

Designing Dynamic Contests

Kostas Bimpikis

Graduate School of Business, Stanford University, kostasb@stanford.edu

Shayan Ehsani

Department of Management Science and Engineering, Stanford University, shayane@stanford.edu

Mohamed Mostagir

Ross School of Business, University of Michigan, mosta@umich.edu

Participants race towards completing an innovation project and learn about its feasibility from their own efforts and their competitors' gradual progress. Information about the status of competition can alleviate some of the uncertainty inherent in the contest, but it can also adversely affect effort provision from the laggards. This paper explores the problem of designing the award structure of a contest and its information disclosure policy in a dynamic framework and provides a number of guidelines for maximizing the designer's expected payoff. In particular, we show that the probability of obtaining the innovation as well as the time it takes to complete the project are largely affected by *when* and *what* information the designer chooses to disclose. Furthermore, we establish that intermediate awards may be used by the designer to appropriately disseminate information about the status of competition. Interestingly, our proposed design matches several features observed in real-world innovation contests.

Key words: Contests, learning, dynamic competition, open innovation, information.

1. Introduction

Innovation contests are a tool that firms and institutions use to outsource innovation to the crowd. An open call is placed for an innovation project that participants compete to finish, and the winners, if any, are awarded a prize.¹ Recent successful examples include The NetFlix Prize and the Heritage Prize², and a growing number of ventures like Innocentive, TopCoder, and Kaggle provide online platforms to connect innovation seekers with potential innovators.

The objective of the contest designer is to maximize the probability of reaching the innovation goal while minimizing the time it takes to complete the project. Obviously, the success of a contest depends crucially on the pool of participants and the amount of effort they decide to provide, and a growing literature considers the question of how to best design a contest. The present paper studies a model with the following three key features. First, in our model, an agent's progress towards the goal is not a deterministic function of effort. As is typically the case in real-world settings, progress

is positively but imperfectly correlated with effort. Second and quite importantly, it is possible that the innovation in question is not attainable, either because the goal is actually infeasible or because it requires too much effort and resources that it makes no economic sense to pursue. We model such a scenario by means of an underlying state that captures whether the innovation is attainable. These two features imply that an agent’s lack of progress may be attributed to either an undesirable underlying state (the innovation is not attainable) or simply to the fact that the agent was unlucky in how her effort was stochastically mapped to progress. Third, we consider a *dynamic* framework to study how competition between agents evolves over time. Critically, we incorporate the fact that they learn from each other’s partial progress about the feasibility of the innovation project. In particular, our modeling setup includes well-defined intermediate milestones that constitute partial progress towards the end goal. Such milestones are usually featured in real-world innovation contests, including the ones we use as motivating examples.

The discussion above implies that news about a participant’s progress has the following interesting dual role: it makes agents more optimistic about the state of the world, as the goal is more likely to be attainable and thus agents have a higher incentive to exert costly effort. We call this the *encouragement effect*.³ At the same time, such news imply that one of the participants has a lead, which might negatively affect effort provision from the remaining agents as the likelihood of them beating the leader and winning the prize becomes slimmer. We refer to this as the *competition effect*. These two effects interact with one another in subtle ways over the duration of the contest, and understanding this interaction is of first-order importance for contest design.

The primary contribution of this paper is twofold. First, while some of the features described above — uncertainty regarding the feasibility of the end goal, stochastic mapping between effort and progress, and intermediate milestones — appear in previous literature, to the best of our knowledge, this framework is the first that explicitly combines all three into a single model. This allows us to focus on the information disclosure policy of the contest designer and show how this policy depends on whether the competition or the encouragement effect dominates. In particular, we consider the question of *whether* and *when* should the contest designer disclose information regarding the competitors’ (partial) progress with the goal of maximizing her expected payoff. Interestingly, we illustrate the benefits of non-trivial information disclosure policies, where the designer withholds information from the agents and only releases it after a certain amount of time has elapsed. Such designs highlight the active role that information may play in incentivizing agents to participate in the contest. As we further elaborate in the literature review, much of the extensive prior work on innovation contests studies static single-shot models that feature no uncertainty regarding the goal (and thus no learning).

Second, we identify the role of intermediate awards as a way for the designer to implement the desired information disclosure policy — the policy that maximizes the effort provision of the agents and consequently the chances of innovation taking place. Intermediate awards are very common in innovation contests (the aforementioned NetFlix and Heritage prizes are examples of contests that have employed intermediate awards), but the exact role they play as information revelation devices has not yet been studied. We show how these awards may serve to both extract private information (from those agents who have made some progress) as well as disseminate this information to the rest of the competition through the public announcement (or not) of giving out an award.

A simple illustration of the main ideas in the paper is the following. Consider an innovation contest that consists of well-defined milestones. For example, the goal of the Netflix prize was to achieve an improvement of 10% over the company’s proprietary algorithm, with a first progress prize set at 1% improvement. In this example, reaching the milestone of 1% improvement constitutes *partial progress* towards the goal, and we assume that the agents and the designer are able to verifiably communicate this. Assume for now that the innovation is attainable with certainty given enough effort, and that agents are fully aware of that. The lack of progress towards the goal is then solely a result of the stochastic return on effort. When no information is disclosed about the agents’ progress, they become progressively more pessimistic about the prospect of them winning, as they believe that someone must have made progress and that they are now lagging behind in the race towards the end goal. This might lead them to abandon the contest, thus decreasing the aggregate level of effort and consequently increasing the time to complete the innovation project.

In contrast, when there is uncertainty about the feasibility of the end goal, agents that have made little or no progress towards the goal become pessimistic about whether it is even possible to complete the contest. If this persists, an agent may drop out of the competition as she believes that it is not worth putting the effort for what is likely an unattainable goal, reducing aggregate experimentation in the process and decreasing the chances of reaching a possibly feasible innovation.

This discussion highlights the complex role that information about the agents’ progress play in this environment. In the first scenario, the competition effect is dominant since there is no uncertainty about the attainability of the end goal. Hence, disclosing that one of the participants is ahead may deter future effort provision as it implies that the probability of winning is lower for the laggards. In the second case, the encouragement effect dominates, and an agent’s progress can be perceived as good news, since it reduces the uncertainty over the feasibility of the end goal.

The information disclosure policy is only one of the levers that the designer has at her disposal to affect the agents’ effort provision decisions. Another is the compensation scheme that in an innovation contest takes the form of an award structure. In a setting with potentially multiple milestones, a design may involve compensating agents for reaching a milestone or having them

compete for a single grand prize given out for completing the entire contest. Our analysis sheds light on the interplay between information disclosure and the contest’s award structure by comparing different mechanisms in terms of their expected payoff for the designer. This brings the information disclosure policy to the forefront as we show that the probability of obtaining the innovation as well as the time it takes to complete the project are largely affected by *when* and *what* information the designer chooses to disclose.

Related Literature: A growing literature explores different aspects of innovation contests. For example, [Taylor \(1995\)](#), in an influential early work, considers a tournament in which agents decide whether to conduct (costly) research and obtain an innovation of value drawn from a known distribution at each of T time periods after which an award is given out to the agent with the highest draw. [Taylor \(1995\)](#) finds that a policy of free and open entry may give rise to low levels of effort at equilibrium; thus, it may be optimal for the sponsor to restrict participation by imposing an entry fee. Relatedly, [Moldovanu and Sela \(2001\)](#) consider the case when the agents’ cost of effort is their private information and show that when the cost is linear or concave in effort, allocating the entire prize sum to the winner is optimal whereas when it is convex several prizes may be optimal. [Che and Gale \(2003\)](#) find that for a set of procurement settings it is optimal to restrict the number of competitors to two and, in the case that the two competitors are asymmetric, handicap the most efficient one. [Moldovanu and Sela \(2006\)](#) explore the performance of contest architectures that may involve splitting the participants among several sub-contests whose winners compete against each other. [Siegel \(2009\)](#) provides a general framework to study such static all-pay contests that allows for features such as different production technologies and attitudes toward risk. [Terwiesch and Xu \(2008\)](#), [Ales et al. \(2016a\)](#), and [Ales et al. \(2016b\)](#) explore static contests in which there is uncertainty regarding the value of an agent’s contribution and explore the effect of the award structure and the number of competitors on performance. Finally, [Boudreau et al. \(2011\)](#) examine related questions empirically using data from software contests.

Unlike the papers mentioned above, a central feature in our model is the fact that there is uncertainty with respect to the attainability of the end goal. In addition, agents *dynamically* adjust their effort provision levels over time responding to the information they receive regarding the status of the competition and the state of the world, i.e., whether the contest can be completed. Early papers that consider the dynamics of costly effort provision in the presence of uncertainty are [Choi \(1991\)](#) and [Malueg and Tsutsui \(1997\)](#). These papers study R&D races and assume that firms can observe each other’s experimentation outcomes, thus abstracting away from using information about relative progress as an incentive mechanism.⁴ In addition, the “award”, which in this case is the value of the innovation in question, is fixed. In contrast, in our setting each agent’s progress,

i.e., the outcomes of her experimentation process, is her private information and a third party, the designer, determines the contest's award structure and information disclosure policy.

There is also a stream of papers (Bolton and Harris (1999), Keller et al. (2005), Keller and Rady (2010), Bonatti and Horner (2011), and Klein and Rady (2011)) that study the dynamics of experimentation within a team of agents that work towards completing a project. Strategic interactions are driven by the fact that experimentation outcomes are observable and information is a public good and the focus is on how the agents' incentives to free-ride affect the level of aggregate experimentation. Bimpikis and Drakopoulos (2016) also consider experimentation incentives within a team and show that having agents work independently and then combine their efforts increases aggregate welfare. Although our model builds on the exponential bandits framework that was introduced in Keller et al. (2005), the setup and focus are considerably different from the strategic experimentation literature. In particular, agents compete with one another for a set of awards that are set ex-ante by the designer. Furthermore, we allow for *imperfect monitoring* of the agents' progress (experimentation outcomes). This, along with the fact that agents dynamically learn about the attainability of the end goal and the status of competition, significantly complicate the analysis as not only do agents form beliefs about whether they can complete the contest but also about their progress relative to their competitors. The latter is not an issue in the strategic experimentation literature since experimentation outcomes are typically assumed to be perfectly observable.

Our paper is also related to the literature on dynamic competition. For example, Harris and Vickers (1987) show that in a one-dimensional model of a race between two competitors, the leader provides more effort than the follower and her effort increases as the gap between the competitors decreases. On the other hand, Hörner (2004) shows that firms invest most in effort provision when they are far ahead in an effort to secure a durable leadership or when they are lagging sufficiently behind to prevent their rival to outstrip them. Furthermore, Moscarini and Smith (2007) consider a two-player dynamic contest with perfect monitoring where the focus is on the design of a *scoring function* in which the leader is appropriately "taxed" whereas the laggard is "subsidized".⁵ Unlike these papers we allow the contest designer to choose what information and when to disclose it, thus putting more emphasis on how the designer can incentivize agents to take a certain set of actions by controlling the information they have access to.

Another paper related to our work is that of Lang et al. (2014) who study a two-player continuous time contest in which there is no uncertainty about the underlying environment but agents exert costly effort to complete as many milestones as they can before a predetermined deadline. They characterize equilibrium behavior and, because of the lack of uncertainty and dynamic learning, they are able to establish a close relation with the outcomes of (static) all-pay auctions thus linking their framework with prior work on static contests (e.g., Siegel (2009)).

Finally, the contemporaneous work of Halac et al. (2016) studies contests that end after the occurrence of a single breakthrough. They do not incorporate the possibility of partial progress and therefore they abstract away from the fact that agents may learn from the progress of others, i.e., the encouragement effect is absent in their model. Our framework shares some features with theirs, particularly the uncertainty regarding the attainability of the end goal, but our main focus is on exploring the interplay between the contest’s award structure and the information disclosure policy that it implies in relation to the encouragement and competition effects. This only becomes relevant in the presence of *partial progress towards the end goal* and *discounting*, which are two features that are unique to our model and that we believe capture realistic aspects of contests. As Halac et al. (2016) consider a contest with no intermediate milestones and assume that there is no discounting, the time it takes to complete the contest is immaterial for their analysis. In addition, incorporating intermediate milestones enables us to study information disclosure policies that involve a (stochastic) delay between (partial) progress and the designer’s announcements. On the other hand, they consider multiple competitors and allow for strategies in which the designer broadcasts a message at time t only if at least k competitors have completed the project by that time. We focus on two competitors and as a result we do not allow for such strategies.

2. Model

Our benchmark model is an innovation contest with two sequential stages, A and B , and two competitors, 1 and 2. Innovation happens if an agent successfully completes Stage A and then Stage B . Stage A is associated with a binary state θ_A that describes whether that stage can be completed ($\theta_A = 1$) or not ($\theta_A = 0$). If $\theta_A = 0$, then Stage A is not feasible (and, consequently, innovation is not possible). If $\theta_A = 1$, then the breakthrough to complete Stage A is feasible and has an arrival rate that is described by a Poisson process with parameter λ .⁶ Once Stage A is complete, there is no residual uncertainty in the model and the arrival rate of the breakthrough to complete Stage B is equal to μ . We assume that agents have a common prior on θ_A and we denote that prior by $p_A = \mathbb{P}(\theta_A = 1)$.




Figure 1 An innovation contest with two stages, A and B .

Agents choose their effort levels continuously over time. Agent $i \in \{1, 2\}$ chooses effort $x_{i,t} \in [0, 1]$ at time t and incurs an instantaneous cost of effort equal to $cx_{i,t}$ for a constant $c > 0$. An agent in Stage A who puts effort x_t at time t obtains a breakthrough, i.e., completes Stage A , with

instantaneous probability $\theta_A \lambda x_t$. We assume that an agent’s effort provision level is not observable by her competitor or by the designer. Agent i is endowed with an information set — described later in this section — that summarizes her information about the contest at time t . Moreover, we assume that although an agent’s effort level is her private information, progress, i.e., completing Stages A or B , is observable by the agent and the designer (but not the agent’s competitor). We relax this assumption towards the end of Section 4 and discuss how the designer may incentivize agents to share their progress. Finally, progress can be verifiably communicated, i.e., we abstract away from “cheap-talk” communication between the designer and the two competitors.⁷

The designer commits to the contest’s award structure and information disclosure policy. In particular, R_A and R_B denote the awards for completing Stages A and B of the contest respectively. Throughout the paper, we assume that the designer announces the completion of Stage B as soon as it happens and gives out award R_B to the agent that completes it. Thus, agents have no incentive to continue exerting effort after such an announcement and the game essentially ends.

Our main focus is on studying how different disclosure policies for the agents’ *partial progress*, i.e., completing Stage A , may affect their effort provision and consequently the designer’s expected payoff. Specifically, we consider a class of policies that — conditional on Stage A being completed at time τ — specify the cumulative probability that the designer publicly announces partial progress by time $t \geq \tau$. In other words, the probability that the designer has disclosed the agents’ partial progress by time t is given as follows as a function of the history up to t :

$$\Phi_t : \{\tau_{1,t}^A, \tau_{2,t}^A\} \rightarrow [0, 1],$$

where $\tau_{i,t}^A \in [0, t] \cup \{\emptyset\}$ denotes the time at which agent i has completed Stage A if she has done so by time t (here, \emptyset denotes that the agent has not completed the stage before time t). We emphasize that the designer’s announcements are public, i.e., we abstract away from asymmetric information disclosure policies that may feature different messages being communicated to the two agents.

Our analysis proceeds by first considering the full and no information disclosure benchmarks in Section 3 (corresponding to $\Phi_t = 1$ if $\tau_{i,t}^A \leq t$ for some i and $\Phi_t = 0$, respectively), whereas Section 4 explores designs in which progress is disclosed with some delay. Finally, in terms of the award R_A , we assume that it is given out to the agent that first completes Stage A or split equally between the two competitors if they both complete the stage before the designer discloses any information about their respective progress.

Payoffs are discounted at a common rate r for both the designer and the agents. Throughout the paper we assume that the expected budget allocated to the contests’ award(s) is kept fixed and we compare different information disclosure policies in terms of the payoff they generate for the designer (thus, focusing our analysis on information disclosure).⁸

On the other hand, an agent's strategy is a mapping from her information set at time t to an effort provision level $x_{i,t} \in [0, 1]$. Agent i 's information set at t includes the calendar time, the contest's award structure and information disclosure policy, the agent's effort levels up to time t , i.e., $\{x_{i,\tau}\}_{0 \leq \tau < t}$, whether the agent has already completed Stage A , i.e., the indicator variable $I_{i,t}^A$, and, last, whether the designer has announced that she or her competitor have already completed Stage A .

