Dynamic Development Contests

Sina Khorasani

Rady School of Management, University of California San Diego, skhorasani@ucsd.edu

Ersin Körpeoğlu

School of Management, University College London, e.korpeoglu@ucl.ac.uk

Vish Krishnan

Rady School of Management, University of California San Diego, vkrishnan@ucsd.edu

Public, private, and not-for-profit organizations find advanced technology and product development projects challenging to manage due to the time and budget pressures, and turn to their development partners and suppliers to address their development needs. We study how dynamic development contests with enriched rank-based incentives and carefully-tailored information design can help these organizations outsource their development projects at the minimum project lead time by stimulating competition among suppliers. We show that an organization can use dynamically-adjusted flexible rewards to achieve the absolute minimum expected project lead time at a significantly lower cost than a fixed-reward policy. Importantly, our flexible-reward policy pays the absolute minimum expected reward (i.e., achieves the first best). We further study the case where the organization does not have sufficient budget to offer a reward that attains the absolute minimum expected lead time. We propose that in this case, the organization can dynamically increase the contest reward until its budget constraint binds and then use information sharing as a strategic tool to incentivize suppliers. Specifically, we propose an easy-to-implement random-update policy where the organization periodically monitors the status of suppliers at random times and immediately discloses any partial progress. We show that such a random-update policy outperforms other canonical information disclosure strategies. Our results indicate that dynamic rewards and strategic information disclosure are powerful tools to help organizations outsource their development needs swiftly and cost effectively.

Key words: Contest, Development, Information Design, Innovation, Technology, Tournament

1. Introduction

Organizations worldwide face the challenges of developing advanced technologies and products under time and budget pressures. Private, public, and not-for-profit organizations aim to develop technologies with the help of Tier 1 suppliers to tackle complex problems that require rigorous work with significant uncertainty and multiple major milestones. Examples include the US Department of Defense racing against other global powers to build hypersonic missiles by leveraging suppliers like Lockheed Martin and Boeing; or Global Alliance of Vaccine and Immunization (GAVI) racing to develop vaccines for deadly and debilitating infectious diseases such as pneumococcal meningitis

or COVID-19 by tapping key competing pharmaceutical suppliers such as Pfizer and Astrazeneca.[1] In this paper, we study how such organizations can go beyond contractual development and resort to a "dynamic development contest" approach - by utilizing advanced elements of contests such as rank-based incentives and strategic information disclosure to stimulate development effort from competing suppliers to minimize project lead time while keeping the incentive budget in check.

There are three distinguishing aspects of a dynamic development contest:

1. **Lead-Time-Based Competition.** The development organization (hereafter, "principal") is engaged in launching a carefully-specified product (e.g., a hypersonic missile) or solving a well-defined problem (e.g., developing a vaccine with a certain efficacy) and is focused on minimizing the highly uncertain project lead time while keeping the cost of incentives in check. Lead time minimization has not been a major priority for the more exploratory innovation contests, where the focus is to obtain the highest quality innovative solutions for a broadly defined problem.

2. **Expert-Sourcing.** Due to complexity of problems and the enormity of investments necessary, the principal in a development contest typically works with a few strategic suppliers (hereafter, "agents") who have a proven expertise in the specific technologies and make significant resource investments - in contrast to crowd-sourced contests that aim to attract a larger pool of solvers making relatively low out of pocket investments. While such expert sourcing is involved in its own ways, it also offers the principal some additional degrees of freedom with the suppliers/agents, which can be put to good use as characterized in this paper to meet the principal's objectives of lead time minimization under budget constraints.

3. **Difficult Multi-phase Problem.** The development setting's complexity means it is subdivided into multiple milestones or stages (such as preclinical and clinical in the case of vaccine development) to achieve the development goals allowing the organization to interact with the suppliers at the milestone points.

Development contests also share similarities with other types of contests in their tendency to engage in and benefit from contestants. We extend and adapt the contest paradigm to development settings and contribute additional novel features such as enriched rank-based incentives and carefully-tailored information design to stimulate competitive effort from suppliers.

Although development contests have the potential for significant economic and societal impact,[2] they are not without challenges. First, as agents tackle a difficult problem, there is a danger that some agents will lose interest in the absence of sufficient progress. Thus, an effective mechanism

---

[1] In GAVI development contests, a small set of vaccine suppliers race to develop vaccines of acceptable efficacy at the shortest lead time to win a reward in the form of a large supply contract with guaranteed prices (GAVI 2020).

[2] GAVI reports that pneumococcal vaccine developed as a result of its development contest has prevented 700,000 children's deaths in 60 developing countries over the last decade (GAVI 2020).

should dynamically keep agents' incentives alive to minimize the development lead time. Second, major development projects require substantial investment from agents so the cost of incentivizing effort from agents is significant. Considering that cost of incentives is usually covered by taxpayers' money or donations, an effective mechanism should avoid overpaying agents. Therefore, we aim to understand how contest organizers can design such contests for maximum efficiency – to minimize lead time of development at the minimum cost of incentives.

To capture the key characteristics of a dynamic development contest, we build a parsimonious model where two agents compete to complete a two-stage development project by exerting costly effort over a continuous time frame. Successful completion of a stage (success) for an agent arrives at a random point in time where the rate of arrival for each success increases with the agent's effort. An agent's success is not observable by another agent so it is up to the principal whether and when to share this information. The contest ends when one of the agents achieves two successes at which point this winning agent is given a pre-determined reward. The principal commits to a *reward schedule* about how this reward changes over time and an *information disclosure policy* which specifies how the principal will disclose information throughout the contest (e.g., no information sharing, full information sharing) at the beginning of the contest.

We first focus on the problem of a principal with no budget constraint. We establish that such a principal can utilize a fixed-reward dynamic contest that does not change the reward over time to solicit agents to exert their best efforts and attain the absolute minimum expected lead time. Yet, we also find that such a principal significantly overpays agents. Specifically, by utilizing a carefully designed flexible-reward schedule that increases the reward over time, the principal can still achieve the absolute minimum expected lead time with 20% less cost of incentives/reward on average, which could mean substantial savings in development settings that can cost hundreds of millions of dollars. Better yet, a principal who makes use of information disclosure can simplify the reward schedule tremendously while reaping all the benefits of the flexible-reward schedule. Specifically, if the principal commits to keep each agent fully apprised of any partial progress (i.e., an agent achieving the first success), the optimal flexible-reward schedule consists of just two levels of rewards where the principal sets a guaranteed reward amount upfront with the promise of increasing the reward amount if multiple agents achieve partial progress. Importantly, we show that the flexible-reward schedule pays the absolute minimum expected reward (i.e., the first-best).

We next analyze a budget-constrained principal that cannot set a sufficiently large reward to achieve the absolute minimum expected lead time. In this case, we first focus on how the principal can utilize information as a strategic commodity by considering a fixed-reward contest that uses all of the reward budget. By harnessing a dynamic Bayesian persuasion approach, we propose a simple to implement "random-update" policy in which the principal commits to not monitoring the status

of agents for a pre-determined initial time window; and then periodically monitoring the status of agents at random times and immediately disclosing any partial progress. We show that this easy-to-implement random-update policy achieves a shorter expected lead time than any canonical information disclosure policy such as no information disclosure, full information disclosure (e.g., Halac et al. 2017, Mihm and Schlapp 2019), cyclic information disclosure with periodic updates every fixed periods of time (e.g., Bimpikis et al. 2019, Ely et al. 2020), and deterministic-delay policy that commits to share partial progress after a fixed delay (e.g., Ely 2017). We then find that by utilizing a flexible-reward policy during the initial no-monitoring period, the principal can reduce the cost of incentives without hindering project lead time. Our results indicate that enriched rank-based incentives and carefully-tailored information design can be powerful tools to incentivize development efforts of suppliers without overpaying them. We now commence with a review of related literature before presenting the model and the results.

2. Related Literature

Our paper is related to three streams of literature: (i) innovation contests with no information sharing, (ii) contests with one-time or dynamic information sharing, and (iii) Bayesian persuasion.

Our work has some connection to innovation contest literature pioneered by Taylor (1995) and Terwiesch and Xu (2008) who study how many participants to let in a contest. In a more general framework, Ales et al. (2017, 2020) derive conditions for the optimality of winner-takes-all and open-entry contests. Extensions to heterogeneous agents, internal innovation contests, and multiple attributes are considered by Körpeoğlu and Cho (2018), Nittala and Krishnan (2016), and Hu and Wang (2020), respectively. More recent work investigates procurement decisions (Chen et al. 2018), participation (Stouras et al. 2019), duration (Korpeoglu et al. 2020), supplier collaboration (Shalpegin et al. 2020), teamwork (Candoğan et al. 2020), and the impact of running parallel contests (Körpeoğlu et al. 2017, Stouras et al. 2020). While this literature has provided valuable insights on innovation contests that aim to crowdsource the best-quality solution to a problem, we are interested in analyzing a dynamic development contest that differs from innovation contests in three key properties discussed in §1. These distinct properties allow us to consider new degrees of freedom available to the development organization including flexible rewards and information design due to their dynamic strategic interactions with a small number of development partners. Moreover, we study the design of incentives in a dynamic multi-stage framework to show how an organization can encourage its suppliers and development partners to reduce project lead time by exerting their best effort throughout the contest without overpaying them.

To study the role of information sharing in contests, a relatively recent stream of research has been focusing on how to provide interim performance feedback in a two-period, two-agent framework (e.g., Aoyagi 2010, Ederer 2010, Goltsman and Mukherjee 2011). Mihm and Schlapp (2019)

applies this framework to innovation contests and adds to this stream by considering private disclosure policy in addition to full and no-information disclosure policies. Schlapp and Mihm (2018) build on this work and prove that, in a wide class of feedback policies, there can be no feedback policy that outperforms pre-committed truthfulness. Private feedback in the form of screening intermediate submissions of $N \geq 2$ solvers is considered by Khorasani et al. (2020) who show that different contest environments require different ways of balancing screening specificity and sensitivity. Our continuous-time framework enables us to add to this literature by providing insights on how an organization should choose the rate and timing of information disclosure, rather than giving a one-time interim feedback, in order to dynamically manage incentives of suppliers.

