least one block behind, would have to outpace all the other miners in solving new puzzles. Provided the attacker has less than half of the network's total computing power, the probability of this event is small, and in fact becomes vanishingly small as the number of blocks needed to catch up increases. Thus, by making the proof of work costly, the platform ensures that attacks are financially untenable.

Observe that the race to add the next block to the ledger is a type of contest where miners expend effort in the form of attempts at a solution. The first miner to succeed wins the block reward, and at that moment all contestants are informed and that specific contest ends. Because the purpose of the contest is to make modifications to the blockchain costly, and at the same time issuing new Bitcoin is inflationary, the contest designer arguably seeks to minimize the prize needed for a given total effort level, or, equivalently, maximize total effort for a given prize.