Finally, agents hold a set of beliefs $\{p_{i,t}, q_{i,t}\}$ that evolve over time, where $p_{i,t}$ denotes agent i 's belief about the feasibility of Stage A and $q_{i,t}$ denotes her belief about whether her competitor has already completed Stage A conditional on the stage being feasible, e.g., $q_{i,t} = 1$ implies that at time t , agent i believes with certainty that her competitor is in Stage B . Note that the agents' beliefs co-evolve through their interaction with the designer's information disclosure policy, since the only way to obtain information about a competitor's progress is through the designer's announcements.

3. Real-time Leaderboard and Grand Prize

We begin our exposition by considering two intuitive contest designs. The first, which we call *real-time leaderboard*, features full information disclosure from the designer, i.e., the agents' progress is publicly disclosed on an online leaderboard by the designer as soon as it happens. The second, which we call (single) *grand prize*, is such that the designer only discloses the completion of the entire contest, i.e., Stage B , and thus an agent determines her effort provision levels over time solely based on observing the outcomes of her own experimentation as well as the beliefs she forms about her competitor's effort levels and progress.

To allow for tractable analysis, we assume that conditional on the innovation being feasible, Stage B takes more time to complete in expectation than Stage A (some of our results are actually stated for the limit $\mu \rightarrow 0$). This assumption together with the assumption that $p_A < 1$ provide a good approximation of the dynamics at the early stages of a contest, when there is both a significant amount of uncertainty as well as plenty of time before a competitor reaches the end goal.

Given that the designer discloses the completion of Stage B as soon as it happens in both designs, the information disclosure policy centers around partial progress, i.e., the completion of Stage A . Below, we describe the agents' belief update process under the two extremes of information disclosure corresponding to the real-time leaderboard and grand prize designs.

- *Full disclosure:* Let $p_{i,t}^{\mathcal{F}}$ denote agent i 's belief about the feasibility of Stage A in the absence of any progress up to time t when agents observe each others' outcomes (superscript \mathcal{F} refers to full disclosure). Then, Lemma 1 provides a characterization for the change in beliefs over period $[t, t + dt)$.

LEMMA 1. *The change in agent i 's beliefs about the feasibility of Stage A in the absence of progress over time period $[t, t + dt)$ is given as:*

$$dp_{i,t}^{\mathcal{F}} = -p_{i,t}^{\mathcal{F}}(1 - p_{i,t}^{\mathcal{F}})(x_{i,t} + x_{-i,t})\lambda dt, \quad (1)$$

where $x_{i,t}$ denotes the effort provision level by agent i at time t and $x_{-i,t}$ denotes the effort level that agent i believes her competitor is exerting at time t .

- *No information disclosure:* In the absence of partial progress (and assuming that her competitor has not claimed the final award), the law of motion of agent i 's belief at t is given by:

$$dp_{i,t}^{\mathcal{N}} = -p_{i,t}^{\mathcal{N}}(1 - p_{i,t}^{\mathcal{N}})(\lambda x_{i,t} + \mu q_{i,t}^{\mathcal{N}}) dt, \quad (2)$$

where $q_{i,t}^{\mathcal{N}}$ denotes the belief agent i assigns to the event that her competitor has already completed Stage A conditional on the stage being feasible (superscript \mathcal{N} refers to no disclosure). Note that the law of motion for $q_{i,t}^{\mathcal{N}}$ takes the following form:

$$dq_{i,t}^{\mathcal{N}} = (1 - q_{i,t}^{\mathcal{N}})(\lambda x_{-i,t} - q_{i,t}^{\mathcal{N}}\mu) dt, \quad (3)$$

where as above $x_{-i,t}$ denotes the effort level that agent i believes her competitor is exerting at t in the absence of progress (Expressions (2) and (3) can be derived similar to Lemma 1).

Intuitively, full information disclosure allows for the fast dissemination of the good news of an agent's progress (completion of Stage A), since it resolves the uncertainty about the feasibility of the end goal and therefore instantly affects the competitors' future effort provision. On the flip side, the absence of progress makes agents pessimistic at a faster rate than when information is not public. Indeed, by comparing Expressions (1) and (2), one can easily deduce that agents' beliefs move downward faster under full disclosure. This comparison clearly highlights one of the designer's main tradeoffs: on the one hand, sharing progress between competitors allows for the timely dissemination of good news and induces agents to exert effort. On the other hand, the absence of partial progress early in the process can make agents pessimistic about the feasibility of the underlying project and adversely affect their effort provision.

Real-Time Leaderboard: The first contest design we study involves full information disclosure, i.e., the designer continuously discloses information about the agents' progress. In addition, the awards for completing Stages A and B are given out to the agent that completes them first.

Consider the subgame that results when one of the agents, the leader, completes Stage A. The leader's *optimal* effort provision takes a very simple form for $t \geq \tau_A$, where τ_A is the random time at which Stage A was completed. In particular, if we index the leader by i , we have

$$x_{i,t}^* = \begin{cases} 1 & \text{if } R_B \geq c/\mu \\ 0 & \text{otherwise} \end{cases}, \quad \text{for } t \geq \tau_A.$$

Thus, the designer should set R_B to be at least as high as c/μ in order to ensure that the contest is going to be completed (recall that Stage B can be completed with probability one, so as soon as an agent breaks through to that stage, she will continue putting an effort until the contest is complete assuming that the value of the award is high enough to cover her cost of effort).

Similarly, the laggard continues putting effort in the contest if her expected payoff is higher than the instantaneous cost of effort. Note that since there is no residual uncertainty after the leader has completed Stage A , the laggard would complete Stage A with probability λdt if she puts effort. In turn, her payoff from that time onwards would be equal to

$$\int_0^\infty \mu e^{-(2\mu+r)t} R_B dt = \frac{\mu}{2\mu+r} R_B,$$

since each one of the agents would complete Stage B at rate μ (recall that r denotes the discount rate). Thus, the laggard would find it optimal to continue putting effort if

$$\lambda \frac{\mu R_B - c}{2\mu+r} \geq c, \text{ or equivalently } R_B \geq \frac{c}{\mu} \left(1 + \frac{2\mu+r}{\lambda} \right).$$

In other words, upon completion of Stage A , both the leader and the laggard remain in the contest and put full effort until one of them completes Stage B if the final award R_B is at least

$$R_B \geq R_B^{\min} \equiv \frac{c}{\mu} \left(1 + \frac{2\mu+r}{\lambda} \right). \quad (4)$$

If, on the other hand, $\frac{c}{\mu} \leq R_B < R_B^{\min}$, the laggard drops out of the contest while the leader continues putting full effort until the end. We let $\Pi(k, \ell, R_B)$ denote the expected payoff of agent i when she is in Stage k , her competitor is in Stage ℓ , and the final award is equal to R_B . Then, it is straightforward to obtain the following:⁹

$$\Pi(A, B, R_B) = \begin{cases} \frac{\lambda}{\lambda+\mu+r} \frac{\mu}{2\mu+r} R_B - \frac{\lambda+2\mu+r}{(\lambda+\mu+r)(2\mu+r)} c & \text{if } R_B \geq R_B^{\min} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

$$\Pi(B, A, R_B) = \begin{cases} \left(\frac{\lambda}{\lambda+\mu+r} \frac{\mu}{2\mu+r} + \frac{\mu}{\lambda+\mu+r} \right) R_B - \frac{\lambda+2\mu+r}{(\lambda+\mu+r)(2\mu+r)} c & \text{if } R_B \geq R_B^{\min} \\ \frac{\mu R_B - c}{\mu+r} & \text{otherwise} \end{cases} \quad (6)$$

Furthermore, given that the designer would only have the incentive to organize the contest if her ex-ante expected payoff was positive, Assumption 1 below states that the utility she obtains from the innovation is sufficiently high, i.e., higher than the size of award R_B^{\min} .

ASSUMPTION 1. *The utility the designer obtains from the innovation is strictly higher than R_B^{\min} .*

Following the discussion above, agent i 's optimization problem can be written as follows:

$$\max_{\{x_{i,\tau}\}_{\tau \geq 0}} \int_0^\infty [x_{i,\tau} (p_{i,\tau}^{\mathcal{F}} \lambda (R_A + \Pi(B, A, R_B)) - c) + x_{j,\tau} p_{i,\tau}^{\mathcal{F}} \lambda \Pi(A, B, R_B)] e^{-\int_0^\tau (p_{i,s}^{\mathcal{F}} \lambda (x_{1,s} + x_{2,s}) + r) ds} d\tau,$$

where the first term between the brackets is equal to the agent's (expected) instantaneous payoff from exerting effort $x_{i,\tau}$ at time τ . On the other hand, the second term captures the agent's expected payoff if her competitor completes Stage A at time τ (which occurs with probability $x_{j,\tau} p_{i,\tau}^F \lambda d\tau$).

We are interested in characterizing the unique symmetric equilibrium in Markovian strategies in this setting. Proposition 1 below states that agents follow a cutoff experimentation policy in Stage A and the aggregate amount of experimentation increases with the size of the intermediate award. The proposition considers awards for completing Stage B that can take one of two values, i.e., $R_B = c/\mu$ or $R_B = R_B^{\min}$, since it is straightforward to establish that given a fixed budget it is optimal for the designer to consider only these two values and allocate her remaining budget to R_A (for a formal argument refer to Lemma EC.1 in the electronic companion).

PROPOSITION 1. *Consider a contest design in which progress is publicly observable and the awards for completing Stages A and B are equal to R_A and R_B respectively. Then, there exists a unique symmetric equilibrium in which agents experiment as follows:*

(i) *Agents follow a cutoff experimentation policy in Stage A , i.e., in the absence of progress they quit the contest at time $t_{\mathcal{F}}$ given below*

$$x_{i,t}^* = \begin{cases} 1 & \text{for } t \leq t_{\mathcal{F}} \equiv \frac{1}{2\lambda} \ln \left(\frac{1-p_{\mathcal{F}}}{p_{\mathcal{F}}} \cdot \frac{p_A}{1-p_A} \right) \\ 0 & \text{otherwise} \end{cases}.$$

The cutoff belief $p_{\mathcal{F}}$ is given as follows

$$p_{\mathcal{F}} = \begin{cases} \frac{c}{\lambda \left(R_A + \frac{\mu R_B - c}{\mu + r} \right)} & \text{if } R_B = \frac{c}{\mu} \\ \frac{c}{\lambda \left(R_A + \left(\frac{\lambda}{\lambda + \mu + r} \frac{\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right) R_B - \frac{2\mu + \lambda + r}{(\lambda + \mu + r)(2\mu + r)} c \right)} & \text{if } R_B = R_B^{\min} \end{cases}. \quad (7)$$

(ii) *If Stage A has been completed, experimentation continues as follows*

(a) *If $R_B = R_B^{\min}$: Both agents experiment with rate one until the end of the contest.*

(b) *If $c/\mu \leq R_B < R_B^{\min}$: The laggard drops out of the contest whereas the leader experiments with rate one until the end.*

Before concluding the discussion on full information disclosure, note that Proposition 1 clearly illustrates the tradeoff the designer faces when she decides how to split her budget between awards R_A and R_B . In particular, setting R_B equal to R_B^{\min} provides an incentive for the laggard to stay active in the contest until it is over. In contrast, setting R_B equal to c/μ implies that the leader is the only agent that puts effort in Stage B (the laggard quits the contest) and thus it may take longer to reach the end goal. On the other hand, a higher fraction of the designer's budget is allocated to R_A when $R_B = c/\mu$ as opposed to when $R_B = R_B^{\min}$. As the cutoff belief $p_{\mathcal{F}}$ is decreasing

in R_A , this implies that the aggregate amount of experimentation in Stage A — and the probability that the contest is going to be completed eventually — is maximized by setting $R_B = c/\mu$.¹⁰ Thus, when setting the sizes of the two awards the designer trades off a higher probability of obtaining the innovation (by setting $R_B = c/\mu$) with getting it sooner (by setting $R_B = R_B^{\min}$).

Grand Prize without Feedback: At the other extreme of information disclosure, we have contests that feature a single final award R for completing the entire contest, i.e., Stage B , and no information disclosure in the interim. Agent i 's optimization problem can be then expressed as:

$$\max_{\{x_{i,\tau}\}} \int_0^\infty x_{i,\tau} \left(\lambda p_{i,\tau}^{\mathcal{N}} \left(q_{i,\tau}^{\mathcal{N}} \Pi(B, B, R) + (1 - q_{i,\tau}^{\mathcal{N}}) V_{i,\tau}(B, A) \right) - c \right) e^{-\int_{s=0}^\tau [p_{i,s}^{\mathcal{N}} (q_{i,s}^{\mathcal{N}} \mu + \lambda x_{i,s}) + r] ds} d\tau,$$

where $V_{i,\tau}(B, A)$ denotes the continuation value for agent i when she has completed Stage A whereas her competitor has not. Agent i cannot observe agent j 's partial progress, and so she forms beliefs about whether her competitor has completed Stage A . Specifically, $q_{i,\tau}^{\mathcal{N}}$ denotes agent i 's belief that her competitor has advanced to Stage B by time τ .

As we show in Proposition 2, equilibrium behavior takes the form of a cutoff experimentation policy as in the case of full information disclosure. However, in this case the time threshold $t_{\mathcal{N}}$ after which an agent stops experimenting in the absence of partial progress depends on her own effort provision as well as her belief about her competitor's progress over time. Like before, as soon as the agent completes Stage A , it is optimal for her to put effort until the contest is over. The proof of the proposition is omitted as it uses similar arguments as those in the proof of Proposition 7.

PROPOSITION 2. *Consider a contest design with a single final award in which no information about partial progress is ever disclosed. Then, there exists a unique symmetric equilibrium in which agents experiment as follows:*

(i) *Agents follow a cutoff experimentation policy in Stage A , i.e.,*

$$x_{i,t}^* = \begin{cases} 1 & \text{for } t \leq t_{\mathcal{N}} \\ 0 & \text{otherwise} \end{cases},$$

where the cutoff time $t_{\mathcal{N}}$ is given as the unique solution to the following equation

$$p_{i,t_{\mathcal{N}}}^{\mathcal{N}} \left(q_{i,t_{\mathcal{N}}}^{\mathcal{N}} \frac{\mu R - c}{2\mu + r} + (1 - q_{i,t_{\mathcal{N}}}^{\mathcal{N}}) \frac{\mu R - c}{\mu + r} \right) = \frac{c}{\lambda}.$$

Here, the posterior beliefs $p_{i,t}^{\mathcal{N}}$ and $q_{i,t}^{\mathcal{N}}$ are given by the following expressions:

$$p_{i,t}^{\mathcal{N}} = \frac{p_A e^{-\lambda t} \left(\frac{\lambda e^{-\mu t} - \mu e^{-\lambda t}}{\lambda - \mu} \right)}{p_A e^{-\lambda t} \left(\frac{\lambda e^{-\mu t} - \mu e^{-\lambda t}}{\lambda - \mu} \right) + (1 - p_A)} \quad \text{and} \quad q_{i,t}^{\mathcal{N}} = \frac{\lambda e^{-\mu t} - \lambda e^{-\lambda t}}{\lambda e^{-\mu t} - \mu e^{-\lambda t}}.$$

(ii) *If an agent completes Stage A , she experiments with rate one until the end of the contest.*

The juxtaposition of Propositions 1 and 2 illustrates the main tradeoff that the designer faces. In particular, equilibrium experimentation takes the form of a cutoff policy under both full and no information disclosure with different time cutoffs $t_{\mathcal{F}}$ and $t_{\mathcal{N}}$ respectively. As we establish below, assuming that both designs consume the same budget in expectation and $R_B = R_B^{\min}$ in the real-time leaderboard design (so that conditional on an agent completing Stage A , both agents compete until Stage B is complete), $t_{\mathcal{N}} > t_{\mathcal{F}}$, i.e., the probability that Stage B is going to be completed (and the designer will obtain the innovation) is higher when no information about partial progress is ever disclosed and the entire budget is allocated to a single final award for completing the project.