Our model is closer to a few papers that study dynamic information disclosure in contests using a continuous-time framework. Halac et al. (2017) consider an experimentation contest where agents compete to obtain a single-stage innovation whose feasibility is initially unknown and the principal wishes to maximize the probability of obtaining the innovation. They show that in a setting with full information disclosure, a winner-takes-all policy dominates others. Building on this framework, Bimpikis et al. (2019) study a two-stage winner-takes-all contest where two agents compete and the feasibility of the first stage is initially unknown. They show that full information disclosure after an initial silent period may dominate no, full, or constant probabilistic disclosure. They also explore a setting with no uncertainty in the first stage and show that a cyclic information disclosure sometimes dominates no or full disclosure. Ely et al. (2020) derive optimal effort-maximizing contests where agents exert effort to obtain a single breakthrough. They propose a cyclic feedback policy that informs agents about peer success at the end of each fixed-length cycle. We contribute to this scant literature in two ways. First, we show the value of flexible rewards in contests which if designed carefully can help the principal achieve the absolute minimum expected lead time by paying the absolute minimum expected reward (i.e., the first-best). This result is important in any development setting where time and budget are of the essence. Second, we highlight the value of strategic probabilistic disclosure in contests and propose a random-update policy that *always* dominates other canonical disclosure policies including the ones studied by the above papers.

Lastly, our work is related to the growing literature of Bayesian persuasion pioneered by Rayo and Segal (2010) and Kamenica and Gentzkow (2011). The bulk of the work on Bayesian persuasion focuses on static information design where principal shares information with each agent only once (e.g., Rayo 2013, Bergemann and Morris 2019, Kamenica 2019). Indeed, there is a growing literature that applies this static framework to operational problems (see Küçükgül et al. 2019, deVericourt et al. 2020, and references therein). Recently, Ely (2017) introduces a dynamic persuasion mechanism where the principal dynamically shares information with agent(s) based on

an exogenously given state of a stochastic process.[3] Inspired by this approach, we construct our random-update policy, and show that it improves upon other canonical information disclosure policies. Importantly, our framework is much more complex because the state of the stochastic process that the principal bases its information design on is no longer exogenous. Instead, it is endogenously determined by actions of both agents as well as any prior information shared with agents. Characterizing a well-performing information disclosure policy in this setting is not an easy endeavor as pointed out in the survey by Horner and Skrzypacz (2017).

3. Model Development

Consider a setting where an organization ("principal") aims to incentivize a small group of expert firms ("agents") to complete a difficult multi-stage task as fast as possible by rewarding the agent who completes the task (i.e., all stages) first. Here we are interested in analyzing the impact of partial progress (see §1 for examples) on agents' incentives, so as common in the related literature (e.g., Bimpikis et al. 2019, Mihm and Schlapp 2019), we take the minimal model with two stages and two agents $\{i, -i\}$. This setting fits our dynamic development contest framework well because such contests feature a small number of expert suppliers and a few major milestones (see §1). Time (indexed by t) runs continuously, and the contest can last over a potentially infinite horizon.

As is standard in the contest literature reviewed in §2, we consider a winner-takes-all contest where the first agent to complete both stages wins the contest and is given a reward R_t. Unlike the standard contest framework where this reward is fixed (i.e., $R_t = R$), we consider a more general rank-based mechanism where the reward can potentially change over time. We refer to the "standard" contest mechanism as a *fixed-reward* contest and ours as a *flexible-reward* contest.

For an agent, successful completion of a stage (hereafter, "success") arrives with a Poisson process, and the agent can boost the arrival rate by exerting costly effort. Specifically, agent i who has achieved $k \in \{0, 1\}$ successes while her opponent has achieved $l \in \{0, 1\}$ successes (if known by agent i) privately chooses effort $x_{k,l,t}^i \in [0, 1]$ at each instant t with an instantaneous cost $cx_{k,l,t}^i$ for a constant $c > 0$, and a success in a stage may arrive with a Poisson process with instantaneous probability $\lambda x_{k,l,t}^i$, where λ is the "achievability" parameter that is inversely proportional to how difficult a stage is. (Our analysis and insights hold if we assume different Poisson arrival rates for different stages of the contest.) The contest ends upon the arrival of the second success (i.e., completion of the second stage) for an agent at any time t and the winner receives the current posted reward R_t. We denote by T the random date at which the contest ends.[4]

---

[3] The key idea in dynamic Bayesian persuasion is to examine how an informed principal can persuade a set of players over time to take desirable actions (in our case, exert their best effort) by influencing their beliefs. The common theme in all such applications is that no information manipulation is required in states when incentives are sufficient and only the bare minimum information is needed in states where incentives are misaligned.

[4] Although the contest can potentially run over an infinite horizon, we show in §4 and §5 that in our model the principal can always design a contest where the second success is attainable in finite time (i.e., $T < \infty$).

Denote by $V^i_{k,l,t}$ the expected utility (hereafter "continuation payoff") of agent i who has achieved $k \in \{0,1\}$ successes while her opponent has achieved $l \in \{0,1\}$ successes (if known by agent i) at any moment t. Then at any time t, each agent i anticipates the efforts of her opponent and chooses her effort levels from time t onward to maximize her expected utility

$$V^i_{k,l,t} = \max_{x^i_{k,l,\tau}} \mathbb{E}\left[R_T.\mathbb{1}_{\{i\ wins\}} - \int_t^T c x^i_{k,l,\tau} d\tau \right]. \tag{1}$$

As is common in the contest literature, we assume that all parties are risk-neutral (e.g., Ales et al. 2017, Halac et al. 2017, Bimpikis et al. 2019, Ely et al. 2020). We also assume that agents do not discount time because considering discounting complicates the expressions without providing any new insights (e.g., Halac et al. 2017, Mihm and Schlapp 2019, Ely et al. 2020). However, in §B.1 of Appendix B we show how our results can be generalized to a setting where all parties discount future payoffs. Also, we normalize the agents' outside option to zero without loss of generality.

Consistent with the literature on dynamic contests with information disclosure (e.g., Halac et al. 2017, Bimpikis et al. 2019, Ely et al. 2020), we assume that successful completion of a stage by an agent is only observable to that agent and the principal, and that only the principal can credibly disclose information about the status of agents' progress (i.e., whether each agent is in the first or second stage). This assumption is sensible given that an agent can easily misrepresent her partial progress (i.e., success in the first stage) to other agents. As a result, the only party who can credibly confirm the completion of a stage is the principal.[5] To ensure that agents can interpret the presence (or lack) of any information, the principal specifies its information disclosure policy to agents at the outset of the contest (e.g., Halac et al. 2017, Bimpikis et al. 2019, Ely et al. 2020). We consider and analyze several mainstream information disclosure policies including full information disclosure where the principal commits to disclose any success upon its arrival, no information disclosure where the principal does not use any information sharing, cyclic information disclosure where the principal stays silent during fixed-length cycles and discloses full information at the end of each cycle, and more strategic disclosure policies with deterministic or stochastic delay. Note that whenever we consider an information disclosure policy where an agent does not know her opponent's state, we drop the corresponding index l (e.g., $x^i_{k,l,t}$ becomes $x^i_{k,t}$).

The principal aims to minimize the expected lead time of the contest while also minimizing the reward necessary to achieve this goal. The standard approach in the contest literature is to assume that the principal can boil down agents' performance and reward components into a single dimensional profit function. Although this assumption makes sense in settings where agents

[5] As we show in §4.2, an agent always prefers her partial progress (i.e., success in the first stage) to be disclosed to her opponent because her opponent is discouraged by this information. Therefore, there is no cause for an agent to conceal partial progress from the principal. For the same reason, an agent has an incentive to falsely disclose partial progress to her opponent so without the approval of the principal, such information is not credible.

compete in solution quality, generating such a single dimensional profit function may be hard to achieve when agents compete in time. Thus, instead of assuming such a single dimensional profit function, we take a lexicographic approach that focuses on expected lead time first and expected reward second. Specifically, we check if the principal can achieve the absolute minimum expected lead time $\underline{T}$ by inducing both agents to exert full effort throughout the contest. If $\underline{T}$ is achievable, we aim to find a reward schedule that yields $\underline{T}$ at the minimum expected reward by solving

$$\min_{R_t} \mathbb{E}\left[R_T.\mathbb{1}_{\{i \ or \ -i \ wins \ at \ T\}}\right] \ \text{s.t.} \ \mathbb{E}\left[T.\mathbb{1}_{\{i \ or \ -i \ wins \ at \ T\}}\right] = \underline{T}. \tag{2}$$

Interestingly, we show in §4 that $\underline{T}$ is achievable if and only if the principal has sufficient funds. If the principle is budget-constrained (formally defined in §5) with a low reward budget $\overline{R}$, then it is not possible to achieve $\underline{T}$. We cover this case in §5, where we compare different information disclosure policies to find the one that yields the minimum expected lead time $\mathbb{E}\left[T.\mathbb{1}_{\{i \ or \ -i \ wins \ at \ T\}}\right]$ by using the whole reward budget. We study a principal with sufficient funds in §4 and a budget-constrained principal in §5. All proofs are presented in the Appendix.

4. A Principal with Sufficient Funds

In §4.1, we present a benchmark (first-best) under which the principal achieves the absolute minimum expected lead time at the lowest possible cost by assuming observable and contractible effort. In the following sections, we use this benchmark to measure the performance of our contest mechanisms with unobservable effort. In §4.2, we characterize the optimal fixed-reward contest under full information disclosure. In §4.3 and §4.4, we derive optimal flexible-reward contests under full information disclosure and no-information disclosure, respectively.

4.1. First-Best Contract with Observable Effort

As a form of benchmark, we first identify the absolute minimum reward the principal should give to achieve the absolute minimum expected lead time. Consider a case in which the principal can observe agents' efforts and specify the effort each agent will exert as long as it is individually rational for the agent (i.e., the agent's expected utility when exerting the designated effort is higher than her outside option normalized to zero). In this case, the principal does not need to run a contest; instead, the principal can pay a sufficient compensation to cover the cost of agents' effort by offering an individually rational contract that induces each agent to exert full effort for all t until one agent achieves two successes. This contract achieves the absolute minimum expected lead time $\underline{T}$. Thus, the principal's objective boils down to offering a contract which pays the smallest expected compensation to agents such that both agents find it individually rational to exert full effort. It is easy to verify that the expected duration of any such contract (i.e., the absolute minimum expected lead time of the task), when both agents exert full effort at all times, is $5/(4\lambda)$ unit of time as

derived in (11) in the Appendix. As a result, if the principal offers each agent $5c/(4\lambda)$, then it will be individually rational for agents to accept such a contract and exert full effort at all times because such effort will make agents' (*ex-ante*) expected utility equal to their outside option. Thus, the minimum required compensation to agents to achieve full effort can be calculated by multiplying $5c/(4\lambda)$ with 2. Proposition 1 formally states this result.

PROPOSITION 1. *There exists an individually rational "first-best" contract that induces full effort at all times with the minimum required compensation of $5c/(2\lambda)$ to agents.*

Proposition 1 characterizes the lower bound for the budget required to induce both agents to exert full effort at all times. Yet, this lower bound is achieved under the assumption that agents' efforts are observable, which is rarely the case in practice. Thus, in the remainder of this section, we seek to understand how much reward is needed to induce both agents to exert their best efforts until one agent achieves two successes by running a development contest without observing agents' efforts and compare the result with the first-best contract.

4.2. Full Information Disclosure with Fixed Reward

In this section, we analyze a contest with a fixed reward R where the principal commits to disclose any success upon its arrival. In other words, each agent is kept apprised of the progress of her opponent at any instant. We analyze the agent's problem by moving backward on the state of the game where the states are defined by the number of successes of the agents. If both agents have already achieved one success, agent i's continuation payoff from any time t onward is given by:

$$V_{1,1,t}^i = \max_{x_{1,1,\tau}^i} \int_t^\infty x_{1,1,\tau}^i (\lambda R - c) e^{-\int_t^\tau \lambda(x_{1,1,s}^i + x_{1,1,s}^{-i})ds} d\tau. \tag{3}$$

To understand the above expression, note that if agent i chooses effort $x_{1,1,\tau}^i$ during interval $(\tau, \tau + dt)$, she incurs a cost $cx_{1,1,\tau}^i dt$ and if a success arrives with probability $\lambda x_{1,1,\tau}^i dt$, she enjoys the fixed reward R, or if her opponent obtains the second success, the contest ends and she gets zero. Furthermore, $e^{-\int_t^\tau \lambda(x_{1,1,s}^i + x_{1,1,s}^{-i})ds}$ captures the probability that none of the agents have obtained the second success by time τ. Now suppose the principal aims to induce both agents to exert full effort at all times to attain the absolute minimum expected lead time $\underline{T}$. From (3), the continuation payoff of each agent i is given by $V_{1,1,t}^i = \frac{1}{2}(R - \frac{c}{\lambda})$.

Next, consider the state of the game with a leader (an agent with one success) and a laggard (an agent with no success). The laggard's continuation payoff from any time t onward is given by:

$$V_{0,1,t}^i = \max_{x_{0,1,\tau}^i} \int_t^\infty x_{0,1,\tau}^i (\lambda V_{1,1,\tau} - c) e^{-\int_t^\tau \lambda(x_{0,1,s}^i + x_{1,0,s}^{-i})ds} d\tau, \tag{4}$$

where the laggard anticipates to receive a continuation payoff $V_{1,1,\tau}$ if she succeeds and zero if her opponent succeeds. Obviously, the laggard is willing to exert any effort only if her continuation

payoff upon success compensates her cost of effort. It follows immediately that the principal needs to specify a fixed reward weakly greater than $3c/\lambda$ (which makes $V_{1,1,t} \geq c/\lambda$) to sufficiently incentivize the laggard. The following proposition establishes that this minimum fixed reward is enough to encourage full effort by both agents at all times.

PROPOSITION 2. *Under full information disclosure, the minimum fixed reward needed to induce both agents to exert full effort at all times is $3c/\lambda$.*

Proposition 2 highlights that when the agents' partial progress is observable, the laggard is at risk of becoming discouraged if the reward is not large enough. This is because an agent i's chance of winning the reward declines once her opponent proceeds to the second stage (and hence agent i becomes a laggard). Therefore, the principal has to offer the minimum fixed reward of $3c/\lambda$ to keep the laggard's continuation payoff upon success equal to c/λ. However, with this large fixed reward, the principal overpays the leader and she receives a continuation payoff equal to $V_{1,0,t}^i = 3c/(2\lambda)$ as derived in the Appendix. This large reward also delivers an ex-ante expected utility equal to $V_{0,0,t}^i = c/(4\lambda)$ to each agent, which shows that the principal leaves money on the table as compared to the first-best contract. In the next section, we investigate whether the principal can do better than the fixed-reward contest by designing a flexible-reward contest that achieves the absolute minimum expected lead time $\underline{T}$ under full information disclosure.

4.3. Full Information Disclosure with Flexible Reward

Notice that under full information disclosure, each agent's effort provision decision and continuation payoff depend solely on the state of the game, rather than on time itself since in the absence of a deadline, the agent's optimization problem is stationary. Thus, as we explain in the proof of Theorem 1, it is without loss of optimality to let the principal choose a reward schedule that depends only on the state of the contest (i.e., whether each agent achieved a success or not). Observe that a contest may end under two states of the world: (i) a case where the leader obtains the second success before the laggard obtains any success and (ii) a case where both agents have already obtained one success and one of them achieves the second success. Denote by $R_{2,0}$ the contest reward in the former case and $R_{2,1}$ the adjusted contest reward in the latter case. This flexible-reward contest enables the principal to motivate full effort at all times by giving sufficient incentives to the laggard without overpaying the leader. The following theorem formalizes this argument and characterizes the optimal level of rewards.

THEOREM 1. *Under full information disclosure, a flexible-reward contest with $R_{2,0} = 2c/\lambda$ and $R_{2,1} = 3c/\lambda$ induces both agents to exert full effort at all times by paying the first-best expected reward of $5c/(2\lambda)$.*

As discussed in §4.2, under full information disclosure, the principal has to offer a minimum continuation payoff of c/λ upon the arrival of a success to motivate the laggard whose opponent advanced to the second stage. Thus, as in the case of fixed reward, this motivation can be achieved by offering $3c/\lambda$ to the winner if the laggard obtains her first success. However, unlike a fixed-reward schedule, the proposed flexible-reward design does not have to overpay the leader, who has larger incentives to exert effort knowing that her opponent is in the first stage. Specifically, the reward a leader will receive before the laggard achieves the first success (i.e., before the leader loses her lead, see case (i) above) is $R_{2,0} = 2c/\lambda$. On the other hand, the reward a winner will receive after both agents achieve the first success is $R_{2,1} = 3c/\lambda$ (i.e., after the leader loses her lead, see case (ii) above). The proposed flexible-reward schedule in Theorem 1 achieves the absolute minimum expected lead time $\underline{T}$ by eliciting full effort at all times from both agents and gives the minimum necessary reward. It also delivers the minimum surplus to the agents while respecting their incentive compatibility constraints (as derived in the Appendix) at all times by offering a continuation payoff equal to zero (recall that agents' outside option is normalized to zero) to any agent with no success (i.e., $V_{0,0,t}^i = V_{0,1,t}^i = 0$) and a continuation payoff equal to c/λ to any agent with one success (i.e., $V_{1,1,t}^i = V_{1,0,t}^i = c/\lambda$).

The optimality of the proposed flexible-reward schedule can be seen from the fact that it incentivizes full effort at all times by just offering the first-best expected reward. This is because with probability $1/2$, the contest ends before the arrival of any success for the laggard and the principal spends $R_{2,0} = 2c/\lambda$; and with probability $1/2$, the contest ends after the arrival of the first success for the laggard (the state when both agents have obtained one success) and the principal pays $R_{2,1} = 3c/\lambda$. Thus, the expected reward of the contest equals $5c/(2\lambda)$, which is the first-best expected reward in Proposition 1. Another key feature of this design which makes it practically appealing is its simplicity. The policy can easily be implemented by offering a guaranteed reward of $2c/\lambda$ with the option to increase the reward if multiple agents progress to the second stage. The flexible-reward schedule is quite impactful as well because it spends 20% less money on average relative to the fixed-reward schedule to achieve the same contest outcome.

In practice, full information disclosure may not always be feasible because it may be difficult for the principal to continuously monitor agents' progress. To capture this setting, in the next section, we investigate the effectiveness of a flexible-reward design in the case of no information disclosure.

4.4. No Information Disclosure with Flexible Reward

In this section, we characterize the optimal flexible-reward schedule in a contest where the principal does not use any information sharing. Because each agent does not receive any information from the principal, she forms a belief about her opponent's progress and updates this belief over time.