On the other hand though, there is a positive probability that in the case when information about partial progress is not disclosed, one of the agents drops out even though the other has completed Stage A . The latter case never occurs under full information disclosure (when $R_B = R_B^{\min}$). Thus, conditional on one agent completing Stage A , the contest is completed earlier in expectation when experimentation outcomes are publicly observable. For additional intuition, note that in the case when $\mu \rightarrow 0$ and conditional on an agent completing Stage A under both designs, the ratio of the expected times the innovation will be obtained under the two designs is equal to $1 + e^{-\lambda t_{\mathcal{N}}}$ (this is precisely due the fact that although under full disclosure both agents will compete until the end of the contest, under no disclosure the laggard will drop out with probability $e^{-\lambda t_{\mathcal{N}}}$).¹¹

PROPOSITION 3. *Consider a design that features full information disclosure with $R_B = R_B^{\min}$ and a design that features no information disclosure and has a single final award. Assume that the two designs consume the same budget in expectation. Then, the probability that an agent will complete the entire contest, i.e., Stage B , is higher for the design that features no information disclosure. On the other hand, conditional on the contest being completed, it takes less time to reach the innovation under full information disclosure.*

We plot the ratio of expected payoffs for the designer under the two designs as a function of the discount rate in Figure 2. When the agents and the designer are sufficiently patient, i.e., the discount rate takes small values, disclosing no information about partial progress outperforms a design where progress is publicly observable.

This tradeoff motivates the search for alternative information disclosure policies that combine the benefits of these two extremes. In the next section, we establish that appropriately timing the designer's announcements about the status of competition between the agents leads to strictly better outcomes for the designer.

4. Delaying the Disclosure of Information

We have established that full disclosure allows for the fast dissemination of good news, but may also adversely affect effort provision as agents become pessimistic about the feasibility of Stage




Figure 2 The ratio of the designer’s payoffs for the designs that induce full information disclosure and no information disclosure respectively as a function of the discount rate for $\lambda = 2, \mu = 0.04$ and $\lambda = 2, \mu = 0.02$ (the prior belief and the budget allocated to awards are $p_A = 0.2$ and $B = 5$ respectively).

A in the absence of partial progress. The fact that “*no news is bad news*” motivates exploring alternative designs that may feature silent periods — time intervals in which the designer does not disclose any information regarding the competitors’ progress (obviously, a special case is the design that features no information disclosure). In particular, we consider designs parameterized by \mathcal{T}, R_A , and R_B in which the designer’s information disclosure policy takes the following form: if any of the agents completes Stage B , then the designer discloses this information and the contest is over. For partial progress, the designer follows a policy that features a silent period: she does not disclose any information until time \mathcal{T} . At \mathcal{T} she discloses any partial progress that has occurred before then. Finally, after \mathcal{T} she discloses any progress as it happens. In other words,

$$\Phi_t = \begin{cases} 0 & \text{if } t < \mathcal{T} \\ 1 & \text{if } t \geq \mathcal{T} \text{ and } \tau_{i,t}^A \leq t \text{ for some } i \end{cases}.$$

On the other hand, the contest’s award scheme is such that R_B is given out to the first agent that completes Stage B whereas R_A is awarded only if Stage A has been completed by time \mathcal{T} . If both agents complete it before \mathcal{T} , then they both get $R_A/2$ (at the time they complete the stage).¹²

Note that the *silent period* design described above combines elements from both designs we studied in the previous section. In particular, before time \mathcal{T} no information is disclosed and agents become pessimistic at a relatively slow rate in the absence of any partial progress. On the other hand, partial progress is disclosed at time \mathcal{T} and thus it is still likely that both agents will continue competing until the contest is complete. Interestingly, as we establish below, when \mathcal{T} is chosen appropriately, a design with a silent period of length \mathcal{T} outperforms both the real-time leaderboard and the grand prize designs. For the remainder of this section, we assume that $R_B = R_B^{\min}$ which ensures that the laggard finds it optimal to compete with the leader until the contest is complete. Apart from simplifying the analysis, this is in line with our focus on the early stages of a contest when there is significant uncertainty and plenty of time until the contest is over.

First, we establish that agents follow a cutoff experimentation policy when information about partial progress is disclosed after an appropriately chosen time t_S .

PROPOSITION 4. *Consider a design with awards R_A and $R_B = R_B^{\min}$ that has a silent period of length t_S such that:*

$$p_{i,t_S}^S \left(q_{i,t_S}^S \left(R_A/2 + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{i,t_S}^S) \left(R_A + \Pi(B, A, R_B^{\min}) \right) \right) = \frac{c}{\lambda}, \quad (8)$$

where the posterior beliefs are such that

$$p_{i,t}^S = \frac{p_A e^{-\lambda t} \left(\frac{\lambda e^{-\mu t} - \mu e^{-\lambda t}}{\lambda - \mu} \right)}{p_A e^{-\lambda t} \left(\frac{\lambda e^{-\mu t} - \mu e^{-\lambda t}}{\lambda - \mu} \right) + (1 - p_A)} \quad \text{and} \quad q_{i,t}^S = \frac{\lambda e^{-\mu t} - \lambda e^{-\lambda t}}{\lambda e^{-\mu t} - \mu e^{-\lambda t}}.$$

Then, there exists a unique symmetric equilibrium in which agents set their effort levels to one until time t_S and quit if the designer does not disclose any partial progress. Otherwise, i.e., if the designer discloses that Stage A has been completed, both agents compete until Stage B is complete.

Proposition 4 allows us to compare silent period designs with full and no information disclosure. First, Proposition 5 states that having a silent period always leads to a higher expected payoff for the designer than full information disclosure.

PROPOSITION 5. *Consider a design that features a silent period of length $\mathcal{T} = t_S$ defined as in Expression (8). Then, this design outperforms one that features full information disclosure when the budgets allocated to awards for the two designs are equal in expectation and $R_B = R_B^{\min}$ for both.*

The main difference between the two designs is the rate at which beliefs drift downward in the absence of progress: under full information disclosure, agents become pessimistic at a faster rate as they can observe the experimentation outcomes of their competitors. The proof of Proposition 5 relies on this observation and establishes that when the budget allocated to awards in the two designs is kept fixed agents stop experimenting earlier under full information disclosure than in a design with a silent period of length t_S .

Furthermore, as we show below, the latter design also outperforms no information disclosure (assuming that Stage B takes relatively longer to complete than Stage A conditional on the innovation being feasible). Note that comparing the two designs is not straightforward as they have different award structures: in the case of a single grand prize, the designer's budget is allocated entirely to the award for completing the contest. On the other hand, silent period designs involve intermediate awards that may be split between the two competitors. As Proposition 6 states, when $\mu \rightarrow 0$, having a silent period of length t_S and offering an intermediate award to the agent(s) that complete Stage A by t_S yields a higher expected payoff for the designer.




Figure 3 The ratio of payoffs for the designer corresponding to the designs that feature a silent period and no information disclosure or full information disclosure respectively as a function of μ . Here, $\lambda = 2$ and $r = 1$ (the prior belief and the budget allocated to awards are $p_A = 0.2$ and $\mathcal{B} = 5$ respectively).

PROPOSITION 6. Assume $\mu \rightarrow 0$ and consider a design that features a silent period of length $\mathcal{T} = t_S$ defined as in Expression (8) and awards R_A and $R_B = R_B^{\min}$. Then, there exists \bar{p}_A such that for $p_A < \bar{p}_A$ the silent period design outperforms a design with a single grand prize when the budget allocated to awards is the same in expectation for the two designs.

The benefit of a silent period design compared to no information disclosure is that it induces a higher probability of both agents competing until the end of the contest. In addition, this benefit is more pronounced when Stage A is a relatively short part of the contest. Figure 3 illustrates that the silent period design dominates full and no disclosure for a wide range of values for μ .

Probabilistic Delay in Announcing Progress: So far, we have established that disclosing no information about the status of competition until some pre-determined time leads to a higher expected payoff for the designer than both full and no information disclosure. Next, we generalize this finding by showing that the silent period design we discussed above outperforms any design in which information about partial progress is disclosed at some (constant) rate ϕ (as opposed to being disclosed at a pre-determined time). In particular, we consider a design that features the following information disclosure policy:

$$\Phi_t = \begin{cases} 0 & \text{if } \tau_{1,t}^A = \tau_{2,t}^A = \emptyset \\ 1 - e^{-\phi(t-\tau)} & \text{otherwise, where } \tau = \min_i \{\tau_{i,t}^A\}. \end{cases}$$

In other words, conditional on at least one agent having completed Stage A by time t , the designer announces that partial progress has been made with rate ϕ , i.e., probability ϕdt in time interval $[t, t + dt)$ (working directly with rate ϕ is more convenient for our analysis). In addition, the design

also features an intermediate award R_A for the agent(s) that complete Stage A as well as an award R_B given out to the first agent that completes Stage B . We assume that award R_A is given out to the first agent that completes Stage A unless both complete it before the designer discloses any information in which case they split it equally. As before, we first establish that when information is disclosed at rate ϕ , agents follow a cutoff experimentation policy.

PROPOSITION 7. *Consider a design with awards R_A and $R_B = R_B^{\min}$ in which the designer discloses partial progress at rate ϕ . Then, there exists a unique symmetric equilibrium such that agents put full effort until time t_ϕ given by:*

$$p_{t_\phi}^\phi \left(q_{t_\phi}^\phi \left(R_A/2 + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{t_\phi}^\phi) \left(R_A + \frac{\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c}{\mu + \phi + r} \right) \right) = \frac{c}{\lambda}, \quad (9)$$

$$\text{where } p_{i,t}^\phi = \frac{p_A e^{-\lambda t} \left(\frac{\lambda e^{-(\mu+\phi)t} - (\phi+\mu)e^{-\lambda t}}{\lambda - \mu - \phi} \right)}{p_A e^{-\lambda t} \left(\frac{\lambda e^{-(\mu+\phi)t} - (\phi+\mu)e^{-\lambda t}}{\lambda - \mu - \phi} \right) + (1 - p_A)} \text{ and } q_{i,t}^\phi = \frac{\lambda e^{-(\mu+\phi)t} - \lambda e^{-\lambda t}}{\lambda e^{-(\mu+\phi)t} - (\mu + \phi)e^{-\lambda t}}.$$

In the absence of partial progress, an agent sets her effort level to zero after t_ϕ until the designer discloses that Stage A has been completed.

Comparing expressions (8) and (9) illustrates the difference of having the designer disclose information about partial progress at a pre-determined time t_S as opposed to doing so at a rate ϕ . In the former case, at any time $t < t_S$ and in the absence of any information about partial progress, agents are relatively more optimistic about the feasibility of Stage A than in the latter case (as when the designer discloses progress at rate ϕ no news is still relatively bad news). On the other hand, in the former case, given that they have no information about their competitors, the probability they assign to the event that their competitor has already completed Stage A and thus they would have to share award R_A in the case of breaking through, is higher. As we establish in Proposition 8 below, the first effect dominates the second and the designer finds it optimal to disclose information about partial progress at a pre-determined (deterministic) time as opposed to making announcements at stochastic times.¹³

PROPOSITION 8. *Assume that $\mu \rightarrow 0$ and consider a design with awards R_A and $R_B = R_B^{\min}$ that has a silent period of length t_S as given in Expression (8). Then, the design with silent period of length t_S leads to a higher expected payoff for the designer than a design at which information about partial progress is disclosed at rate $\phi > 0$ with awards R'_A and $R_B = R_B^{\min}$ such that the budget allocated to awards is the same in expectation for the two designs.*

Incentivizing Agents to Disclose their Progress: So far we have assumed that agents' progress is observable to the designer. But in many practical settings it may only be the agents' private information and the designer would have to incentivize them to disclose it. Expression (8) suggests that, when agents' progress is only privately observable, the design with a silent period of length t_S remains incentive compatible if the intermediate award given out for disclosing such information is high enough. In particular, note that the relevant incentive constraint is given by:¹⁴

$$q^S \left(\frac{R_A}{2} + \Pi(B, B, R_B^{\min}) \right) + (1 - q^S) \left(R_A + \Pi(B, A, R_B^{\min}) \right) \geq q^S \Pi(B, B, R_B^{\min}) + \frac{(1 - q^S) \mu R_B^{\min} - c}{\mu + r}, \quad (10)$$

where $q^S = \frac{\lambda e^{-\mu t_S} - \lambda e^{-\lambda t_S}}{\lambda e^{-\mu t_S} - \mu e^{-\lambda t_S}}$, denotes the probability an agent assigns to the event that her competitor has completed Stage A by time t_S . The right hand side of inequality (10) describes an agent's expected payoff assuming that she completes Stage A at time t_S and decides not to report her progress to the designer. In that case, the agent does not claim award R_A but in the case that her competitor has not already completed Stage A she can continue in the contest with no competition. Inequality (10) implies that if R_A satisfies the inequality below reporting her progress is always optimal for an agent

$$R_A \geq \left(1 - \frac{q^S/2}{1 - q^S/2} \right) \left(\frac{\mu R_B^{\min} - c}{\mu + r} - \Pi(B, A, R_B^{\min}) \right). \quad (11)$$

Belief q^S is increasing in R_A and, in turn, the first term on the right hand side of the inequality is decreasing in R_A . Thus, there exists \bar{R}_A such that the design described in Proposition 4 with $R_A \geq \bar{R}_A$ remains incentive compatible even when agents' progress is their private information.¹⁵ Thus, assuming that the budget allocated to awards is large enough, setting the contest's intermediate award R_A to a value higher than \bar{R}_A guarantees that agents disclose their progress to the designer as soon as it happens.

Asymmetric Information Disclosure Policies: As noted in Section 2, we restrict attention to information disclosure policies that involve public announcements from the contest designer, i.e., we abstract away from policies that could feature revealing information to only one of the two agents. Apart from being more challenging to implement, such disclosure rules do not lead to better outcomes for the designer. Below we provide a simple argument to illustrate this.

First, to facilitate the comparison with the designs we have discussed so far, we consider the following variation of the silent period design: the designer offers an intermediate award R_A if an agent completes Stage A before some pre-determined time \bar{t} . She also offers an award R_B for completing the entire contest. If both agents complete Stage A before time \bar{t} , they both receive $R_A/2$. Finally, the designer at time \bar{t} announces whether agent 1 has completed Stage A (but does

not disclose any information about agent 2’s progress — payments are deferred to the end of the contest and, thus, agent 1 cannot infer any information about how well agent 2 is doing). Then, it is straightforward to see that such a design is either dominated by the silent period design we discussed above (in which the designer discloses information about both agents’ progress) or by a design that features no information disclosure.

In particular, assume for the sake of contradiction that the asymmetric design we describe here led to a higher payoff for the designer than the silent period design. Comparing the two designs implies that the first, i.e., asymmetric disclosure, will feature more experimentation from agent 2 (as her competitor might quit in the absence of progress) but lower probability of having both agents compete until the end (since in the silent period design if one of the agents completes Stage *A* both will compete until the end). Given our assumption that asymmetric disclosure benefits the designer, we obtain that the additional experimentation from agent 2 outweighs the potential loss from having agent 1 quit the contest although agent 2 has completed Stage *A* before time \bar{t} .

Note though that not disclosing any information about agent 1 (in addition to not disclosing information about agent 2) results in the same benefit/loss for the designer. Thus, if asymmetric disclosure dominates the silent period design, then no information disclosure would dominate both. A similar argument can be used to show that if a design with asymmetric information disclosure dominates no information disclosure, then they would be both dominated by the silent period design. In summary, designs that feature asymmetric information disclosure policies do not yield any benefit in our environment and the focus to symmetric disclosure is without loss of generality.

We conclude the section by noting that the silent period design we describe here resembles the structure of many real-world tournaments. As an example, apart from the final grand prize, participants in the Netflix prize competed for intermediate awards that were given out at pre-determined times. In particular, Netflix was offering an annual progress prize to the team that showed the most improvement during the year, as long as this improvement was above a given threshold. This mirrors the design with a silent period: the designer gives out an intermediate award to the agent(s) that has completed Stage *A*, i.e., has progressed above a threshold, by some pre-determined time. Interestingly, the Netflix design allowed participants to disclose their progress as it happened in a public real-time leaderboard. However, since the awards were given out once a year, i.e., at pre-determined times, most of the teams posted their progress in the proximity of the deadline, effectively implementing a silent period until the intermediate award was given out.