Specifically, let p_t^i be the probability that agent i assigns at time t to the event that her opponent has already progressed to the second stage (i.e., achieved a success in the first stage). As time passes, the only information that agent i receives is whether the contest is still ongoing which signals to agent i that her opponent has not yet completed the second stage. Given this information update, and by Bayes' rule, p_t^i evolves according to:

$$dp_t^i = \lambda(1 - p_t^i)(x_{0,t}^{-i} - p_t^i x_{1,t}^{-i})dt, \tag{5}$$

with the boundary condition $p_0^i = 0$ where $x_{0,t}^{-i}$ denotes the opponent's effort at time t conditional on not having achieved a success yet, and $x_{1,t}^{-i}$ denotes her effort at time t conditional on having achieved the first success.[6] Intuitively, and as shown in (5), the probability that each agent's opponent already advanced to the second stage, p_t^i, is increasing in the opponent's first-stage effort $x_{0,t}^{-i}$ and decreasing in the opponent's second-stage effort $x_{1,t}^{-i}$.

Recall that $V_{k,t}^i$ is the continuation payoff of each agent i who has achieved $k \in \{0, 1\}$ successes by time t. Then, the maximization problem for agent i after obtaining the first success is given by:

$$V_{1,t}^i = \max_{x_{1,\tau}^i} \int_t^\infty x_{1,\tau}^i (\lambda R_\tau - c) e^{-\int_t^\tau \lambda(x_{1,s}^i + p_s^i x_{1,s}^{-i})ds} d\tau, \tag{6}$$

where R_τ is the specified reward if the contest ends at time τ. The above expression can be interpreted as follows: if agent i chooses effort $x_{1,\tau}^i$ during interval $(\tau, \tau + dt)$, she incurs a cost $cx_{1,\tau}^i dt$ and if a success arrives with probability $\lambda x_{1,\tau}^i dt$, she enjoys a reward R_τ. On the other hand, the probability that her opponent is in the second stage and achieves a success in this time interval is $p_\tau^i \lambda x_{1,\tau}^{-i} dt$, and in that case, the agent receives zero reward. Finally, $e^{-\int_t^\tau \lambda(x_{1,s}^i + p_s^i x_{1,s}^{-i})ds}$ captures the probability that neither the agent achieves a success, nor does her opponent achieve the second success by time τ. Anticipating a continuation payoff of $V_{1,\tau}^i$ upon the arrival of the first success at time τ, agent i's continuation payoff from time t onward before obtaining any success can be expressed as follows:

$$V_{0,t}^i = \max_{x_{0,\tau}^i} \int_t^\infty x_{0,\tau}^i (\lambda V_{1,\tau}^i - c) e^{-\int_t^\tau \lambda(x_{0,s}^i + p_s^i x_{1,s}^{-i})ds} d\tau. \tag{7}$$

The above expression can be interpreted similar to (6). As derived in condition (31) in the Appendix, an agent with no success finds it optimal to exert full effort if and only if her additional utility upon the arrival of her first success is sufficiently large. Proposition 3 characterizes the optimal flexible-reward schedule that induces both agents to exert full effort at all times.

---

[6] Note that, by Bayes' rule, the probability that agent i assigns at time $t + dt$ to the event that her opponent has succeeded once, given p_t^i, can be expressed as follows:

$$p_{t+dt}^i = \frac{p_t^i(1 - x_{1,t}^{-i}\lambda dt) + (1 - p_t^i)x_{0,t}^{-i}\lambda dt}{p_t^i(1 - x_{1,t}^{-i}\lambda dt) + 1 - p_t^i},$$

where the numerator is the probability that the opponent is in the second stage given that the game has not ended yet, and the denominator is the total probability that the contest has not finished yet. The law of motion can be obtained by subtracting p_t^i from both sides, dividing by dt, and taking the limit as $dt \to 0$.

PROPOSITION 3. *Under no information disclosure, a flexible-reward contest with $R_t = (2 + p_t)c/\lambda$ where $p_t = \lambda t/(1 + \lambda t)$ induces both agents to exert full effort at all times by paying the first-best expected reward of $5c/(2\lambda)$.*

Under no information disclosure, an agent gradually increases the probability she assigns to her opponent's partial progress. This update in the agent's beliefs reduces her incentives to exert full effort. Thus, to mitigate reduced incentives, the optimal design with no information disclosure offers a gradually increasing reward schedule over time which achieves the first-best expected reward despite unobservable effort. To understand the intuition for this reward structure better, recall that under full information disclosure, the principal has to offer a minimum continuation payoff of c/λ to encourage an agent with no success to keep working when her opponent already progressed to the second stage. In the absence of any information about the opponent's progress, an agent i's belief about her opponent agent $-i$'s progress gradually increases over time, according to (5), as she anticipates full effort provision in equilibrium. Thus, as agent i fails to achieve the first success over time, she will anticipate that agent $-i$ is more and more likely to be in the second stage. As a result, agent i will think that even if she advances to the second stage, she is likely to compete with agent $-i$ in the second stage for the reward. This reduces the expected utility of progressing to the second stage over time, and hence reduces the incentives for an agent with no success to exert effort in the first stage. In order to give sufficient incentives to agents, the principal offers a flexible-reward schedule according to $R_t = (2 + p_t)c/\lambda$, where p_t is the equilibrium belief of each agent about the partial progress of her opponent.

Notice that agents' and principal's combined surplus is the same across any design that achieves the absolute minimum expected lead time $\underline{T}$. Therefore, the principal's surplus is maximized when the agents' surplus is minimized. The proposed flexible-reward schedule in the above proposition with no information disclosure minimizes the agents' surplus by keeping the continuation payoff of an agent with no success equal to zero (her outside option) and the continuation payoff of an agent with one success equal to c/λ which is the bare minimum utility to incentivize first-stage effort. Therefore, this design maximizes the principal's surplus and hence is optimal.

Our results suggest that under both full information and no-information disclosure policies, the principal should pay attention to two things. First, the principal should give sufficient reward to keep the agent with no success incentivized to work when her opponent already progressed (or is likely to have progressed) to the second stage. Second, the principal should avoid overpaying agents when they have sufficient incentives to work. Under full information disclosure, this can be achieved by increasing the reward one time from $2c/\lambda$ to $3c/\lambda$ when both agents complete the first stage. Under no information disclosure, a gradually increasing reward in the form of $(2 + p_t)c/\lambda$

sustains full effort with the lowest expected reward. Therefore, for any information disclosure policy, the principal can find an appropriate flexible-reward schedule that attains the absolute minimum expected lead time at the minimal cost of incentives. This indicates that although the flexible-reward schedule is easier to implement under a full information disclosure policy, the specific information disclosure policy does not make a material difference in the principal's objective. However, this result assumes that the principal has a minimum budget of $3c/\lambda$ to encourage the agent with no success to keep working when her opponent is in the second stage. (Note that the reward under no information disclosure increases to $3c/\lambda$ as the opponent's probability of advancing to the second stage approaches 1). However, when the principal is budget-constrained (i.e., her maximum budget $\overline{R}$ is strictly less than $3c/\lambda$), these reward schedules will not be directly implementable and information disclosure may play a more central role as a non-monetary incentive mechanism in designing dynamic contests. Therefore, in the following section, we focus on the information disclosure policy for a budget-constrained principal with a budget less than $3c/\lambda$.

5. Using Information Design for Development Contests

In this section, we study another potential degree of freedom for principals engaging in development contests, namely the use of information design. Such an approach can come in handy for a budget-constrained principal who may fall short of the optimal reward amount derived in the previous section. As we discussed in §4, the principal needs a minimum budget of $2c/\lambda$ for the second success to be attainable (c/λ for each success) with no competition, and $3c/\lambda$ is a sufficient budget to implement full effort at all times with two competing agents. If the principal is resource constrained and falls short of this sufficient budget amount, we study how they can use information disclosure to incentivize suppliers. We begin by considering a fixed reward and then consider information disclosure jointly with a flexible reward in §5.3. For the rest of the analysis, we assume $2c/\lambda < \overline{R} < 3c/\lambda$, and we name a principal facing this limitation a *budget-constrained* principal. In §5.1, we discuss the role of strategic information disclosure and propose a novel probabilistic design which improves upon extreme cases of disclosure policies. Then in §5.2, we study the comparative benefit of probabilistic disclosure over other canonical disclosure policies.

5.1. Information as a Strategic Commodity

An agent who has already achieved one success is easy to incentivize because she is already encouraged by the fact that she needs only one more success to obtain the reward. For such an agent, as we will show throughout this section, the information disclosure policy has no impact on incentives. In contrast, an agent with no success is at risk of becoming discouraged over time once she realizes (or believes) that her opponent has already progressed to the second stage. To analyze the impact of

partial progress on the incentives of agents and show the strategic value of information disclosure, we first analyze the problem faced by a budget-constrained principal that offers no information disclosure throughout the contest. Notice that unlike §4.4, the principal does not have sufficient funds to gradually increase the reward up to $3c/\lambda$ to mitigate an agent's reduced incentives caused by the threat that her opponent has already progressed to the second stage. Therefore, in the unique symmetric equilibrium of the contest, an agent with no success only exerts full effort as long as monetary incentives are sufficient, and lowers her effort level after a certain threshold when incentives are missing. The equilibrium is formally characterized below.

PROPOSITION 4. *When the principal is budget-constrained, and under no information disclosure, there exists a unique symmetric equilibrium in which an agent with no success exerts full effort until some time $t_r = \frac{p_r}{\lambda(1-p_r)}$ where $p_r = \frac{\lambda\overline{R}}{c} - 2$. After t_r, she reduces her effort level to $p_r(<1)$. An agent who has achieved one success exerts full effort until the end.*

To understand Proposition 4, we need to first understand that exerting full effort is incentive compatible for an agent with no success if and only if she earns an additional utility of at least c/λ upon the arrival of her first success (see condition (31) in the Appendix). While the agent is working, the expected utility of progressing to the second stage diminishes over time as she puts more weight to the event that her opponent already achieved the first success. Therefore, a budget-constrained principal can only induce an agent with no success to exert full effort until time t_r where her belief reaches $p_r = \frac{\lambda\overline{R}}{c} - 2$. After that, monetary incentives are not sufficient to justify full effort, so each agent with no success reduces her effort to p_r and keeps this effort level. To understand why the equilibrium effort becomes p_r, we need to understand that p_r is the belief level that keeps an agent indifferent between exerting any effort level. When an agent i's belief is below p_r (which happens if the opponent exerts effort smaller than p_r after t_r), she exerts full effort but when it is above p_r (which happens if the opponent exerts effort larger than p_r after t_r), she exerts zero effort. Thus, after t_r, the unique symmetric equilibrium effort is p_r, which uniquely keeps the agent's belief at the threshold and holds her expected continuation payoff after obtaining the first success at c/λ such that the agent remains indifferent between exerting any effort at each instant.

It is worth noting that both the time threshold t_r and the belief threshold and reduced effort level p_r are increasing in the size of the budget $\overline{R}$ and the achievability parameter λ, and decreasing in the marginal cost of effort c. Also, as we discuss above, each agent has sufficient incentives to exert full effort until the end after obtaining the first success (see condition (33) in the Appendix).

While no-information design elicits full effort before time t_r, the reduced effort provision by an agent with no success after time t_r hints the benefit of using information strategically to improve contest outcome. Let us first consider the full information disclosure policy. As discussed in §4.3,

the principal should offer a reward of $3c/\lambda$ to the laggard once she realizes that her opponent is in the second stage. Obviously, a budget-constrained principal does not have access to such a high reward. Therefore, upon the arrival of the first success, the laggard stops exerting effort immediately (see Proposition B1 in Appendix B). This is in contrast to no-information disclosure policy where the principal never loses the laggard, but an agent with no success reduces her effort level to $p_r < 1$ after time t_r even if her opponent does not achieve any partial progress. Our result seems consistent with the empirical findings of Lemus and Marshall (2021) who study information disclosure (no versus full), and show that disclosing information discourages a laggard, but the lack of information creates uncertainty regarding how much effort is needed to remain competitive.

In a recent study, Bimpikis et al. (2019) highlight this observation and suggest a cyclic information disclosure design with *silent periods* in between and full disclosure at the end of each cycle. Specifically, in their mechanism, the principal stays silent during cycles of t_r periods of time and discloses full information about the status of the contest at the end of each cycle. This design somehow blends the key benefits of the two extreme cases by: i) hiding any partial progress during each silent period and inducing the agents to keep working, and ii) disclosing full information at the end of each silent period to replenish incentives of the agents and to avoid reduced effort levels. We next analyze the equilibrium under cyclic information disclosure policy in our setting.

PROPOSITION 5. *When the principal is budget-constrained, and commits to cyclic information disclosure with silent periods of length $t_r = \frac{p_r}{\lambda(1-p_r)}$ where $p_r = \frac{\lambda \overline{R}}{c} - 2$, there exists a unique symmetric equilibrium in which both agents exert full effort until time t_r. At the end of the first cycle:*
(i) If no agent has made any partial progress by time t_r, the contest resets and the next silent period with length t_r begins in which both agents exert full effort.
(ii) If only one agent has made partial progress at the end of one cycle, the agent with no success quits and the agent with one success exerts full effort until the end.
(iii) If both agents have made partial progress at the end of one cycle, both agents exert full effort until the end.

The intuition of Proposition 5 is as follows. A cyclic design with silent periods of length t_r ensures that no information is provided to the agents during each cycle where the incentives are sufficient to elicit full effort. At the end of each cycle, the belief of any agent with no success reaches the threshold p_r. In order to mitigate the reduced effort by the agents with no progress, the principal provides full information disclosure so that if both agents have zero success, they reset their beliefs and keep exerting full effort for another cycle. However, the major caveat of this information disclosure is that if one of the agents has progressed to the second stage by the end of one cycle, an agent with no success immediately quits once this information is disclosed.

While the cyclic information policy improves upon the full information policy by delaying the laggard's stopping time with the help of the silent periods and hence reducing the expected lead time of the contest, it does not necessarily improve upon the no-information policy (see Figure 2). Indeed, no-information policy provides no information after t_r when incentives are missing, and cyclic information policy may be providing too much information at the end of each cycle which may hurt the principal by losing one of the agents. Thus, information is a strategic commodity and the principal must ensure not to provide too much or too little information.

As the main result of this section, we will provide a better information disclosure policy that improves upon all of the above policies. To address the issues of too little information under no information disclosure and too much information under full (or cyclic) information disclosure, we propose an information disclosure policy in which the principal commits to disclose any partial progress (arrival of the first success) with time-dependent instantaneous rate $\gamma_t \geq 0$. If we study the problem faced by agent i who has not achieved any success under this information policy, we observe that the belief of agent i about the partial progress of her opponent evolves according to:

$$dp_t^i = (1 - p_t^i)(x_{0,t}^{-i}\lambda - p_t^i x_{1,t}^{-i}\lambda - p_t^i \gamma_t)dt. \tag{8}$$

This equality takes into account that the principal sharing no news with agent i has a lowering effect on the probability that her opponent is in the second stage. Under this design, the continuation payoff of agent i with no success from any time t onward is the following:

$$V_{0,t}^i = \max_{x_{0,\tau}^i} \int_t^\infty x_{0,\tau}^i (\lambda V_{1,\tau}^i - c)e^{-\int_t^\tau (x_{0,s}^i\lambda + p_s^i x_{1,s}^{-i}\lambda + p_s^i \gamma_s)ds}d\tau, \tag{9}$$

where by choosing effort $x_{0,\tau}^i$ during interval $(\tau, \tau + dt)$, the agent incurs a cost $cx_{0,\tau}^i dt$ and if a success arrives, she enters the second stage and enjoys a continuation payoff of $V_{1,\tau}^i$. If her opponent completes both stages during interval $(\tau, \tau + dt)$, agent i receives zero reward. Moreover, if the principal discloses partial progress of agent i's opponent, agent i quits and receives zero utility because her continuation payoff upon the arrival of her first success falls below c/λ. Also, recall that an agent with one success always exerts full effort until the end.

Previous analysis indicates that the principal can induce agent i to exert full effort with no information sharing until her belief reaches p_r. After that, the principal's budget is not sufficient to incentivize full effort. When no information is provided (i.e., $\gamma_t = 0$), the agent reduces her effort level to p_r which holds the agent's equilibrium belief constant. If full information is provided ($\gamma_t = \infty$), the agent keeps working if her belief jumps to zero and quits if her belief jumps to one. Using the idea of dynamic Bayesian Persuasion (e.g., Ely 2017), we introduce the following design.

DEFINITION 1. The "***probabilistic encouragement design***" prescribes no information to the agents (i.e., $\gamma_t = 0$) up to time $t_r = \frac{p_r}{\lambda(1-p_r)}$ where $p_r = \frac{\lambda \overline{R}}{c} - 2$. After that it discloses any partial progress with rate $\gamma_r = \frac{\lambda(1-p_r)}{p_r} = \frac{1}{t_r}$.

The probabilistic encouragement design features a silent period of length t_r in the beginning where the principal discloses no information and the agents' beliefs drift upward to p_r according to (5). After that, the principal commits to inform each agent of her opponent's partial progress at constant rate γ_r. Proposition 6 describes the equilibrium under this design.

PROPOSITION 6. *When the principal is budget-constrained, and commits to probabilistic encouragement design, an agent who has not achieved a success exerts full effort until she obtains her first success, or her opponent obtains her second success, or the principal discloses the opponent's partial progress. An agent who has achieved one success exerts full effort until the end.*

The key benefit of the probabilistic encouragement design is to fine-tune the rate of information disclosure after the silent period t_r. In this case, (8) implies that the principal can persuade an agent with no success to exert higher effort than the case of no information disclosure, $x_{0,t} > p_r$, by committing to disclose partial progress with rate $\lambda(x_{0,t} - p_t)/p_t$. This rate of information disclosure holds the agent's belief at the threshold p_r when her opponent exerts effort $x_{0,t} > p_r$ as long as no partial progress is disclosed. This belief keeps the agent's continuation payoff upon success at c/λ for all $t \geq t_r$, and hence the unique symmetric equilibrium is where an agent with no success continues exerting effort $x_{0,t}$ until she succeeds, the contest ends, or partial progress is disclosed.

Notice that the principal is facing a trade-off by sharing information regarding partial progress. A higher rate of disclosure stimulates greater effort from an agent with no success after t_r which reduces the expected lead time of the contest, but it also increases the probability of losing the laggard and prolonging the contest's expected lead time if partial progress is disclosed. The probabilistic encouragement design works such that the bare minimum of information is disclosed by the principal to sustain full effort (i.e., $x_{0,t} = 1$) after t_r as illustrated in Figure 1(c). This policy enables the principal to not disclose too much information like the case of full information disclosure and also significantly improves the reduced effort level of the agents under no information disclosure. Figure 1 depicts the evolution of beliefs and the continuation payoffs $V_{0,t}$ and $V_{1,t}$ over time for an agent with no success under probabilistic encouragement design. As illustrated in Figures 1(a) and 1(b), the agent's belief remains at p_r and her incentive compatibility condition (36) is binding after t_r. We are now ready to state the main result of this section.