5. When Competition is Dominant

Sections 3 and 4 consider contest design in the presence of uncertainty regarding the end goal. In this section, we discuss a complementary case when, given enough time and effort, innovation

occurs with probability one ($p_A = p_B = 1$). We also assume that Stage B is shorter (in expectation) than Stage A ($\mu > \lambda$). These two assumptions can be thought of as providing an approximation of the dynamics towards the end of the contest when uncertainty has been largely resolved and competition between the agents intensifies. In that case, the contest designer is facing a different tradeoff: when no information is disclosed about the agents' progress, they become progressively more pessimistic about the prospect of them winning, as they believe that someone must have made progress already and that they are lagging behind. On the other hand, making information about partial progress readily available may not benefit the designer either, as the agent who is lagging behind may find it optimal to quit. As we argue in the remainder of the section, a design with silent periods yields a higher expected payoff to the designer.

In particular, since innovation is certain, the interest of the contest designer is in achieving it as quickly as possible. Having both agents actively participating in the contest, i.e., exerting effort towards its completion, naturally expedites innovation compared to the case when only one of them remains active in the race, and thus the focus of the designer is on providing the right informational incentives for agents to continue experimenting and not drop out of the contest. These incentives may involve signaling to the agents that, relative to the competition, they are not lagging behind.

In particular, assume that the designer's budget and consequently the award structure is such that if an agent observes her competitor completing Stage A she has no incentive to continue putting effort in the contest (this is the most interesting case when $p_A = p_B = 1$). Then, unlike the case we study in Sections 3 and 4, no progress by her competitor is actually good news for an agent. This fact leads to a different tradeoff for the designer: her objective is to delay announcing progress by either of the agents as long as she can maximize the probability that both complete Stage A (and thus compete until the end). Agents, on the other hand, form beliefs about the progress of their competitors and, in the absence of any announcement from the designer, become pessimistic about their prospects of winning as they find it more likely that they are lagging behind.

For the remainder of the section, we assume that the designer's budget is entirely allocated to a single award given out to the first agent that completes the entire contest and focus on the impact of different disclosure policies on the agents' incentives for an effort provision. As before, we first compare the effort provision under full and no information disclosure of partial progress. When agents can perfectly observe each other's progress, they compete in Stage A by exerting full effort until one of them advances to Stage B at which point the laggard finds it optimal to quit. On the other hand, in the absence of any announcements, an agent's belief that her competitor is already in Stage B (which would imply that she should quit) increases. As a response, agents drop their effort levels to strike a balance between quitting the competition early and persisting in an attempt to win, without losing too much if it turns out they were lagging behind.

Although the agents' incentives for effort provision and the designer's tradeoff are quite different than in the case we covered in Sections 3 and 4, it turns out that a design that features silent periods again outperforms both full and no information disclosure (under some assumptions on the budget the designer allocates to awards). Instead of stating the results formally, we provide some intuition of why this is the case by describing how equilibrium behavior evolves over time depending on whether agents observe their competitors' progress.

In particular, under full information disclosure, agents race towards completing Stage A by exerting full effort. Upon the stage's completion, the leader continues towards completing Stage B whereas the laggard finds it optimal to quit. On the other hand, when agents cannot observe their competitors' partial progress, they form beliefs about whether they have already completed Stage A . As before, if we let $q_{1,t}$ denote the probability that agent 1 assigns at time t to the event that her competitor is in Stage B , then we have $dq_{1,t} = (1 - q_{1,t})(x_{2,t}\lambda - q_{1,t}\mu)dt$. Here, $x_{2,t}$ is the amount of effort that agent 1 believes agent 2 would allocate if she is still in Stage A (she would allocate effort equal to one if she is in Stage B). Interestingly, there exists a symmetric equilibrium in Markovian strategies that takes a simple form: in the absence of progress, agents put full effort up to some time \bar{t}_N , after which they drop their effort level to $q_{i_N}\mu/\lambda$, where q_{i_N} denotes an agent's belief that her competitor has completed Stage A by time \bar{t}_N . In other words, neither of the agents quits, but instead they continue exerting effort until one of them completes the entire contest (albeit with a rate lower than one after time \bar{t}_N).

Finally, as we mention above, a design that features silent periods leads to a higher expected payoff for the designer than both full and no information disclosure. In particular, consider the designer announcing the status of competition every \bar{t}_S time periods for some \bar{t}_S . As in the case when there is no disclosure, agents form beliefs regarding the likelihood that their competitors have already advanced to Stage B before the designer's announcement. Beliefs are reset at time \bar{t}_S if no progress is announced and the game essentially restarts. The probability that both agents progress to Stage B is positive (unlike the case when progress is publicly observable) and if the silent period is sufficiently short, effort levels are higher than in the case of no information disclosure (since beliefs are reset every \bar{t}_S). These two observations for the design with silent periods, i.e., the probability that agents will compete until the end of the contest is positive and beliefs are reset after each of the designer's announcements, imply that it outperforms both full and no disclosure.

As a final comment, note that the proposed design could be implemented in a straightforward way even when agents' progress is only privately observed. In particular, agents have the incentive to disclose their (partial) progress to the designer as soon as they complete a stage since such information would induce their competitors to quit the contest. Thus, unlike Sections 3 and 4 implementing the design does not require an intermediate award of a sufficiently high value to incentivize agents to disclose their progress.

6. Concluding Remarks

This paper studies the role of information in innovation contests and how it is inextricably linked to the encouragement and competition effects present in this setting. In particular, we examine the role of intermediate awards as information revelation devices that can be used to improve the performance of contests both in terms of the probability of reaching the end goal as well as the time it takes to complete the project. Interestingly, the role of an intermediate award depends on which of the two effects dominates: for the competition effect, an intermediate award that is *not handed out* is good news for the agents and increases their willingness to put in effort, since they believe they are still in the running to win the contest. When the encouragement effect dominates, an award that is handed out makes agents more optimistic about the feasibility of the project and provides an incentive for them to continue experimenting. This implies that the designer has to balance the trade-off between a higher level of aggregate experimentation in the early stages of the contest with a larger number of participants in later stages (i.e., faster completion of the contest) when determining the sizes of the awards.

We use a two-stage contest to provide an approximation of the dynamics in multi-stage contests. Naturally, the more progress being made, the less uncertain agents become about the feasibility of the end goal. Thus, a multi-stage contest can be thought of as having two distinct phases. First, during the early stages, uncertainty regarding the attainability of the end goal is the main driving force behind the competitors' actions. Competition is of secondary importance as there is still plenty of time for the laggards to catch up. We capture this situation as a two-stage contest in which the feasibility of the discovery required to complete the first stage is uncertain, i.e., $p_A < 1$. On the other hand, the second stage – which models the remainder of the contest – takes on average a much longer time to complete, i.e., the arrival rate associated with stage *B* is much lower than that of stage *A*. As the contest draws to an end, the dynamics become quite different. Agents are more optimistic about the feasibility of the end goal, but the chances for the laggards to catch up with the leader are slimmer. Thus, the agents' behavior is mainly prescribed by the competition effect. We capture this scenario by examining two successive stages that feature little or no uncertainty.

We show that a design that features silent periods — time intervals in which there is no information disclosure about the status of competition — as well as appropriately sized and timed intermediate awards for partial progress outperforms both the design when information about progress is not shared among competitors (implemented as a single grand prize for reaching the end goal) and the design that has a real-time leaderboard and gives out awards for partial progress as it happens (in which case agents are certain about the status of competition at all times). Silent periods have been implemented explicitly and implicitly as parts of real-world innovation

contests. For example, although the Netflix Prize had an online real-time leaderboard most of the activity was recorded close to the deadline of the annual progress prizes (the intermediate awards for partial progress), effectively imposing a silent period between two consecutive such deadlines. Matlab programming contests organized by Mathworks explicitly feature a silent period early on in the contest, the so-called “Darkness Segment,” after which participants are allowed to share their progress with their competitors (in what are known as the “Twilight” and “Daylight” segments).

The modeling framework in the paper can be used as a foundation for subsequent work that investigates the role of information disclosure in dynamic competition settings. More generally, our work is applicable to settings that involve a designer or a social planner selectively provides feedback to the agents involved. Finally, although we believe our setting captures the most important features of a dynamic contest, it has a number of limitations. Exploring optimal dynamic policies in the presence of both learning and competition is quite challenging and this is among very few recent papers that incorporate both of these features. Below we provide a list of potentially interesting directions for future research along with our thoughts on how they might affect our findings.

Uncertainty in Both Stages: The first part of the paper considers the early stages of an innovation contest. For the sake of tractability, we assume that there is uncertainty only about the feasibility of Stage A . Our analysis indicates that there is a trade-off between more experimentation in Stage A (and, thus, a higher probability of having at least one agent move to Stage B) and the time it takes to complete the contest. Even in the absence of discounting, a similar trade-off exists if there is uncertainty regarding the feasibility of both stages in the contest. The designer may find it optimal to incentivize agents to remain active in the contest even after a competitor completes Stage A in order to increase the aggregate amount of experimentation in the Stage B . Thus, we expect that our main qualitative insights regarding the optimality of designs that feature silent periods will continue to hold. The analysis becomes quite challenging however, with the main difficulty being that in addition to the agents’ beliefs about the feasibility of Stage A and the status of competition, the agents also have to form beliefs about the feasibility of Stage B , which in turn depend not only on whether a competitor has completed Stage A but also on when exactly this happened.

Skill Heterogeneity: We assume that agents are symmetric with respect to their skills as captured by rates λ and μ . An interesting direction for future research would be to relax this assumption and instead consider a setting in which agents are privately informed about their skills. In that case, giving out an intermediate award introduces an additional trade-off. The completion or not of Stage A by a competitor provides a signal regarding her skills and may further affect effort provision. The choice of the timing and size of awards becomes even more involved as the designer has to take this additional signal into account.

Contests with Many Stages: A typical contest may involve several milestones. As we have already argued, our analysis aims to capture the dynamics near the beginning and towards the end of the contest, where the encouragement and the competition effects, respectively, dominate. A contest consisting of a large number of stages may involve multiple intermediate awards. We conjecture that the interval between two consecutive awards increases at the beginning of the contest, thus reflecting the fact that uncertainty is gradually resolved. On the other hand, when the competition effect is dominant, agents become pessimistic about their relative progress; thus, the optimal interval between consecutive announcements by the designer decreases. At any given stage one of the two effects will be dominant and so our analysis would still apply. However, figuring out the optimal timing for the designer's announcements is quite challenging.

Appendix

Preliminaries: In the appendix we use $\Pi(k, \ell, R_B)$ to denote the expected payoff that agent i obtains when she is in Stage k , her competitor is in Stage ℓ , and the final award is R_B . We have:

- The expected payoff for an agent when both have completed Stage A (but not B) is given by

$$\Pi(B, B, R_B) = \begin{cases} (\mu R_B - c)/(2\mu + r) & \text{if } R_B \geq c/\mu \\ 0 & \text{otherwise} \end{cases}. \quad (12)$$

- The expected payoff for agent i when she is in Stage A and j is in Stage B is given by

$$\Pi(A, B, R_B) = \begin{cases} (\lambda \Pi(B, B, R_B) - c)/(\lambda + \mu + r) & \text{if } R_B \geq R_{\min}^B \\ 0 & \text{otherwise} \end{cases}. \quad (13)$$

- The expected payoff for agent i when she is in Stage B and j is in Stage A is given by

$$\Pi(B, A, R_B) = \begin{cases} \frac{\lambda \Pi(B, B, R_B) - c}{\lambda + \mu + r} + \frac{\mu R_B}{\lambda + \mu + r} & \text{if } R_B \geq R_{\min}^B \\ \frac{\mu R_B - c}{\mu + r} & \text{if } \frac{c}{\mu} \leq R_B < R_{\min}^B \end{cases}. \quad (14)$$

Notation: In the proofs that follow we use expressions for an agent's beliefs that Stage A is feasible and that her competitor has already completed Stage A denoted by p and q respectively. The superscripts in the beliefs refer to the respective contest designs. In particular, we have:

- Real-time leaderboard: Along the equilibrium path agents have a common belief about the feasibility of Stage A which in the absence of progress can be expressed as:

$$p_{i,t}^{\mathcal{F}} = \frac{p_A e^{-\int_{\tau=0}^t \lambda(x_{1,\tau} + x_{2,\tau}) d\tau}}{p_A e^{-\int_{\tau=0}^t \lambda(x_{1,\tau} + x_{2,\tau}) d\tau} + (1 - p_A)}, \quad (\text{superscript } \mathcal{F} \text{ refers to full information disclosure}).$$

- Grand prize without feedback: Assuming that both agents experiment with rate one up to time t , agent i 's belief about the feasibility of Stage A can be expressed as:

$$p_{i,t}^{\mathcal{N}} = \frac{p_A e^{-\lambda t} \left(\frac{\lambda e^{-\mu t} - \mu e^{-\lambda t}}{\lambda - \mu} \right)}{p_A e^{-\lambda t} \left(\frac{\lambda e^{-\mu t} - \mu e^{-\lambda t}}{\lambda - \mu} \right) + (1 - p_A)}.$$

Here, superscript \mathcal{N} refers to no information disclosure. Also, when agents cannot directly observe their competitor's progress, they form beliefs about whether they have already completed Stage A . For the grand prize design we have $q_{i,t}^{\mathcal{N}} = \frac{\lambda e^{-\mu t} - \lambda e^{-\lambda t}}{\lambda e^{-\mu t} - \mu e^{-\lambda t}}$.

- Silent periods: The expressions for beliefs $p_{i,t}^{\mathcal{S}}$ and $q_{i,t}^{\mathcal{S}}$ before the time at which the designer makes her first announcement about the status of competition take the same form as those for the grand prize design (superscript \mathcal{S} refers to silent period).

- Announcing progress at rate ϕ : Assuming that both agents experiment with rate one up to time t , agent i 's belief about the feasibility of Stage A can be expressed as:

$$p_{i,t}^{\phi} = \frac{p_A e^{-\lambda t} \left(\frac{\lambda e^{-(\mu+\phi)t} - (\phi+\mu)e^{-\lambda t}}{\lambda - \mu - \phi} \right)}{p_A e^{-\lambda t} \left(\frac{\lambda e^{-(\mu+\phi)t} - (\phi+\mu)e^{-\lambda t}}{\lambda - \mu - \phi} \right) + (1 - p_A)}.$$

Here, superscript ϕ refers to the rate at which the designer discloses partial progress. Also, agent i 's belief about whether her competitor has already completed Stage A is equal to:

$$q_{i,t}^{\phi} = \frac{\lambda e^{-(\mu+\phi)t} - \lambda e^{-\lambda t}}{\lambda e^{-(\mu+\phi)t} - (\mu + \phi)e^{-\lambda t}}.$$

Proof of Proposition 1

Assume that agent 2 is using strategy $\{x_{2,t}\}_{t \geq 0}$. We establish that the best response for agent 1 takes the form of a cutoff, i.e., she sets her effort level to one up to some time and then quits the contest if neither of the agents has completed Stage A . Consider agent 1's optimization problem:

$$\max_{\{x_{1,t}\}_{t \geq 0}} \int_0^{\infty} [x_{1,t}(p_{1,t}^{\mathcal{F}} \lambda (R_A + \Pi(B, A, R_B)) - c) + x_{2,t} p_{i,t}^{\mathcal{F}} \lambda \Pi(A, B, R_B)] e^{-\int_0^t (p_{1,\tau}^{\mathcal{F}} \lambda (x_{1,\tau} + x_{2,\tau}) + r) d\tau} dt,$$

where the term $e^{-\int_{\tau=0}^t p_{1,\tau}^{\mathcal{F}} \lambda (x_{1,\tau} + x_{2,\tau}) d\tau}$ is equal to the probability that neither of the agents has completed Stage A by time t .

Given that the final award for completing the contest is such that $R_B \leq R_B^{\min} = \frac{c}{\mu} \left(1 + \frac{2\mu+r}{\lambda}\right)$, it is straightforward to see that $\Pi(A, B, R_B^{\min}) = 0$, i.e., the continuation value of agent 1 if she is the laggard in the contest is equal to zero. Thus, we can rewrite the agent's optimization problem as

$$\max_{\{x_{1,t}\}_{t \geq 0}} \int_0^{\infty} [x_{1,t}(p_{1,t}^{\mathcal{F}} \lambda (R_A + \Pi(B, A, R_B)) - c)] e^{-\int_0^t (p_{1,\tau}^{\mathcal{F}} \lambda (x_{1,\tau} + x_{2,\tau}) + r) d\tau} dt. \quad (15)$$

The coefficient of $x_{1,t}$ in the expression above is decreasing over time since in the absence of progress $p_{1,t}^{\mathcal{F}}$ is non-increasing in t . This implies that agent 1's best response to any strategy from agent 2 is to set her effort level to one up to some time and then quit in the absence of progress.