THEOREM 2. *Probabilistic encouragement design yields an expected lead time of $\left(5 + e^{-2\lambda t_r}\right)/(4\lambda)$, which is strictly lower than the expected lead times under no, full or cyclic information disclosure policies.*

As we discussed in §4.1, the absolute minimum expected lead time of the contest for an unconstrained principal is $\underline{T} = 5/(4\lambda)$, and this minimum is attainable by inducing both agents to exert
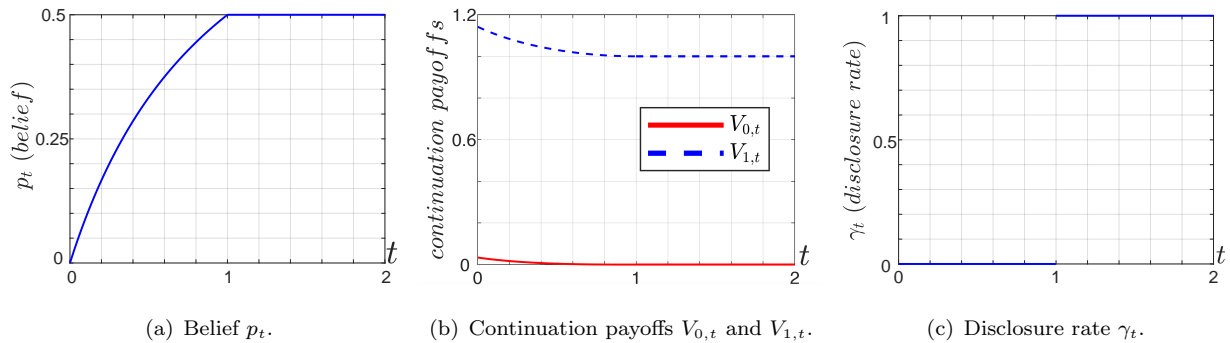
(a) Belief p_t.   (b) Continuation payoffs $V_{0,t}$ and $V_{1,t}$.   (c) Disclosure rate γ_t.

Figure 1   Probabilistic encouragement design. Setting: $c = \lambda = 1$ and $\overline{R} = 2.5$. Here, $t_r = 1$ and $p_r = 0.5$.

full effort at all times. When the principal is budget-constrained, inducing agents to exert full effort is no longer possible. Yet, the probabilistic encouragement design performs relatively well and prolongs the expected lead time by only $e^{-2\lambda t_r}/(4\lambda)$. This additional lead time gets smaller if the reward size $\overline{R}$ or the achievability parameter λ increases, or the cost of effort c decreases.

As cyclic information disclosure clearly dominates full information disclosure, we will explain the intuition for why probabilistic encouragement design dominates cyclic information disclosure policy. To understand the intuition, we need to compare agents' effort provision under different states in the contest. Both policies are able to elicit full effort in states with no partial progress or after both agents achieve partial progress. Thus, the subtlety arises in states where only one success is obtained. There are two different cases depending on whether the first success is achieved before or after time t_r. If there is one partial progress before t_r, say at time t, cyclic information disclosure immediately shares this information at time t_r and loses the laggard. However, probabilistic encouragement design delays the relaying of this information and keeps the laggard incentivized for t_r more periods of time on average. This is because $\gamma_r = 1/t_r$ and disclosure starts at t_r. If the first success arrives at time $t > t_r$, then cyclic information disclosure policy releases this information at the end of the current cycle with a delay of at most t_r periods of time. Probabilistic encouragement design on the other hand allows a delay of t_r periods on average. Thus, in both cases, probabilistic encouragement design delays the disclosure of partial progress more (on average) than the cyclic information disclosure, and keeps the laggard incentivized longer.

Next, to see how probabilistic encouragement design dominates no information disclosure, we show in the Appendix that the expected lead time of the contest is the same across both designs once the first success is obtained. However, probabilistic encouragement design significantly outperforms no information disclosure by sustaining full effort when both agents have zero success, whereas agents reduce their effort level after t_r under no information disclosure.

In the probabilistic encouragement design, the principal commits to disclose any partial progress probabilistically. For cases where it is difficult to justify such a probabilistic disclosure policy, we

propose an equivalent and easy-to-implement "random-update" policy. In this policy, the principal commits to not monitoring the status of agents up to time t_r and then periodically monitoring the status of agents at random times (with arrival rate $\gamma_r = \frac{1}{t_r}$) and immediately disclosing any partial progress. This way the principal does not need to constantly observe the status of the game and selectively obfuscate information. Instead, in an economically equivalent approach, the principal can deliberately choose to only monitor agents' progress at random times after time t_r and inform agents immediately if any progress is made.

In the following section, we further highlight the value of probabilistic disclosure design by showing how it improves upon a deterministic information disclosure with delay.

5.2. The Value of Probabilistic Information Disclosure

In the previous section, we proposed a probabilistic information disclosure policy and showed that it dominates other mainstream information disclosure policies. The key idea of probabilistic encouragement design is to probabilistically delay the disclosure of partial progress and hence keep an agent with no success incentivized for a longer period of time. Thus, one may wonder whether the same benefit can be achieved by deterministically delaying the disclosure of partial progress. In this section, we will show that this is not the case.

Consider a case where the principal commits to disclose any partial progress with t_d periods of delay. This mechanism leads to an initial silent period of length t_d during which agents' beliefs drift upward to p_d according to (5). During interval $(t_d, t_d + d\tau)$, if the principal announces partial progress, an agent with no success (i.e., the laggard) quits. Otherwise, each agent's belief remains constant at p_d as she realizes that no progress has been made during the initial $d\tau$ period of the contest, akin to a contest that starts at $d\tau$ instead of time 0. Here, to compare the contest outcome with our proposed probabilistic encouragement design (and equivalently our random-update policy), we restrict our attention to mechanisms in which an agent with no success puts full effort until information is disclosed by the principal. Intuitively, the principal wishes to extend this delay as long as possible to keep the laggard in the contest. The following theorem derives the unique symmetric equilibrium under this deterministic delay policy and compares the expected lead time of the contest under this policy with that under our probabilistic encouragement design.

THEOREM 3. *Suppose that the principal is budget-constrained. In the deterministic delay design, the principal commits to disclose partial progress after a delay of length t_d given by*

$$\left(1 + e^{-2\lambda t_d}\right)\left(\overline{R} - \frac{c}{\lambda}\right) = 2\left(1 + \lambda t_d\right)\left(\frac{3c}{\lambda} - \overline{R}\right). \tag{10}$$

Furthermore, under the deterministic delay design:

(i) An agent who has not achieved a success exerts full effort until she obtains her first success, or

her opponent obtains her second success, or the principal discloses the opponent's partial progress.
An agent who has achieved one success exerts full effort until the end.

(ii) Delay length $t_d < t_r = \frac{p_r}{\lambda(1-p_r)}$ where $p_r = \frac{\lambda \overline{R}}{c} - 2$.

(iii) The expected lead time of the contest is given by $(5 + e^{-2\lambda t_d})/(4\lambda)$, which is strictly larger than
the expected lead time under the probabilistic encouragement design.

To elicit full effort from agents, the principal should commit to disclose any progress at most
t_d periods of time after its arrival. Theorem 3 shows that when information is disclosed with a
deterministic delay, the principal cannot keep the initial silent period as long as the one in prob-
abilistic encouragement design (i.e., $t_d < t_r$). The intuition is as follows. When deciding whether
to exert effort at any instant, an agent i with no success trades off the benefit of achieving partial
progress with the cost of additional effort. The more likely it is for agent i's opponent to give up
after agent i's partial progress, the higher agent i's benefit from partial success. Under probabilistic
encouragement design, there is a chance that the principal discloses agent i's partial progress at
any instant, whereas under deterministic delay, agent i knows that there is no such chance because
the principal never discloses success until t_d periods after its arrival. Thus, under probabilistic
encouragement design, an agent i with no success expects to benefit more from partial progress
and hence has larger incentives to exert full effort. Somewhat counter-intuitively, although prob-
abilistic encouragement design on average delays information disclosure more than deterministic
delay design (i.e., $t_r > t_d$), the agent is still willing to work longer under this design.

This observation highlights the value of probabilistic information disclosure as it smoothens
incentives over time, and helps the principal extend the period over which agents are willing to
work. Consequently, as depicted in Figure 2, our probabilistic encouragement design achieves a
shorted average lead time than its deterministic counterpart. Figure 2 also illustrates that the
probabilistic design (and equivalently our random-update policy discussed in §5.1) outperforms
other canonical information disclosure policies unconditionally.

5.3. Probabilistic Information Disclosure with Flexible Reward

So far in §5, we focus on the role of information disclosure in the contest design by holding the
reward fixed. We can further improve our proposed probabilistic encouragement design without
changing agents' equilibrium effort levels by incorporating flexible reward to reduce the expected
reward of the contest. Observe that in the probabilistic encouragement design, the principal dis-
closes no information up to time t_r where monetary incentives are sufficient to motivate full effort.
Only after t_r, when the belief of an agent with no success reaches the threshold p_r, the principal
starts disclosing information. Proposition B2 in Appendix B states that the principal can offer a
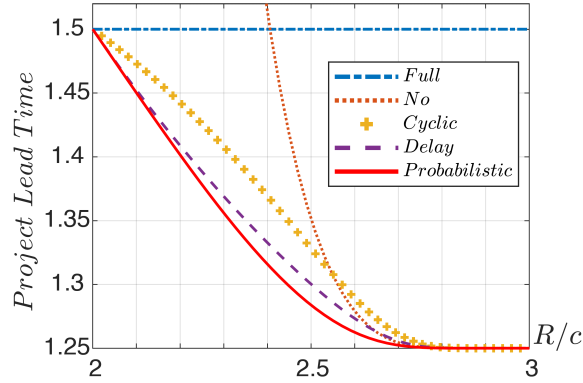
Figure 2   Probabilistic encouragement design reduces expected project lead time. Setting: $\lambda = 1$.

flexible reward according to $R_t = (2 + p_t)c/\lambda$ where $p_t = \lambda t/(1 + \lambda t)$ if $t < t_r$, and $R_t = \overline{R}$ if $t \geq t_r$ and achieve the same equilibrium outcome as in the original probabilistic encouragement design.

The intuition for this result is that the principal does not need to overpay agents during early stages of the contest where incentives are sufficiently high. As discussed in §4.4, the principal can guarantee sufficient incentives to agents by offering a gradually increasing reward schedule over time as each agent finds it more likely for her opponent to achieve partial progress. This increasing flexible-reward structure up to time t_r helps the principal to set the reward at the right amount before time t_r at which the reward schedule hits the budget constraint $\overline{R}$. After that, the principal spends the whole reward budget. Notice that the expected lead time is strictly decreasing in $\overline{R}$ as a higher reward budget helps the principal to start the probabilistic information disclosure phase later. For the same reason, the principal benefits from allocating her entire budget after time t_r. Combining probabilistic disclosure with flexible reward can help the principal reduce both the expected lead time and the cost of reward in the contest.

6. Conclusion

Organizations worldwide face the challenges of developing advanced technologies and products under time and budget pressures, and turn to their expert suppliers to tackle complex problems. In this paper, we study how such organizations can effectively organize a dynamic development contest to stimulate development effort from a small set of competing suppliers to minimize the lead time of a multi-stage project while keeping the incentive budget in check.

We sought to build a parsimonious model of a dynamic development contest where two agents compete to complete a two-stage development project by exerting costly effort over a continuous time frame. Successful completion of a stage for an agent arrives at a random point in time where the rate of arrival for each success increases with the agent's effort. An agent's success is not observable by another agent so it is up to the principal whether and when to share this information. The contest ends when one of the agents successfully completes both stages and wins a pre-determined

reward. At the outset of the contest, the principal commits to a reward schedule that determines how this winner reward will change over time and an information disclosure policy which specifies how the principal will disclose information throughout the contest.

We establish that a principal with no budget constraint can utilize a flexible-reward schedule to achieve the absolute minimum expected lead time by giving 20% less reward on average than a fixed-reward schedule. Under full-information disclosure, the principal can easily implement the optimal flexible-reward schedule by setting a guaranteed reward amount upfront with the promise of increasing the reward amount if multiple agents successfully complete the first stage. Importantly, the flexible-reward schedule pays the absolute minimum expected reward (i.e., the first-best). We next analyze how a budget-constrained principal can utilize information as an incentive tool. By harnessing a dynamic Bayesian persuasion approach, we characterize an easy-to-implement "random-update" policy in which the principal does not monitor the status of agents for a pre-determined initial time window; and then periodically monitors the status of agents at random times and discloses any partial progress immediately. We show that this random-update policy achieves a shorter expected lead time than any canonical information disclosure policy. We then find that by utilizing a flexible-reward policy during the initial no-monitoring period, the principal can reduce the cost of incentives without hindering project lead time. Our results indicate that enriched rank-based incentives and carefully-tailored information design can be powerful tools to incentivize development efforts of suppliers without overpaying them. With development contests such as vaccines and hypersonic missiles running into billions of dollars and racing against time, the flexible reward coupled with information disclosure can potentially achieve substantial savings in project budget (to the tune of hundreds of millions of dollars) and lead time.

Our analysis opens up several interesting future research directions. First, as a first step to understanding dynamic development contests, we have abstracted away from features such as skills heterogeneity, learning by doing, or the uncertainty regarding the feasibility of the first or second stage, but extending our work by containing such features can be interesting research avenues. Adding these features would make the analysis more involved but can provide useful further insights on designing development contests. Second, an interesting research to pursue would be to see how our proposed methods work when success has a quality measure that can be improved over time rather than taking the form of a breakthrough. Finally, our random-update policy also provides new opportunities for empirical and experimental research. Recently Mostagir et al. (2021) run a laboratory experiment to investigate the impact of full and no-information disclosure policies on the agents' behavior and the principal's outcome in a framework similar to Bimpikis et al. (2019). In our model, besides considering full and no information, we analyze and compare more strategic disclosure policies including cyclic, deterministic delay and probabilistic policies to show the value

of information design and highlight its impact on incentives. Studying these policies empirically or experimentally could be an interesting research direction.

References

Ales L, Cho SH, Körpeoğlu E (2017) Optimal award scheme in innovation tournaments. *Operations Research* 65(3):693–702.

Ales L, Cho SH, Körpeoğlu E (2020) Innovation tournaments with multiple contributors. *Production and Operations Management* Forthcoming.

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

Bergemann D, Morris S (2019) Information design: A unified perspective. *Journal of Economic Literature* 57(1):44–95.

Bimpikis K, Ehsani S, Mostagir M (2019) Designing dynamic contests. *Operations Research* 67(2):339–356.

Candoğan ST, Korpeoglu CG, Tang CS (2020) Team collaboration in innovation contests. *Working Paper, University College London* .

Chen Z, Mihm J, Schlapp J (2018) Sourcing innovation: Integrated system or individual components? *INSEAD Working Paper* .

deVericourt F, Gurkan H, Wang S (2020) Informing the public about a pandemic. *Management Science* Forthcoming.

Ederer F (2010) Feedback and motivation in dynamic tournaments. *Journal of Economics & Management Strategy* 19(3):733–769.

Ely J, Georgiadis G, Khorasani S, Rayo L (2020) Optimal feedback in contests. *Working paper, Northwestern University* .

Ely JC (2017) Beeps. *American Economic Review* 107(1):31–53.

GAVI (2020) What is an advance market commitment and how could it help beat covid-19? https://www.gavi.org/vaccineswork/what-advance-market-commitment-and-how-could-it-help-beat-covid-19, accessed on March 07, 2021.

Goltsman M, Mukherjee A (2011) Interim performance feedback in multistage tournaments: The optimality of partial disclosure. *Journal of Labor Economics* 29(2):229–265.

Halac M, Kartik N, Liu Q (2017) Contests for experimentation. *Journal of Political Economy* 125(5):1523–1569.

Horner J, Skrzypacz A (2017) Learning, experimentation, and information design, volume 1 of econometric society monographs.

Hu M, Wang L (2020) Joint vs. separate crowdsourcing contests. *Management Science* Forthcoming.

Kamenica E (2019) Bayesian persuasion and information design. *Annual Review of Economics* 11:249–272.

Kamenica E, Gentzkow M (2011) Bayesian persuasion. *American Economic Review* 101(6):2590–2615.

Khorasani S, Nittala L, Krishnan V (2020) Screening in complex contests. *Working Paper, University of California San Diego* .

Korpeoglu CG, Körpeoğlu E, Tunç S (2020) Optimal duration of innovation contests. *Manufacturing & Service Operations Management* Forthcoming.

Körpeoğlu E, Cho SH (2018) Incentives in contests with heterogeneous solvers. *Management Science* 64(6):2709–2715.

Körpeoğlu E, Korpeoglu CG, Hafalir IE (2017) Parallel innovation contests. *Working Paper, University College London* .

Küçükgül C, Özer Ö, Wang S (2019) Engineering social learning: Information design of time-locked sales campaigns for online platforms. *Working Paper, University of Texas at Dallas.* .

Lemus J, Marshall G (2021) Dynamic tournament design: Evidence from prediction contests. *Journal of Political Economy* 129(2):383–420.

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

Mostagir M, Chen Y, Yeckehzaare I (2021) Information provision in dynamic contests: An experimental study. *Working Paper, University of Michigan* .

Nittala L, Krishnan V (2016) Designing internal innovation contests. *Working Paper, University of California San Diego* .

Rayo L (2013) Monopolistic signal provision. *The BE Journal of Theoretical Economics* 13(1):27–58.

Rayo L, Segal I (2010) Optimal information disclosure. *Journal of Political Economy* 118(5):949–987.

Schlapp J, Mihm J (2018) Optimal stochastic feedback in asymmetric dynamic contests. *INSEAD Working Paper* .

Shalpegin T, Sommer SC, Van Delft C (2020) Supplier collaboration in collaborative product development with internal competition. *Working Paper, University of Auckland.* .

Stouras KI, Erat S, Lichtendahl Jr KC (2020) Prizes on crowdsourcing platforms: An equilibrium analysis of competing contests. *Proceedings of the 21st ACM Conference on Economics and Computation*, 875–876.

Stouras KI, Hutchison-Krupat J, Chao RO (2019) The role of participation in innovation contests. *Working Paper, University College Dublin* .

Taylor CR (1995) Digging for golden carrots: An analysis of research tournaments. *The American Economic Review* 872–890.

Terwiesch C, Xu Y (2008) Innovation contests, open innovation, and multiagent problem solving. *Management Science* 54(9):1529–1543.

Appendix

A. Proofs

Proof of Proposition 1: We first calculate the expected duration of a contract in which both agents exert full effort at all times. We solve the problem by backward induction on the state of the game where the states are defined by the number of success for each agent. First, consider the state when both agents have already achieved one success, then the expected arrival time for the second success is given by:

$$\int_0^\infty 2\lambda t e^{-2\lambda t}dt = -te^{-2\lambda t}\Big|_0^\infty + \int_0^\infty e^{-2\lambda t}dt = -\frac{1}{2\lambda}e^{-2\lambda t}\Big|_0^\infty = \frac{1}{2\lambda}.$$

Next, consider the state of the game with a leader (an agent with one success) and a laggard (an agent with no success). Then the expected arrival time for the second success can be expressed as:

$$\int_0^\infty \left[\lambda t + \lambda\left(t + \frac{1}{2\lambda}\right)\right]e^{-2\lambda t}dt = \frac{1}{2\lambda} + \int_0^\infty \frac{1}{2}e^{-2\lambda t}dt = \frac{1}{2\lambda} + \frac{1}{4\lambda} = \frac{3}{4\lambda},$$

where with instantaneous probability λ the leader may obtain the second success at time t or the laggard may hit the first success at time t (proceeding to the above-mentioned state) in which case the expected duration of the contract is $t + 1/2\lambda$. Finally, considering the state when neither of the agents has one success, the expected duration of the contract is as follows:

$$\int_0^\infty 2\lambda(t + \frac{3}{4\lambda})e^{-2\lambda t}dt = \frac{1}{2\lambda} + \int_0^\infty \frac{3}{2}e^{-2\lambda t}dt = \frac{1}{2\lambda} + \frac{3}{4\lambda} = \frac{5}{4\lambda}. \tag{11}$$

Clearly, given the cost of effort at each instant, the principal has to offer each agent at least $5c/(4\lambda)$ so that each agent's ex-ante expected payoff is non-negative. ∎