Finally, to complete the claim we show that putting effort up to time $t_{\mathcal{F}}$, where $t_{\mathcal{F}}$ is given as in the statement of the proposition, constitutes a symmetric equilibrium. To see this assume that agent 2 puts full effort up to time $t_{\mathcal{F}}$. Then according to the first part of the proof the best response strategy for agent 1 takes the form of a time cutoff. Optimization problem (15) implies that the time at which agent 1 stops putting effort satisfies $\lambda p_{1,t}^{\mathcal{F}} (R_A + \Pi(B, A, R_B)) = c$, which together with Expression (14) for $\Pi(B, A, R_B)$ completes the proof. Q.E.D.

Proof of Proposition 6

First, we provide expressions for the expected payoffs for the designer in the two contest designs we consider. Then, we establish that under the assumptions of the proposition the contest with a silent period of length t_S leads to a higher expected payoff for the designer than the contest that features no information disclosure.

Consider a design that features a silent period of length t_S given by Expression (8). Then, the expected payoff for the designer is given by

$$\begin{aligned} U^S &= U \cdot \int_0^{t_S} 2\lambda e^{-(2\lambda+r)t} \left(\frac{\lambda}{\lambda+\mu+r} \frac{2\mu}{2\mu+r} + \frac{\mu}{\lambda+\mu+r} \right) dt \\ &= \frac{2\lambda\mu \cdot U}{(\lambda+r+\mu)(2\mu+r)} \frac{2\lambda+2\mu+r}{2\lambda+r} (1 - e^{-(2\lambda+r)t_S}), \end{aligned} \quad (16)$$

where U denotes the instantaneous utility that the designer derives from obtaining the innovation.

Next, for the design with a single final award and no information disclosure we have

$$\begin{aligned} U^N &= U \cdot \int_{t=0}^{t_N} e^{-(2\lambda+r)t} 2\lambda \left(\int_{\tau=0}^{t_N-t} e^{-(\lambda+\mu+r)\tau} \left(\lambda \frac{2\mu}{2\mu+r} + \mu \right) d\tau + e^{-(\lambda+\mu+r)(t_N-t)} \frac{\mu}{\mu+r} \right) dt \\ &= \frac{2\lambda\mu \cdot U}{(\mu+r)(\lambda+\mu+r)(2\mu+r)} \left[\left(1 - e^{-(2\lambda+r)t_N} \right) \frac{(\mu+r)(2\lambda+2\mu+r)}{2\lambda+r} - e^{-(2\lambda+r)t_N} \frac{\lambda r \left(e^{(\lambda-\mu)t_N} - 1 \right)}{\lambda - \mu} \right]. \end{aligned} \quad (17)$$

Similarly, we obtain the following expressions for the (expected) budgets \mathcal{B}^S and \mathcal{B}^N :

$$\begin{aligned} \mathbb{E}[\mathcal{B}^S] &\leq \int_{t=0}^{t_S} e^{-(2\lambda+r)t} 2\lambda \left(R_A + \left(\frac{\mu}{\lambda+\mu+r} + \frac{\lambda}{\lambda+\mu+r} \frac{2\mu}{2\mu+r} \right) R_B^{\min} \right) dt \\ &= \frac{2\lambda \left(1 - e^{-(2\lambda+r)t_S} \right)}{2\lambda+r} \left(R_A + \left(\frac{\mu}{\lambda+\mu+r} + \frac{\lambda}{\lambda+\mu+r} \frac{2\mu}{2\mu+r} \right) R_B^{\min} \right), \end{aligned} \quad (18)$$

where the inequality is due to the fact that the intermediate award may be split between the two agents (when they both complete the stage before time t_S). Also, we have:

$$\mathbb{E}[\mathcal{B}^N] = \frac{2\lambda\mu R \left[\left(1 - e^{-(2\lambda+r)t_N} \right) \frac{(\mu+r)(2\lambda+2\mu+r)}{2\lambda+r} - e^{-(2\lambda+r)t_N} \frac{\lambda r \left(e^{(\lambda-\mu)t_N} - 1 \right)}{\lambda - \mu} \right]}{(\mu+r)(\lambda+\mu+r)(2\mu+r)}. \quad (19)$$

Expressions (16), (17), (18), and (19) along with the fact that $\mathbb{E}[\mathcal{B}^S] = \mathbb{E}[\mathcal{B}^N]$ yield the following inequality for the ratio of the expected payoffs for the designer

$$\frac{U^S}{U^N} \geq \frac{(2\lambda+2\mu+r)\mu}{(\lambda+\mu+r)(2\mu+r)} \frac{R}{R_A + \left(\frac{\mu}{\lambda+\mu+r} + \frac{\lambda}{\lambda+\mu+r} \frac{2\mu}{2\mu+r} \right) R_B^{\min}}. \quad (20)$$

The rest of the proof establishes that the right hand side of (20) is strictly greater than one, i.e.,

$$R \geq \frac{(\lambda+\mu+r)(2\mu+r)}{(2\lambda+2\mu+r)\mu} R_A + R_B^{\min}. \quad (21)$$

To establish (21), we turn our attention to the cutoff times t_S and t_N . Note that if $t_S \geq t_N$, the claim follows directly as then trivially the expected payoff for the designer is higher in the silent period design. So, we consider the case when $t_S < t_N$, and equivalently $p_{t_S}^S > p_{t_N}^N$. Then, by the characterization of the cutoff times we obtain:

$$q_{t_S}^S \left(\frac{R_A}{2} + \frac{\mu R_B^{\min} - c}{2\mu + r} \right) + (1 - q_{t_S}^S) \left(R_A + \frac{\mu R_B^{\min}}{\lambda + \mu + r} \right) < q_{t_N}^N \frac{\mu R - c}{2\mu + r} + (1 - q_{t_N}^N) \frac{\mu R - c}{\mu + r}.$$

Note that $q_{t_S}^S < q_{t_N}^N$ when $t_S < t_N$, thus, we can rewrite the inequality above as:

$$q_{t_S}^S \left(\frac{R_A}{2} + \frac{\mu R_B^{\min} - c}{2\mu + r} \right) + (1 - q_{t_S}^S) \left(R_A + \frac{\mu R_B^{\min}}{\lambda + \mu + r} \right) < q_{t_S}^S \frac{\mu R - c}{2\mu + r} + (1 - q_{t_S}^S) \frac{\mu R - c}{\mu + r}. \quad (22)$$

The definition of R_B^{\min} implies that $\frac{c}{\mu + r} = \frac{\lambda \mu}{(2\mu + \lambda + r)(\mu + r)} R_B^{\min}$. Thus, we can cancel out and rearrange the terms involving c and rewrite (22) as

$$\begin{aligned} q_{t_S}^S \left(\frac{R_A}{2} + \frac{\mu R_B^{\min}}{2\mu + r} \right) + (1 - q_{t_S}^S) \left(R_A + \left(\frac{\mu}{\lambda + \mu + r} + \frac{\lambda}{2\mu + \lambda + r} \frac{\mu}{\mu + r} \right) R_B^{\min} \right) \\ < q_{t_S}^S \frac{\mu R}{2\mu + r} + (1 - q_{t_S}^S) \frac{\mu R}{\mu + r} < \frac{\mu R}{\mu + r}. \end{aligned} \quad (23)$$

For the remainder of the proof, we assume that $p_A < \bar{p}_A$ where \bar{p}_A is such that

$$q_{t_S}^S < \frac{\lambda r}{(2\lambda + 2\mu + r)(\mu + r)}.$$

Note that both t_S and $q_{t_S}^S$ are decreasing in p_A , thus such \bar{p}_A always exists. Inequality (23) along with the upper bound on $q_{t_S}^S$ yield

$$R > \left(1 - \frac{\lambda r}{2(2\lambda + 2\mu + r)(\mu + r)} \right) \frac{\mu + r}{\mu} R_A + \frac{\mu + r}{2\mu + r} R_B^{\min}. \quad (24)$$

Inequality (24) implies that the ratio given in (20) is strictly greater than one when

$$R_A > \frac{2c\mu(2\lambda + 2\mu + r)(2\mu + \lambda + r)}{\lambda^2 r(2\mu + r)}. \quad (25)$$

Finally, note that inequality (25) always holds when $\mu \rightarrow 0$ as the right hand side also goes to zero at the limit, which completes the proof of the claim. Q.E.D.

Proof of Proposition 8

The proof consists of two steps. We first show that when the designer uses the same intermediate award R_A and final award R_B^{\min} in both information disclosure policies, then agents exert more cumulative effort under the silent period design than under the ϕ -design, i.e., when partial progress is disclosed at rate ϕ . In the second step, we show that the result from the first step still holds when the two designs use the same budget.

Step 1: According to Proposition 7 when the designer discloses partial progress at rate ϕ and the design has awards R_A and R_B^{\min} , agents put full effort up to time t_ϕ , which is defined by:

$$\lambda p_{t_\phi}^\phi \left(q_{t_\phi}^\phi \left(\frac{R_A}{2} + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{t_\phi}^\phi) \left(R_A + \frac{\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c}{\mu + \phi + r} \right) \right) - c = 0. \quad (26)$$

Next, consider a design featuring a silent period with length equal to t_ϕ defined above. The remainder of the proof establishes that when $\mu \rightarrow 0$ both agents exert full effort under this disclosure policy at least up to time t_ϕ , thus implying that the probability of completing the first stage (and the contest) in that design is higher than that of the disclosure induced by rate ϕ . Furthermore, disclosing information at time t_ϕ ensures that the probability that both agents complete Stage A and compete until the end is higher than in the equilibrium induced in the ϕ -design.

Denote by $p_{t_\phi}^S$ and $q_{t_\phi}^S$ the beliefs that agents have about the feasibility of Stage A and about whether the competitor is in Stage B, respectively, under the silent period design with length t_ϕ . Note that the claim follows if $\lambda p_{t_\phi}^S \left(q_{t_\phi}^S \left(\frac{R_A}{2} + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{t_\phi}^S) \left(R_A + \Pi(B, A, R_B^{\min}) \right) \right) \geq c$, or equivalently by (26) if

$$p_{t_\phi}^S \left(q_{t_\phi}^S \left(\frac{R_A}{2} + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{t_\phi}^S) \left(R_A + \Pi(B, A, R_B^{\min}) \right) \right) \geq p_{t_\phi}^\phi \left(q_{t_\phi}^\phi \left(\frac{R_A}{2} + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{t_\phi}^\phi) \left(R_A + \frac{\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c}{\mu + \phi + r} \right) \right). \quad (27)$$

We multiply both sides of (27) by $(1 - q_{t_\phi}^S + p_{t_\phi}^\phi (q_{t_\phi}^S - q_{t_\phi}^\phi)) (1 - q_{t_\phi}^\phi)$ and use (26) to rewrite (27) as:

$$p_{t_\phi}^S \left(\frac{1 - q_{t_\phi}^S + p_{t_\phi}^\phi (q_{t_\phi}^S - q_{t_\phi}^\phi)}{(1 - q_{t_\phi}^\phi) p_{t_\phi}^\phi} \right) (1 - q_{t_\phi}^\phi) \left(q_{t_\phi}^S \left(\frac{R_A}{2} + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{t_\phi}^S) \left(R_A + \Pi(B, A, R_B^{\min}) \right) \right) \geq (1 - q_{t_\phi}^S) \left(q_{t_\phi}^\phi \left(\frac{R_A}{2} + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{t_\phi}^\phi) \left(R_A + \frac{\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c}{\mu + \phi + r} \right) \right) + (q_{t_\phi}^S - q_{t_\phi}^\phi) c / \lambda. \quad (28)$$

Note that $R_B = R_B^{\min}$ implies that $\Pi(B, B, R_B^{\min}) = c/\lambda$ and thus we can replace the last term in the right hand side of the inequality above by $(q_{t_\phi}^S - q_{t_\phi}^\phi) \Pi(B, B, R_B^{\min})$. Also, note that we have:

$$\frac{p_{t_\phi}^\phi (1 - q_{t_\phi}^\phi)}{1 - q_{t_\phi}^S - p_{t_\phi}^\phi (q_{t_\phi}^\phi - q_{t_\phi}^S)} = p_{t_\phi}^S. \quad (29)$$

Finally, using Equation (29) we can rewrite (28) as:

$$(q_{t_\phi}^S - q_{t_\phi}^\phi) \frac{R_A}{2} - (1 - q_{t_\phi}^\phi) (1 - q_{t_\phi}^S) \left(\frac{\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c}{\mu + \phi + r} \right) - \Pi(B, A, R_B^{\min}) \geq 0. \quad (30)$$

Note that the above inequality holds since when $\phi > 0$ we always have $q_{t_\phi}^S > q_{t_\phi}^\phi$ and, in addition, when $\mu \rightarrow 0$ it holds that $\lim_{\mu \rightarrow 0} \left(\frac{\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c}{\mu + \phi + r} - \Pi(B, A, R_B^{\min}) \right) = 0$.

Step 2: In step 1 we established that if both information disclosure policies use the same intermediate and final awards then agents exert more effort under the design featuring a silent period of appropriately determined length t_S than in the ϕ -design. This step completes the proof by showing that the same holds even when the two designs consume the same budget in expectation.

To this end, assume that in the design with a silent period we decrease the intermediate award from R_A to R'_A , such that the two designs consume the same budget. Then, there are two cases to consider. First, note that if agents stop exerting effort in the absence of any news under the ϕ -design earlier than under the design with a silent period, i.e., if $t_\phi < t_S$, the result follows directly. Thus, for the remainder of the proof we show that the claim holds also when $t_\phi \geq t_S$. In that case, note that since $R'_A < R_A$, the designer spends a higher fraction of her budget on intermediate awards under the ϕ -design than under a silent period design. This holds since the expected budget allocated to the intermediate award in the ϕ -design is lower bounded as follows:

$$\begin{aligned} & \int_{t=0}^{t_\phi} e^{-(2\lambda+r)t} 2\lambda \left(\int_{\tau=0}^{t_\phi-t} e^{-(\lambda+\mu+\phi+r)\tau} \left((\mu+\phi)e^{r\tau} R_A + \lambda \left(\frac{R_A}{2} + e^{r\tau} \frac{R_A}{2} \right) \right) + e^{-(\lambda+\mu+\phi+r)(t_\phi-t)} e^{r(t_\phi-t)} R_A \right) dt \\ & \geq \int_{t=0}^{t_\phi} e^{-(2\lambda+r)t} 2\lambda \left(\int_{\tau=0}^{t_\phi-t} e^{-(\lambda+\mu+r)\tau} \left(\mu e^{r\tau} R_A + \lambda \left(\frac{R_A}{2} + e^{r\tau} \frac{R_A}{2} \right) \right) + e^{-(\lambda+r)(t_\phi-t)} e^{r(t_\phi-t)} R_A \right) dt \\ & \geq \int_{t=0}^{t_S} e^{-(2\lambda+r)t} 2\lambda \left(\int_{\tau=0}^{t_S-t} e^{-(\lambda+\mu+r)\tau} \left(\mu e^{r\tau} R'_A + \lambda \left(\frac{R'_A}{2} + e^{r\tau} \frac{R'_A}{2} \right) \right) + e^{-(\lambda+r)(t_S-t)} e^{r(t_S-t)} R'_A \right) dt, \end{aligned}$$

where the first inequality holds, since for a fixed stopping time t_ϕ , the expression is decreasing in ϕ (and, thus, it is minimized for $\phi = 0$). The second inequality holds since we $t_\phi \geq t_S$ and $R_A > R'_A$.

Finally, note that the last expression is equal to the budget allocated to the intermediate award for the silent period design in expectation. Thus, in order to have both designs use the same budget in expectation, it must be the case that the designer allocates a higher fraction of her budget to the final award R_B^{\min} in the silent period design. In turn, this implies that the designer's utility from obtaining the innovation is also higher under the silent period design since the latter is discounted at the same rate as the final award R_B^{\min} . Q.E.D.

Endnotes

1. We use the terms “participant”, “competitor”, and “agent” interchangeably throughout.
2. The NetFlix Prize offered a million dollars to anyone who succeeded in improving the company's recommendation algorithm by a certain margin and was concluded in 2009. The Heritage Prize was a multi-year contest whose goal was to provide an algorithm that predicts patient readmissions to hospitals. A successful breakthrough was obtained in 2013.
3. We note that the usage of the term “encouragement” is different from the literature on strategic experimentation (e.g., [Bolton and Harris \(1999\)](#), [Keller et al. \(2005\)](#)), where an agent may find it optimal to exert effort in order to “encourage” other agents to experiment in the future and,

therefore, generate valuable information about the state (since experimentation outcomes are perfectly observable by all). In the present paper, an agent becomes encouraged to exert more effort as she updates her belief about the state of the world because of progress made by others.

4. We remark that there is a conflation of the terms “contest” and “race” in our work. The two agents compete in an innovation race towards a pre-determined goal whose feasibility is unknown. We use the term contest to be in line with the literature, e.g., [Halac et al. \(2016\)](#).

5. Relatedly, [Seel and Wasser \(2014\)](#) consider the design of an optimal “head start” that is given to a player whereas [Seel \(2014\)](#) analyzes all-pay auctions where one agent is uncertain about the size of her competitor’s head start ([Siegel \(2009\)](#) refers to an agent’s effort as her score and allows for head starts in his quite general framework for all-pay auctions).

6. Assuming that λ is constant over time is essential for making the analysis tractable (and is in line with the prior work on strategic experimentation referenced on Page 5). That said, an important aspect of innovation contests is learning-by-doing, which could be incorporated in our model through a breakthrough rate that is an increasing function of the effort an agent has allocated to the contest. Although we conjecture that extending our model towards this direction would not affect our qualitative insights, the corresponding analysis is significantly more involved.

7. We assume that the two competitors cannot communicate directly with one another. Communicating an agent’s own progress as soon as it occurs would not yield any benefit to the agent unless the competition effect dominates (as in Section 5).

8. To optimize over the designer’s net payoff, i.e., the utility from obtaining the innovation minus the budget allocated to the award structure, one could use our analysis (that studies how to optimally use a fixed budget) and then optimize over the size of the budget. As it turns out, when the value of obtaining the innovation is sufficiently high for the designer, our findings illustrate that the optimal design takes qualitatively the same form irrespective of the exact size of the budget.

9. [Choi \(1991\)](#) provides similar expressions for the expected payoffs of the leader and the laggard in a model where outcomes are publicly observable, effort levels are binary, quitting the race is irreversible, and the value of winning an innovation race is fixed ex-ante. It is worthwhile to note that although in principle our model allows for interior levels of effort and the decision to quit the race is not irreversible, it turns out that under the information disclosure policies we study, agents’ effort levels are binary and quitting is irreversible at equilibrium (much like in [Choi \(1991\)](#)).

10. Lemma [EC.2](#) establishes that setting $R_B = c/\mu$ leads to more aggregate experimentation in Stage A than $R_B = R_B^{\min}$.

11. When $\mu \rightarrow 0$ and conditional on the contest being completed under both designs, the time it takes to complete it is approximately equal to the expected time it will take to complete Stage B if one or both agents exert effort to complete the project respectively, weighted by the probability

that an agent will quit the contest under no information disclosure. Given that the rate at which each agent completes Stage B is the same, the expected time to complete it is twice as long when only one agent exerts effort than when both do so.

12. As a remark, Equation (8) clarifies why it is optimal for the designer to split award R_A to the two agents if both complete Stage A before time \mathcal{T} (as opposed to giving out a higher fraction of the award to the agent that completed the stage first). As time goes by, the probability that a given agent will be the first to complete Stage A conditional on both completing it decreases and thus splitting R_A in half guarantees that the agents' effort provision is maximized.

13. Although we state Proposition 8 for a constant rate $\phi > 0$, we mainly use this fact to ensure that agents put full effort up to some time at equilibrium after which they stop and wait for an announcement from the designer. In other words, we strongly conjecture that the proposition holds for any disclosure policy as long as agents follow a cutoff experimentation policy at equilibrium. Although the latter is not generally true, i.e., there exist disclosure policies under which it is not optimal for agents to experiment using a cutoff policy, we believe that such disclosure policies will be suboptimal for the designer as they introduce delay to the agents' experimentation. Obtaining such a result seems quite challenging, since it requires a general mapping between any disclosure policy to equilibrium effort. Therefore, we chose to restrict our attention to constant rates.

14. It is straightforward to see that when the size of the intermediate award is fixed over time the binding incentive constraint is the one that corresponds to time t_S .

15. An upper bound for \bar{R}_A is obtained by setting $q^S = 0$, i.e., $\bar{R}_A < \left(\frac{\mu R_B^{\min} - c}{\mu + r} - \Pi(B, A, R_B^{\min}) \right)$.

Acknowledgments

We are thankful to David Gamarnik, an Associate Editor, and three referees for very constructive feedback. We are also grateful to Alessandro Bonatti, Kimon Drakopoulos, Karim Lakhani, Mohammad Mousavi, Philipp Strack, Alireza Tahbaz-Salehi, and seminar participants at Columbia, Chicago Booth, Caltech, UCSD, CMU, UC Berkeley, University of Michigan, USC, UT Austin, NYU, UPF, and Kaggle.

References

- Ales, Laurence, Soo-Haeng Cho, Ersin Körpeoğlu. 2016a. Innovation tournaments with multiple contributors. *Working paper* .
- Ales, Laurence, Soo-Haeng Cho, Ersin Körpeoğlu. 2016b. Optimal award scheme in innovation tournaments. *Forthcoming in Operations Research* .
- Bimpikis, K, K Drakopoulos. 2016. Disclosing information in strategic experimentation. *Working paper* .
- Bolton, P, C Harris. 1999. Strategic experimentation. *Econometrica* **67**(2) 349–374.
- Bonatti, A, J Horner. 2011. Collaborating. *American Economic Review* **101**(2) 632–663.

- Boudreau, K J, N Lacetera, K R Lakhani. 2011. Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science* **57**(5) 843–863.
- Che, Yeon-Koo, Ian Gale. 2003. Optimal design of research contests. *American Economic Review* **93**(3) 646–671.
- Choi, Jay P. 1991. Dynamic R&D competition under “hazard rate” uncertainty. *The RAND Journal of Economics* **22**(4) 596–610.
- Halac, Marina C, Qingmin Liu, Navin Kartik. 2016. Contests for experimentation. *Forthcoming in the Journal of Political Economy* .
- Harris, Christopher, John Vickers. 1987. Racing with uncertainty. *The Review of Economic Studies* **54**(1) 1–21.
- Hörner, Johannes. 2004. A perpetual race to stay ahead. *The Review of Economic Studies* **71**(4) 1065–1088.
- Keller, G, S Rady. 2010. Strategic experimentation with poisson bandits. *Theoretical Economics* **5**(2) 275–311.
- Keller, G, S Rady, M Cripps. 2005. Strategic experimentation with exponential bandits. *Econometrica* **73**(1) 39–68.
- Klein, Nicolas, Sven Rady. 2011. Negatively correlated bandits. *The Review of Economic Studies* **78**(2) 693–732.
- Lang, Matthias, Christian Seel, Philipp Strack. 2014. Deadlines in stochastic contests. *Journal of Mathematical Economics* **52** 134–142.
- Malueg, David A, Shunichi O Tsutsui. 1997. Dynamic R&D competition with learning. *The RAND Journal of Economics* 751–772.
- Moldovanu, Benny, Aner Sela. 2001. The optimal allocation of prizes in contests. *American Economic Review* 542–558.
- Moldovanu, Benny, Aner Sela. 2006. Contest architecture. *Journal of Economic Theory* **126**(1) 70–96.
- Moscarini, Giuseppe, Lones Smith. 2007. Optimal dynamic contests. *Working Paper* .
- Seel, Christian. 2014. The value of information in asymmetric all-pay auctions. *Games and Economic Behavior* **86** 330 – 338.
- Seel, Christian, Cédric Wasser. 2014. On optimal head starts in all-pay auctions. *Economics Letters* **124**(2) 211 – 214.
- Siegel, Ron. 2009. All-pay contests. *Econometrica* **77**(1) 71–92.
- Taylor, C. 1995. Digging for golden carrots: An analysis of research tournaments. *American Economic Review* **85**(4) 872–90.
- Terwiesch, C, Y Xu. 2008. Innovation contests, open innovation, and multiagent problem solving. *Management Science* **54**(9) 1529–1543.

Supporting Material

Proof of Lemma 1

In the absence of progress over time period $[t, t + dt)$, agent i updates her belief about the feasibility of Stage A by applying Bayes' rule as follows:

$$p_{i,t+dt}^{\mathcal{F}} = \frac{p_{i,t}^{\mathcal{F}}(1 - \lambda(x_{i,t} + x_{-i,t})dt)}{p_{i,t}^{\mathcal{F}}(1 - \lambda(x_{i,t} + x_{-i,t})dt) + (1 - p_{i,t}^{\mathcal{F}})},$$

since the probability that either of the agents completes Stage A over $[t, t + dt)$ when the stage is feasible is given by $\lambda(x_{i,t} + x_{-i,t})dt$. Therefore, we have $dp_{i,t}^{\mathcal{F}} = -p_{i,t}^{\mathcal{F}}(1 - p_{i,t+dt}^{\mathcal{F}})(x_{i,t} + x_{-i,t})\lambda dt$, which completes the proof of the lemma. Q.E.D.

For completeness, Figure EC.1 illustrates the evolution of the posterior beliefs in the absence of progress under full and no information disclosure.




Figure EC.1 Posterior belief over time in the absence of progress (here $\lambda = 3$ and $\mu = 1$).

LEMMA EC.1. Consider a contest design in which progress is publicly observable and the designer's expected budget for awards R_A and R_B is fixed and sufficiently high. Then, it is optimal for the designer to set $R_B = c/\mu$ or $R_B = R_B^{\min}$.

Proof: Assume for the sake of contradiction that the designer's expected utility under full information disclosure is higher for some $R'_B > R_B^{\min}$ than when $R_B = R_B^{\min}$. In particular, assume that

$$\begin{aligned} \int_{t=0}^{\infty} e^{-rt} e^{-\lambda \int_{\tau=0}^t (x_{1,\tau}^* + x_{2,\tau}^*) d\tau} (x_{1,t}^* + x_{2,t}^*) \left[\frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right] dt \\ > \int_{t=0}^{t^{\mathcal{F}}} e^{-(r+2\lambda)t} 2\lambda \left[\frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right] dt, \end{aligned} \quad (\text{EC.1})$$

where $\{x_{1,t}^*\}, \{x_{2,t}^*\}$ denote the equilibrium effort levels of agents 1 and 2 respectively when $R_B = R'_B$.

Next, we compare the (expected) budget allocated to awards for the two designs described above. Note that for the design for which $R_B = R'_B$ we have:

$$\int_{t=0}^{\infty} e^{-rt} e^{-\lambda \int_{\tau=0}^t (x_{1,\tau}^* + x_{2,\tau}^*) d\tau} \lambda (x_{1,t}^* + x_{2,t}^*) \left(R'_A + \left[\frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right] R'_B \right) dt, \quad (\text{EC.2})$$

whereas for the design for which $R_B = R_B^{\min}$ we have

$$\int_{t=0}^{t_{\mathcal{F}}} e^{-(r+2\lambda)t} 2\lambda \left(R_A + \left[\frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right] R_B^{\min} \right) dt. \quad (\text{EC.3})$$

Expressions (EC.1), (EC.2), and (EC.3) along with the fact that the budget allocated to awards is the same in the two designs (i.e., Expressions (EC.2) and (EC.3) are equal) imply that

$$R_A + \left[\frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right] R_B^{\min} > R'_A + \left[\frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right] R'_B. \quad (\text{EC.4})$$

The proof of the claim follows from showing that inequality (EC.4) implies that the belief at which agents stop experimenting when $R_B = R_B^{\min}$ is lower than the one that they stop experimenting when $R_B = R'_B$, i.e., agents experiment more when $R_B = R_B^{\min}$. In particular, agents stop experimenting when their expected instantaneous payoff is equal to c which, in turn, implies that the corresponding cutoff beliefs at which agents stop experimenting are such that:

$$p' \left(R'_A + \left(\frac{\lambda}{\lambda + \mu + r} \frac{\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right) R'_B - \frac{2\mu + \lambda + r}{(\lambda + \mu + r)(2\mu + r)} c \right) = c, \quad \text{and} \quad (\text{EC.5})$$

$$p \left(R_A + \left(\frac{\lambda}{\lambda + \mu + r} \frac{\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right) R_B - \frac{2\mu + \lambda + r}{(\lambda + \mu + r)(2\mu + r)} c \right) = c. \quad (\text{EC.6})$$

From Equations (EC.5) and (EC.6) we obtain that $p < p'$ if

$$R'_A + \left[\frac{\lambda}{\lambda + \mu + r} \frac{\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right] R'_B < R_A + \left[\frac{\lambda}{\lambda + \mu + r} \frac{\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right] R_B^{\min}. \quad (\text{EC.7})$$

Inequality (EC.7) follows from (EC.4) and the fact that $R'_B > R_B^{\min}$ which, in turn, leads to a contradiction. Using similar arguments we can also show that the designer's expected utility is the same for any $c/\mu \leq R_B < R_B^{\min}$.

LEMMA EC.2. *Consider a contest design in which progress is publicly observable and the designer's expected budget for awards R_A and R_B is fixed and sufficiently high. Then, setting $R_B = c/\mu$ leads to more aggregate experimentation in Stage A than setting $R_B = R_B^{\min}$.*

Proof: Consider the following two contest designs in both of which progress is publicly observable: the first features awards for completing Stages A and B that are equal to some R_A and $R_B = c/\mu$ respectively. The second is such that the corresponding awards are R'_A and $R_B = R_B^{\min}$. Furthermore, let $t_{\mathcal{F}}$ and $t'_{\mathcal{F}}$ denote the times at which agents stop putting effort in the absence of progress under the two designs. For the sake of contradiction, assume that agents exert more aggregate effort in stage A under the design with final award R_B^{\min} , i.e., $t'_{\mathcal{F}} > t_{\mathcal{F}}$ and, in turn, the corresponding stopping beliefs are such that $p'_{\mathcal{F}} < p_{\mathcal{F}}$. Recall that according to Expression (7), we have:

$$\begin{aligned} p'_{\mathcal{F}} \left(R'_A + \Pi(B, A, R_B^{\min}) \right) &= \frac{c}{\lambda} \quad \text{and} \\ p_{\mathcal{F}} \left(R_A + \frac{\mu R_B - c}{\mu + r} \right) &= \frac{c}{\lambda}. \end{aligned} \tag{EC.8}$$

The equalities above imply that:

$$R'_A + \frac{\mu R_B^{\min}}{\lambda + \mu + r} > R_A. \tag{EC.9}$$

Since both designs have to use the same budget in expectation, we have:

$$\int_{t=0}^{t_{\mathcal{F}}} e^{-(2\lambda+r)t} 2\lambda \left(R_A + \frac{\mu}{\mu+r} \frac{c}{\mu} \right) dt = \int_{t=0}^{t'_{\mathcal{F}}} e^{-(2\lambda+r)t} 2\lambda \left(R'_A + \left(\frac{\mu}{\lambda + \mu + r} + \frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} \right) R_B^{\min} \right) dt. \tag{EC.10}$$

According to our assumption that $t'_{\mathcal{F}} > t_{\mathcal{F}}$ we obtain that $R_A + \frac{\mu}{\mu+r} \frac{c}{\mu} > R'_A + \left(\frac{\mu}{\lambda + \mu + r} + \frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} \right) R_B^{\min}$. Substituting the expression for R_B^{\min} yields

$$R_A > R'_A + \frac{\mu R_B^{\min}}{\lambda + \mu + r} + \left(\frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} \frac{2\mu + \lambda + r}{\lambda} - \frac{\mu}{\mu + r} \right) \frac{c}{\mu},$$

which contradicts inequality (EC.9).

Proof of Proposition 3

First, it is straightforward to see that conditional on the contest being completed under both full and no information disclosure, the time it would take to reach the end goal is shorter in expectation under full disclosure. This is a direct consequence of the fact that the real-time leaderboard design incentivizes both agents to compete by putting full effort until one completes the entire contest (when $R_B = R_B^{\min}$). This is not necessarily the case in the grand prize design since even conditional on completing the contest there is a positive probability that the laggard quits.

In what follows we establish the first claim, i.e., $t_{\mathcal{N}} > t_{\mathcal{F}}$. To this end, note that Proposition 2 states that in the design with a single grand prize of size R and no information disclosure about partial progress, agents put effort with rate one up to time $t_{\mathcal{N}}$ such that

$$\lambda p_{t_{\mathcal{N}}}^{\mathcal{N}} \left(q_{t_{\mathcal{N}}}^{\mathcal{N}} \Pi(B, B, R) + (1 - q_{t_{\mathcal{N}}}^{\mathcal{N}}) \frac{\mu R - c}{\mu + r} \right) = c. \tag{EC.11}$$

The proof consists of two steps. In the first step, we consider the full information disclosure design that uses final award R_B^{\min} and consumes the same budget in expectation as the no disclosure design. If agents quit earlier than t_N in the full disclosure design, then the statement of the proposition follows. Otherwise, we assume by way of contradiction that $t_F \geq t_N$ and find an upper bound on the value of the intermediate award R_A . In the second step, we show that

$$\lambda p_{t_N}^F \left(R_A + \Pi(B, A, R_B^{\min}) \right) < c, \quad (\text{EC.12})$$

which contradicts our assumption that under full information disclosure with final award R_B^{\min} agents put full effort up to time t_F .

Step 1: Let \mathcal{B}^N and \mathcal{B}^F denote the total budgets allocated to awards in the designs that feature no and full information disclosure respectively. In particular, note that the budget allocated to final award R in the design with no information disclosure is equal to the following in expectation:

$$\begin{aligned} \mathbb{E}[\mathcal{B}^N] &= \int_{t=0}^{t_N} e^{-(2\lambda+r)t} 2\lambda \left(\int_{\tau=0}^{t_N-t} e^{-(\mu+\lambda+r)\tau} \left(\mu R + \lambda \frac{2\mu}{2\mu+r} R \right) d\tau + e^{-(\mu+\lambda+r)(t_N-t)} \frac{\mu}{\mu+r} R \right) dt \\ &= \frac{2\lambda\mu R}{(\mu+r)(\lambda+\mu+r)(2\mu+r)} \left(\left(1 - e^{-(2\lambda+r)t_N} \right) \frac{(\mu+r)(2\lambda+2\mu+r)}{2\lambda+r} - \frac{\lambda r e^{-(2\lambda+r)t_N} \left(e^{(\lambda-\mu)t_N} - 1 \right)}{\lambda-\mu} \right). \end{aligned} \quad (\text{EC.13})$$

Next, we obtain an upper bound on R_A by calculating the budget allocated to awards R_A, R_B^{\min} under full disclosure when agents stop at t_N . This gives an upper bound on R_A under the assumption that $t_F \geq t_N$ – note that we impose that the budgets allocated to awards are equal under these two disclosure schemes. We have

$$\begin{aligned} \mathbb{E}[\mathcal{B}^F] &\geq \int_{t=0}^{t_N} e^{-(2\lambda+r)t} 2\lambda \left(R_A + \left(\frac{\lambda}{\lambda+\mu+r} \frac{2\mu}{2\mu+r} + \frac{\mu}{\lambda+\mu+r} \right) R_B^{\min} \right) dt \\ &= \frac{2\lambda}{2\lambda+r} \left(1 - e^{-(2\lambda+r)t_N} \right) \left(R_A + \left(\frac{\lambda}{\lambda+\mu+r} \frac{2\mu}{2\mu+r} + \frac{\mu}{\lambda+\mu+r} \right) R_B^{\min} \right). \end{aligned} \quad (\text{EC.14})$$

Thus, since the two designs consume the same budget in expectation we obtain:

$$\mathbb{E}[\mathcal{B}^N] \geq \frac{2\lambda}{2\lambda+r} \left(1 - e^{-(2\lambda+r)t_N} \right) \left(R_A + \left(\frac{\lambda}{\lambda+\mu+r} \frac{2\mu}{2\mu+r} + \frac{\mu}{\lambda+\mu+r} \right) R_B^{\min} \right). \quad (\text{EC.15})$$

Finally, using Equation (EC.13) yields the following upper bound on R_A :

$$\begin{aligned} R_A &\leq \frac{(2\lambda+2\mu+r)\mu R}{(\lambda+\mu+r)(2\mu+r)} - \frac{(2\lambda+r)e^{-(2\lambda+r)t_N}}{(1-e^{-(2\lambda+r)t_N})} \frac{\mu R}{(\mu+r)(\lambda+\mu+r)(2\mu+r)} \frac{\lambda r}{\lambda-\mu} \left(e^{(\lambda-\mu)t_N} - 1 \right) \\ &\quad - R_B^{\min} \left(\frac{\lambda}{\lambda+\mu+r} \frac{2\mu}{2\mu+r} + \frac{\mu}{\lambda+\mu+r} \right). \end{aligned} \quad (\text{EC.16})$$

Step 2: In this step, we show that inequality (EC.12) holds and thus conclude that $t_{\mathcal{F}} < t_{\mathcal{N}}$. Note that by using Equation (EC.11) we can replace the right hand side of inequality (EC.12) with

$$\lambda p_{t_{\mathcal{N}}}^{\mathcal{N}} \left(q_{t_{\mathcal{N}}}^{\mathcal{N}} \Pi(B, B, R) + (1 - q_{t_{\mathcal{N}}}^{\mathcal{N}}) \frac{\mu R - c}{\mu + r} \right).$$

Our goal then is to show that

$$p_{t_{\mathcal{N}}}^{\mathcal{F}} \left(R_A + \Pi(B, A, R_B^{\min}) \right) < p_{t_{\mathcal{N}}}^{\mathcal{N}} \left(q_{t_{\mathcal{N}}}^{\mathcal{N}} \Pi(B, B, R) + (1 - q_{t_{\mathcal{N}}}^{\mathcal{N}}) \frac{\mu R - c}{\mu + r} \right).$$

First, note that $p_{t_{\mathcal{N}}}^{\mathcal{F}} = \frac{p_{t_{\mathcal{N}}}^{\mathcal{N}}(1 - q_{t_{\mathcal{N}}}^{\mathcal{N}})}{1 - p_{t_{\mathcal{N}}}^{\mathcal{N}} q_{t_{\mathcal{N}}}^{\mathcal{N}}}$. Thus, we can rewrite inequality (EC.12) as

$$(1 - q_{t_{\mathcal{N}}}^{\mathcal{N}}) \left(R_A + \Pi(B, A, R_B^{\min}) \right) < (1 - p_{t_{\mathcal{N}}}^{\mathcal{N}} q_{t_{\mathcal{N}}}^{\mathcal{N}}) \left(q_{t_{\mathcal{N}}}^{\mathcal{N}} \Pi(B, B, R) + (1 - q_{t_{\mathcal{N}}}^{\mathcal{N}}) \frac{\mu R - c}{\mu + r} \right).$$

Using Equation (EC.11) once again, we can further simplify the right hand side and rewrite the above inequality as follows:

$$(1 - q_{t_{\mathcal{N}}}^{\mathcal{N}}) \left(R_A + \Pi(B, A, R_B^{\min}) \right) < \left(q_{t_{\mathcal{N}}}^{\mathcal{N}} \Pi(B, B, R) + (1 - q_{t_{\mathcal{N}}}^{\mathcal{N}}) \frac{\mu R - c}{\mu + r} R \right) - q_{t_{\mathcal{N}}}^{\mathcal{N}} \frac{c}{\lambda}.$$

Note that

$$R_B^{\min} \left(\frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} + \frac{\mu}{\lambda + \mu + r} \right) - \Pi(B, A, R_B^{\min}) = \frac{\lambda}{\lambda + \mu + r} \frac{2\mu}{2\mu + r} R_B^{\min},$$

Also, recall that $\Pi(B, B, R) = \frac{\mu R - c}{2\mu + r}$, and

$$q_{t_{\mathcal{N}}}^{\mathcal{N}} = \frac{\lambda(e^{-\mu t_{\mathcal{N}}} - e^{-\lambda t_{\mathcal{N}}})}{\lambda e^{-\mu t_{\mathcal{N}}} - \mu e^{-\lambda t_{\mathcal{N}}}} \text{ and } R_B^{\min} = \frac{c}{\mu} \left(1 + \frac{2\mu + r}{\lambda} \right).$$

Substituting the upper bound for R_A from (EC.16), using the expressions above, and some straightforward algebra yields

$$\begin{aligned} \frac{\lambda \mu r R}{(\lambda + \mu + r)(2\mu + r)(\mu + r)} & \left(\frac{(\lambda + \mu + r)(\mu + r)}{r} + \frac{(2\lambda + r)e^{-(2\lambda + r)t_{\mathcal{N}}}}{(1 - e^{-(2\lambda + r)t_{\mathcal{N}}})} - (\lambda - \mu) \frac{e^{-(\lambda - \mu)t_{\mathcal{N}}}}{1 - e^{-(\lambda - \mu)t_{\mathcal{N}}}} \right) > \\ & \frac{c(2\mu + \lambda + r)}{(2\mu + r)} - (\lambda - \mu) \frac{e^{-(\lambda - \mu)t_{\mathcal{N}}}}{1 - e^{-(\lambda - \mu)t_{\mathcal{N}}}} \frac{c(\lambda r + r^2 + 3r\mu + 2\mu^2)}{(\mu + r)(\lambda + \mu + r)(2\mu + r)}. \end{aligned} \quad (\text{EC.17})$$

Next, we show that the term in the parenthesis in the left hand side of (EC.17) is positive, i.e.,

$$\frac{(\lambda + \mu + r)(\mu + r)}{r} + \frac{(2\lambda + r)e^{-(2\lambda + r)t_{\mathcal{N}}}}{(1 - e^{-(2\lambda + r)t_{\mathcal{N}}})} - (\lambda - \mu) \frac{e^{-(\lambda - \mu)t_{\mathcal{N}}}}{1 - e^{-(\lambda - \mu)t_{\mathcal{N}}}} > 0. \quad (\text{EC.18})$$

Ignoring term $\frac{\mu + r}{r}$ (since it is greater than one) and simplifying the above expression, we obtain

$$\frac{(2\lambda + r)(1 - e^{-(\lambda - \mu)t_{\mathcal{N}}}) - (\lambda - \mu)(1 - e^{-(2\lambda + r)t_{\mathcal{N}}})}{(1 - e^{-(2\lambda + r)t_{\mathcal{N}}})(1 - e^{-(\lambda - \mu)t_{\mathcal{N}}})} > 0.$$

The denominator is positive since we assume that $\lambda > \mu$ and thus it is enough to show that the numerator is positive as well. The latter follows since for $t_{\mathcal{N}} = 0$ the numerator is equal to 0 and its derivative with respect to $t_{\mathcal{N}}$ is positive.

Finally, note that $R \geq R_B^{\min}$ and thus the left hand side of inequality (EC.17) is minimized when $R = R_B^{\min}$. Thus, the claim follows by establishing that (EC.17) holds when we substitute R_B^{\min} for R . Straightforward algebra establishes this last step and concludes the proof. Q.E.D.

Proof of Proposition 4

First, we show that an agent's best response takes the form of a cutoff experimentation policy before time t_S . In particular, assume that agent 2 follows some strategy $\{x_{2,t}\}_{t \geq 0}$ in the absence of partial progress. Then, agent 1's optimization problem takes the following form:

$$\begin{aligned} & \max_{\{x_{1,t}\}_{t \geq 0}} \int_{t=0}^{t_S} x_{1,t} \left(\lambda p_{1,t}^S \left(q_{1,t}^S (R_A/2 + \Pi(B, B, R_B^{\min})) \right) + (1 - q_{1,t}^S) V_{1,t}(B, A) - c \right) e^{-\int_{\tau=0}^t p_{1,\tau}^S (q_{1,\tau}^S \mu + \lambda x_{1,\tau}) d\tau} e^{-rt} dt \\ & + e^{-\int_{\tau=0}^{t_S} p_{1,\tau}^S (q_{1,\tau}^S \mu + \lambda x_{1,\tau}) d\tau} (1 - p_{1,t_S}^S q_{1,t_S}^S) \int_{t=t_S}^{\infty} x_{1,t} (\lambda p_{1,t}^S \Pi(B, A, R_B^{\min}) - c) e^{-\int_{\tau=t_S}^t \lambda p_{1,\tau}^S (x_{1,\tau} + x_{2,\tau}) dt} e^{-rt} dt, \end{aligned} \quad (\text{EC.19})$$

where for any $t \leq t_S$, we have $p_{1,t}^S = p_{1,t}^{\mathcal{N}}$ and $q_{1,t}^S = q_{1,t}^{\mathcal{N}}$, whereas for $t > t_S$ we have $p_{1,t}^S = p_{1,t}^{\mathcal{F}}$.

First, note that the instantaneous payoff for agent 1 is decreasing in time before t_S (first integral in Expression (EC.19)). Thus, it is optimal for agent 1 to employ a cutoff experimentation policy before t_S . As a result, for the remainder of the proof, we only focus on cutoff experimentation policies for both agents 1 and 2. In other words, we assume that in the absence of any partial progress agent 2 sets her effort level to one up to some point $T_2 \leq t_S$ and then to zero until time t_S , i.e., we let

$$x_{2,\tau} = \begin{cases} 1 & \text{for } \tau \leq T_2 \\ 0 & \text{for } T_2 \leq \tau \leq t_S \\ x_{2,\tau} & \text{for } \tau > t_S \end{cases}.$$

To complete the proof, we show that agent 1's best response to $\{x_{2,t}\}_{t \geq 0}$ described above involves setting $x_{1,t} = 1$ for $t \leq t_S$. Note that the instantaneous payoff for agent 1 inside each integral is decreasing in time. Thus, establishing that agent 1's instantaneous payoff just before t_S is higher than her instantaneous payoff after t_S implies that if it is optimal to put any effort after t_S (in the absence of any announcement about partial progress), it is also optimal to put full effort up to t_S . Assume that agent 1 sets $x_{1,\tau} = 1$ for $\tau \leq T_1 \leq t_S$. Then, her instantaneous payoff for putting effort just before t_S is given by:

$$\lambda p_{1,t_S}^S \left(q_{1,t_S}^S (R_A/2 + \Pi(B, B, R_B^{\min})) \right) + (1 - q_{1,t_S}^S) \Pi(B, A, R_B^{\min}) - c, \quad (\text{EC.20})$$

with

$$p_{1,t_S}^S = \frac{e^{-\lambda T_1} \left(\frac{\lambda e^{-\mu t_S}}{\lambda - \mu} + e^{-\lambda T_2} \right) p_A}{1 - p_A + e^{-\lambda T_1} \left(\frac{\lambda e^{-\mu t_S}}{\lambda - \mu} + e^{-\lambda T_2} \right) p_A} \quad \text{and} \quad q_{1,t_S}^S = \frac{\lambda e^{-\mu t_S} (1 - e^{-(\lambda - \mu) T_2})}{\lambda e^{-\mu t_S} (1 - e^{-(\lambda - \mu) T_2}) + (\lambda - \mu) e^{-\lambda T_2}}.$$

On the other hand, her instantaneous payoff for putting effort just after t_S is given by:

$$(1 - p_{1,t_S}^S q_{1,t_S}^S) (\lambda p_{1,t_S^+} \Pi(B, A, R_B^{\min}) - c), \quad (\text{EC.21})$$

where $p_{1,t_S^+} = \frac{p_A e^{-\lambda(T_1+T_2)}}{1-p_A+p_A e^{-\lambda(T_1+T_2)}}$ where p_{1,t_S^+} denotes agent 1's belief just after the designer switches to full information disclosure at time t_S . Moreover, note that $\Pi(B, B, R_B^{\min}) = \frac{c}{\lambda}$ and

$$(1 - p_{1,t_S}^S q_{1,t_S}^S) p_{1,t_S^+} = p_{1,t_S}^S (1 - q_{1,t_S}^S).$$

Putting these together yields the desired result, i.e., that the instantaneous payoff for agent 1 is higher before t_S than after t_S if the designer does not announce that Stage A has been completed. Finally, note that t_S is chosen so that the instantaneous payoff for agent 1 is non-negative for $t \leq t_S$ for any $T_2 \leq t_S$. Thus, agent 1's best response to agent 2's strategy is to put full effort up to time t_S . The claim follows by noting that if both agents put full effort up to t_S their instantaneous payoff at t_S is exactly equal to zero. Thus, in the absence of any announcement by the designer they set their effort levels to zero after t_S . Q.E.D.

Proof of Proposition 5

Consider a full information disclosure design and recall that $t_{\mathcal{F}}$ denotes the time at which agents stop putting effort in Stage A if none of them has completed it. Recall also that with intermediate award R_A and final award R_B^{\min} , $t_{\mathcal{F}}$ is the solution of the following equation:

$$\lambda p_{t_{\mathcal{F}}}^{\mathcal{F}} (R_A + \Pi(B, A, R_B^{\min})) = c, \quad (\text{EC.22})$$

where $p_{t_{\mathcal{F}}}^{\mathcal{F}} = \frac{p_A e^{-2\lambda t_{\mathcal{F}}}}{p_A e^{-2\lambda t_{\mathcal{F}}} + (1-p_A)}$ is the agents' common posterior belief about the feasibility of Stage A at time $t_{\mathcal{F}}$. The proof follows by showing that t_S given by Expression (8) is such that $t_S > t_{\mathcal{F}}$. This directly implies that the designer's expected payoff is higher for the silent period design than under full information disclosure (since Proposition 4 implies that both agents will experiment with rate one until t_S). In particular, the claim follows by showing that

$$\lambda p_{t_S}^{\mathcal{F}} (R_A + \Pi(B, A, R_B^{\min})) < \lambda p_{t_S}^S (q_{t_S}^S (R_A/2 + \Pi(B, B, R_B^{\min})) + (1 - q_{t_S}^S) (R_A + \Pi(B, A, R_B^{\min}))).$$

First, recall that by definition t_S satisfies:

$$\lambda p_{t_S}^S (q_{t_S}^S (R_A/2 + \Pi(B, B, R_B^{\min})) + (1 - q_{t_S}^S) (R_A + \Pi(B, A, R_B^{\min}))) = c.$$

Also note that since $\lambda (R_A/2 + \Pi(B, B, R_B^{\min})) \geq c$, by substituting it in the right hand side of the above equation, and rearranging terms we obtain the following inequality:

$$\frac{p_{t_S}^S q_{t_S}^S (1 - q_{t_S}^S)}{1 - p_{t_S}^S q_{t_S}^S} (R_A + \Pi(B, A, R_B^{\min})) < q_{t_S}^S (R_A/2 + \Pi(B, B, R_B^{\min})). \quad (\text{EC.23})$$

After some algebra we can rewrite inequality (EC.23) as:

$$\frac{p_{t_S}^S (1 - q_{t_S}^S) (R_A + \Pi(B, A, R_B^{\min}))}{1 - p_{t_S}^S q_{t_S}^S} < p_{t_S}^S (q_{t_S}^S (R_A/2 + \Pi(B, B, R_B^{\min})) + (1 - q_{t_S}^S) (R_A + \Pi(B, A, R_B^{\min}))). \quad (\text{EC.24})$$

Finally, noting that

$$p_{t_S}^{\mathcal{F}} = \frac{p_{t_S}^S (1 - q_{t_S}^S)}{1 - p_{t_S}^S q_{t_S}^S},$$

implies that inequality (EC.24) can be rewritten as

$$\lambda p_{t_S}^{\mathcal{F}} (R_A + \Pi(B, A, R_B^{\min})) < \lambda p_{t_S}^S (q_{t_S}^S (R_A/2 + \Pi(B, B, R_B^{\min})) + (1 - q_{t_S}^S) (R_A + \Pi(B, A, R_B^{\min}))),$$

which completes the claim. Q.E.D.

Proof of Proposition 7

The proof relies on showing that the best response of agent 1 to any strategy $\{x_{2,t}\}_{t \geq 0}$ from agent 2 is a cutoff experimentation policy. To this end, consider agent 1's optimization problem:

$$\max_{\{x_{1,t}\}_{t \geq 0}} \int_{t=0}^{\infty} \left(\lambda x_{1,t} p_{1,t}^{\phi} V_{1,t}(B) + p_{1,t}^{\phi} q_{1,t}^{\phi} \phi \Pi(A, B, R_B^{\min}) - c x_{1,t} \right) \Psi_t e^{-\int_{\tau=0}^t p_{1,\tau}^{\phi} \lambda x_{1,\tau} d\tau} e^{-rt} dt, \quad (\text{EC.25})$$

where (with some abuse of notation) $V_{1,t}(B)$ denotes the continuation value of agent 1 at time t when she has completed Stage A and the designer has not disclosed any information about agent 2's progress. Also, we let Ψ_t denote the probability that the designer has not disclosed any information about agent 2's progress, i.e., Ψ_t is given by the following expression

$$\Psi_t = e^{-\int_{\tau=0}^t p_{1,\tau}^{\phi} q_{1,\tau}^{\phi} \mu d\tau} e^{-\int_{\tau=0}^t p_{1,\tau}^{\phi} q_{1,\tau}^{\phi} \phi d\tau}. \quad (\text{EC.26})$$

Finally, note that since $\Pi(A, B, R_B^{\min}) = 0$ we can ignore the second term in the parenthesis and rewrite optimization problem (EC.25) as follows:

$$\max_{\{x_{1,t}\}_{t \geq 0}} \int_{t=0}^{\infty} x_{1,t} \left(\lambda p_{1,t}^{\phi} V_{1,t}(B) - c \right) \Psi_t e^{-\int_{\tau=0}^t p_{1,\tau}^{\phi} \lambda x_{1,\tau} d\tau} e^{-rt} dt. \quad (\text{EC.27})$$

Next, we show that the coefficient of $x_{1,t}$, i.e., expression $(\lambda p_{1,t}^{\phi} V_{1,t}(B) - c) \Psi_t e^{-\int_{\tau=0}^t p_{1,\tau}^{\phi} \lambda x_{1,\tau} d\tau} e^{-rt} dt$, is decreasing in t . This, in turn, implies that agent 1's optimal strategy is to follow a cutoff experimentation policy, i.e., put full effort up to some time and, in the absence of progress or positive news from the designer, stop exerting effort altogether. Note that since $e^{-\int_{\tau=0}^t p_{1,\tau}^{\phi} \lambda x_{1,\tau} d\tau} e^{-rt} dt$ is decreasing in t , it is enough to show that $(\lambda p_{1,t}^{\phi} V_{1,t}(B) - c) \Psi_t$ is also decreasing in t . In addition, $V_{1,t}(B)$ is given as follows:

$$V_{1,t}(B) = q_{1,t}^{\phi} \left(R_A/2 + \Pi(B, B, R_B^{\min}) \right) + (1 - q_{1,t}^{\phi}) V_{1,t}(B, A), \quad (\text{EC.28})$$

where the first term is the continuation value for agent 1 conditional on agent 2 having already completed Stage B whereas the second term is the continuation value for agent 1 when agent 2 has not completed Stage A up to t . For the rest of the analysis, it is helpful to further split $V_{1,t}(B, A)$ into $V_{1,t}^{RA}(B, A)$ and $V_{1,t}^{RB}(B, A)$ that denote the expected payoff for agent 1 from awards R_A and R_B respectively (note that agents split award R_A if they both complete Stage A before the designer discloses the progress). Thus, the claim follows if the derivative below is negative

$$\frac{d}{dt} \left(\lambda p_{1,t}^\phi \left(q_{1,t}^\phi (R_A/2 + \Pi(B, B, R_B^{\min})) \right) + (1 - q_{1,t}^\phi) \left(V_{1,t}^{RA}(B, A) + V_{1,t}^{RB}(B, A) \right) \right) - c \Big) \Psi_t < 0. \quad (\text{EC.29})$$

First, we obtain expressions for the following derivatives $\dot{q}_{1,t}^\phi$, $\dot{p}_{1,t}^\phi$, $\dot{\Psi}_t$, $\dot{V}_{1,t}^{RA}(B, A)$, and $\dot{V}_{1,t}^{RB}(B, A)$. In particular, we have

$$\dot{p}_{1,t}^\phi = -p_{1,t}^\phi(1 - p_{1,t}^\phi)(q_{1,t}^\phi\phi + q_{1,t}^\phi\mu + \lambda x_{1,t}), \text{ and} \quad (\text{EC.30})$$

$$\dot{q}_{1,t}^\phi = (1 - q_{1,t}^\phi) (\lambda x_{2,t} - q_{1,t}^\phi\phi - q_{1,t}^\phi\mu). \quad (\text{EC.31})$$

Furthermore, from Expression (EC.26) we obtain

$$\dot{\Psi}_t = -\Psi_t p_{1,t}^\phi q_{1,t}^\phi (\mu + \phi). \quad (\text{EC.32})$$

Continuation values $V_{1,t}^{RA}(B, A)$ and $V_{1,t}^{RB}(B, A)$ are differentiable and their derivatives can be obtained by noting that

$$V_{1,t}^{RA}(B, A) = \left(\mu + \lambda x_{2,t}/2 + \phi \right) R_A dt + (1 - \mu dt - \lambda x_{2,t} dt - \phi dt) V_{1,t+dt}^{RA}(B, A),$$

and

$$\begin{aligned} V_{1,t}^{RB}(B, A) &= \left(\mu R_B^{\min} + \lambda x_{2,t} \Pi(B, B, R_B^{\min}) + \phi \Pi(B, A, R_B^{\min}) - c \right) dt \\ &\quad + (1 - \mu dt - \lambda x_{2,t} dt - \phi dt - r dt) V_{1,t+dt}^{RB}(B, A). \end{aligned}$$

In particular, we have

$$\dot{V}_{1,t}^{RA}(B, A) = (\mu + \lambda x_{2,t} + \phi) V_{1,t}^{RA}(B, A) - (\mu R_A + \lambda x_{2,t} R_A/2 + \phi R_A), \quad (\text{EC.33})$$

and

$$\dot{V}_{1,t}^{RB}(B, A) = (\mu + \lambda x_{2,t} + \phi + r) V_{1,t}^{RB}(B, A) - (\mu R_B^{\min} + \lambda x_{2,t} \Pi(B, B, R_B^{\min}) + \phi \Pi(B, A, R_B^{\min}) - c). \quad (\text{EC.34})$$

Using the expressions for $\dot{q}_{1,t}^\phi$, $\dot{p}_{1,t}^\phi$, $\dot{\Psi}_t$, $\dot{V}_{1,t}^{RA}(B, A)$, and $\dot{V}_{1,t}^{RB}(B, A)$ from Equations (EC.30)–(EC.34) we can rewrite inequality (EC.29) as:

$$\Psi_t \left(\lambda p_{1,t}^\phi \left((1 - q_{1,t}^\phi) (\lambda x_{2,t} - q_{1,t}^\phi\phi - q_{1,t}^\phi\mu) \right) \left(R_A/2 + \Pi(B, B, R_B^{\min}) - V_{1,t}^{RA}(B, A) - V_{1,t}^{RB}(B, A) \right) \right)$$

$$\begin{aligned}
& + (1 - q_{1,t}^\phi) \left((\mu + \lambda x_{2,t} + \phi) V_{1,t}^{RA}(B, A) - (\mu R_A + \lambda x_{2,t} R_A / 2 + \phi R_A) \right. \\
& \left. + (\mu + \lambda x_{2,t} + \phi + r) V_{1,t}^{RB}(B, A) - (\mu R_B^{\min} + \lambda x_{2,t} \Pi(B, B, R_B^{\min}) + \phi \Pi(B, A, R_B^{\min}) - c) \right) \\
& \left. - p_{1,t}^\phi (1 - p_{1,t}^\phi) (q_{1,t}^\phi \phi + q_{1,t}^\phi \mu + \lambda x_{1,t}) \lambda V_{1,t}(B) \right) \\
& - \Psi_t p_{1,t}^\phi q_{1,t}^\phi (\mu + \phi) \left(\lambda p_{1,t}^\phi (q_{1,t}^\phi (R_A / 2 + \Pi(B, B, R_B^{\min}))) + (1 - q_{1,t}^\phi) (V_{1,t}^{RA}(B, A) + V_{1,t}^{RB}(B, A)) \right) - c < 0.
\end{aligned} \tag{EC.35}$$

We can further simplify Expression (EC.35) by canceling out all terms involving $x_{2,t}$. Also, since the coefficient of $x_{1,t}$ is negative, it is sufficient to show the following inequality:

$$\begin{aligned}
& \Psi_t \left(\lambda p_{1,t}^\phi \left((1 - q_{1,t}^\phi) q_{1,t}^\phi (-\phi - \mu) (R_A / 2 + \Pi(B, B, R_B^{\min})) - V_{1,t}^{RA}(B, A) - V_{1,t}^{RB}(B, A) \right) \right. \\
& \left. + (1 - q_{1,t}^\phi) \left((\mu + \phi) V_{1,t}^{RA}(B, A) - (\mu R_A + \phi R_A) \right) \right. \\
& \left. + (\mu + \phi + r) V_{1,t}^{RB}(B, A) - (\mu R_B + \phi \Pi(B, A) - c) \right) - p_{1,t}^\phi (1 - p_{1,t}^\phi) q_{1,t}^\phi (\mu + \phi) \lambda V_{1,t}(B) \\
& - \Psi_t p_{1,t}^\phi q_{1,t}^\phi (\mu + \phi) \left(\lambda p_{1,t}^\phi (q_{1,t}^\phi (R_A / 2 + \Pi(B, B, R_B^{\min}))) + (1 - q_{1,t}^\phi) (V_{1,t}^{RA}(B, A) + V_{1,t}^{RB}(B, A)) \right) - c < 0.
\end{aligned}$$

Next, replacing $V_{1,t}(B)$ using (EC.28) and canceling out common terms yields

$$\begin{aligned}
& \Psi_t \left(\lambda p_{1,t}^\phi \left((1 - q_{1,t}^\phi) (-q_{1,t}^\phi \phi - q_{1,t}^\phi \mu) (R_A / 2 + \Pi(B, B, R_B^{\min})) + (1 - q_{1,t}^\phi) \left((\mu + \phi) V_{1,t}^{RA}(B, A) - (\mu R_A + \phi R_A) \right) \right. \right. \\
& \left. \left. + (\mu + \phi + r) V_{1,t}^{RB}(B, A) - (\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c) \right) - \lambda p_{1,t}^\phi q_{1,t}^\phi (\mu + \phi) \left(q_{1,t}^\phi (R_A / 2 + \Pi(B, B, R_B^{\min})) \right) \right) \\
& + \Psi_t p_{1,t}^\phi q_{1,t}^\phi (\mu + \phi) c < 0.
\end{aligned}$$

Finally, since the coefficients of $V_{1,t}^{RA}(B, A)$ and $V_{1,t}^{RB}(B, A)$ are positive, we can replace them with their respective upper bounds (obtained when agent 2 does not put effort towards completing Stage A unless the designer announces that agent 1 has already completed it). In other words, we bound $V_{1,t}^{RA}(B, A)$ and $V_{1,t}^{RB}(B, A)$ by

$$\begin{aligned}
V_{1,t}^{RA}(B, A) & \leq R_A, \quad \text{and} \\
V_{1,t}^{RB}(B, A) & \leq \int_{\tau=0}^{\infty} e^{-\tau(\phi+\mu+r)} (\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c) \\
& = \frac{1}{\mu + \phi + r} (\mu R_B^{\min} + \phi \Pi(B, A, R_B^{\min}) - c).
\end{aligned} \tag{EC.36}$$

Substituting the upper bounds obtained in (EC.36) and canceling out common terms we obtain:

$$\Psi_t \left(\lambda p_{1,t}^\phi \left((-q_{1,t}^\phi \phi - q_{1,t}^\phi \mu) (\Pi(B, B, R_B^{\min}) + R_A / 2) \right) \right) + \Psi_t p_{1,t}^\phi (q_{1,t}^\phi \mu + q_{1,t}^\phi \phi) c < 0,$$

and since $\Pi(B, B, R_B^{\min}) = \frac{c}{\lambda}$ (recall that $R_B^{\min} = c/\mu(1 + (2\mu + r)/\lambda)$) we get:

$$- \Psi_t \left(p_{1,t}^\phi q_{1,t}^\phi (\phi + \mu) \left(\lambda \left(\frac{c}{\lambda} + R_A/2 \right) - c \right) \right) < 0, \quad (\text{EC.37})$$

which holds and, thus, completes the proof.

As a final remark, note that Proposition 2 follows from the same line of arguments when setting $\phi = R_A = 0$ and allocating the entire budget to the award for completing Stage B . Q.E.D.