Proof of Proposition 2: To derive the symmetric pure-strategy Nash equilibrium with full effort, let us fix agent $-i$'s effort $x_{k,l,t}^{-i} = 1$ for all k, l, and t and find conditions under which agent i optimally chooses $x_{k,l,t}^i = 1$ for all k, l, and t. For notational simplicity, we drop the superscript i. Consider the state of the game where both agents have already achieved one success, using (3) we can write:

$$V_{1,1,t} = \max_{x_{1,1,\tau}} \int_t^\infty x_{1,1,\tau}(\lambda R - c)e^{-\int_t^\tau \lambda(x_{1,1,s}+1)ds}d\tau.$$

The agent's problem is an infinite horizon problem, so it is stationary. Thus, we can drop the subscript t and write the equivalent Bellman equation for the agent's problem as follows:

$$V_{1,1} = \max_{x_{1,1}} \left\{x_{1,1}(\lambda R - c)dt + (1 - \lambda x_{1,1}dt - \lambda dt)V_{1,1}\right\},$$

Note that $-cx_{1,1}dt$ denotes the agent's cost of effort within the time interval $(t, t+dt)$, while $x_{1,1}\lambda dt$ denotes the probability that a success arrives within $(t, t+dt)$, in which case the agent receives

R. On the other hand, the probability that her opponent achieves the second success in that time interval is λdt, and in this case, the agent receives zero reward. With probability $(1 - \lambda x_{1,1}dt - \lambda dt)$, neither the agent nor her opponent achieves the second success, in which case the contest continues and the agent anticipates to receive a continuation payoff of $V_{1,1}$ due to stationarity. Simplifying the above expression and dividing both sides by dt, we obtain the following Hamilton-Jacobi- Bellman (hereafter HJB) equation for the agent's problem:

$$0 = \max_{x_{1,1}} \left\{ \underbrace{-cx_{1,1}}_{cost} + \underbrace{\lambda x_{1,1}\left(R - V_{1,1}\right)}_{benefit} - \underbrace{\lambda V_{1,1}}_{externality} \right\}. \tag{12}$$

First, second, and third terms reflect the agent's flow cost of effort, her flow benefit from effort, and the externality imposed by her opponent's effort, respectively. Since the HJB in (12) is linear in $x_{1,1}$, it can be concluded that $x_{1,1} = 1$ is optimal if and only if

$$R - V_{1,1} \geq \frac{c}{\lambda}. \tag{13}$$

The above condition implies that each agent finds it optimal to work if the principal rewards the agent with additional utility of at least c/λ upon the arrival of a success.

Next, consider the state of the game where agent i is the leader with one success and agent $-i$ is the laggard with no success. Bellman and HJB equations for agent i can be expressed as follows:

$$V_{1,0} = \max_{x_{1,0}} \left\{ x_{1,0}(\lambda R - c)dt + \lambda V_{1,1}dt + (1 - \lambda x_{1,0}dt - \lambda dt)V_{1,0} \right\}$$

$$\Rightarrow 0 = \max_{x_{1,0}} \left\{ x_{1,0}\left(\lambda R - c - \lambda V_{1,0}\right) + \lambda\left(V_{1,1} - V_{1,0}\right) \right\}, \tag{14}$$

The first line admits a similar interpretation as in the previous case, except that if the laggard (agent $-i$) obtains a success, the leader agent i receives a continuation payoff equal to $V_{1,1}$. From (14), we can derive the Incentive Compatibility (hereafter IC) constraint for agent i which tells us that $x_{1,0} = 1$ is incentive compatible if and only if

$$R - V_{1,0} \geq \frac{c}{\lambda}. \tag{15}$$

When agent i is the laggard with no success and agent $-i$ is the leader with one success, we can rewrite agent i's problem in (4) as follows:

$$V_{0,1} = \max_{x_{0,1}} \left\{ x_{0,1}(\lambda V_{1,1} - c)dt + (1 - \lambda x_{0,1}dt - \lambda dt)V_{0,1} \right\}$$

$$\Rightarrow 0 = \max_{x_{0,1}} \left\{ x_{0,1}\left(\lambda V_{1,1} - c - \lambda V_{0,1}\right) - \lambda V_{0,1} \right\}, \tag{16}$$

which implies that exerting full effort for the laggard is optimal if and only if the following IC constraint holds:

$$V_{1,1} - V_{0,1} \geq \frac{c}{\lambda}. \tag{17}$$

Finally, before the arrival of any success, the continuation payoff of agent i is given by:

$$V_{0,0} = \max_{x_{0,0}} \{x_{0,0}(\lambda V_{1,0} - c)dt + \lambda V_{0,1}dt + (1 - \lambda x_{0,0}dt - \lambda dt)V_{0,0}\}$$

$$\Rightarrow 0 = \max_{x_{0,0}} \{x_{0,0}(\lambda V_{1,0} - c - \lambda V_{0,0}) + \lambda(V_{0,1} - V_{0,0})\}. \tag{18}$$

From (18), exerting $x_{0,0} = 1$ is incentive compatible for agent i if and only if

$$V_{1,0} - V_{0,0} \geq \frac{c}{\lambda}. \tag{19}$$

We are now ready to show that $R = 3c/\lambda$ is the minimum required fixed reward to induce agent i (and by symmetry agent $-i$ as well) to exert full effort at all times. From (17), $V_{1,1} \geq c/\lambda$ since $V_{0,1}$ has to be non-negative. Also, from (12), under full effort, one can verify that $V_{1,1} = \frac{1}{2}\left(R - \frac{c}{\lambda}\right)$. Combining these together, we require $\frac{1}{2}\left(R - \frac{c}{\lambda}\right) \geq \frac{c}{\lambda}$, which boils down to $R \geq 3c/\lambda$. Thus, we need $R = 3c/\lambda$ at the minimum to ensure that (17) is satisfied, and hence it is incentive compatible for the laggard to exert full effort.

It remains to show that $R = 3c/\lambda$ satisfies all IC constraints. It is straightforward to check that (13) is satisfied, that is $R - V_{1,1} = \frac{3c}{\lambda} - \frac{c}{\lambda} > \frac{c}{\lambda}$. Plugging in the value of $V_{1,1} = c/\lambda$ into (16), we find that $V_{0,1} = 0$ and so the IC constraint in (17) for the laggard is binding. Similarly, plugging in the value of $V_{1,1} = c/\lambda$ into (14), it can be concluded that $V_{1,0} = 3c/(2\lambda)$ and so the IC constraint in (15) for the leader is satisfied since $R - V_{1,0} = \frac{3c}{\lambda} - \frac{3c}{2\lambda} = \frac{3c}{2\lambda} > \frac{c}{\lambda}$. Finally, plugging in the values of $V_{1,0} = 3c/(2\lambda)$ and $V_{0,1} = 0$ into (18), one can verify that $V_{0,0} = c/(4\lambda)$ and so the IC constraint in (19) for each agent is satisfied as $V_{1,0} - V_{0,0} = \frac{3c}{2\lambda} - \frac{c}{4\lambda} = \frac{5c}{4\lambda} > \frac{c}{\lambda}$. ∎

Proof of Theorem 1: Consider a flexible-reward contest with $R_{2,0} = 2c/\lambda$ and $R_{2,1} = 3c/\lambda$ where the principal commits to disclose any success upon its arrival. Similar to the previous case, we analyze the problem by moving backward on the state of the game where the states are defined by the number of successes of the agents. Let us fix agent $-i$'s effort $x_{k,l,t}^{-i} = 1$ for all k, l, and t and find conditions under which agent i optimally chooses $x_{k,l,t}^{i} = 1$ for all k, l, and t. Consider the state of the game where both agents have already achieved one success. The Bellman equation and the corresponding HJB for agent i's problem can be expressed as follows:

$$V_{1,1} = \max_{x_{1,1}} \{x_{1,1}(\lambda R_{2,1} - c)dt + (1 - \lambda x_{1,1}dt - \lambda dt)V_{1,1}\}$$

$$\Rightarrow 0 = \max_{x_{1,1}} \{x_{1,1}(\lambda R_{2,1} - c - \lambda V_{1,1}) - \lambda V_{1,1}\}, \tag{20}$$

where we use the fact that the winner receives $R_{2,1}$ in this state of the game. From (20), we can derive that $x_{1,1} = 1$ is optimal if and only if

$$R_{2,1} - V_{1,1} \geq \frac{c}{\lambda}. \tag{21}$$

Next, consider the state of the game with a leader and a laggard. The Bellman equation and the corresponding HJB for the leader's problem (which we assume to be agent i) can be written as:

$$V_{1,0} = \max_{x_{1,0}} \left\{ x_{1,0}(\lambda R_{2,0} - c)dt + \lambda V_{1,1}dt + (1 - \lambda x_{1,0}dt - \lambda dt)V_{1,0} \right\}$$

$$\Rightarrow 0 = \max_{x_{1,0}} \left\{ x_{1,0} \left(\lambda R_{2,0} - c - \lambda V_{1,0} \right) + \lambda \left(V_{1,1} - V_{1,0} \right) \right\}, \tag{22}$$

where we use the fact that the winner receives $R_{2,0}$ in this state of the game. From (22), we can derive the IC constraint for the leader which tells us that $x_{1,0} = 1$ is incentive compatible if and only if

$$R_{2,0} - V_{1,0} \geq \frac{c}{\lambda}. \tag{23}$$

Similarly, we can express the Bellman equation and the corresponding HJB for the laggard's problem (assuming to be agent i) as follows:

$$V_{0,1} = \max_{x_{0,1}} \left\{ x_{0,1}(\lambda V_{1,1} - c)dt + (1 - \lambda x_{0,1}dt - \lambda dt)V_{0,1} \right\}$$

$$\Rightarrow 0 = \max_{x_{0,1}} \left\{ x_{0,1} \left(\lambda V_{1,1} - c - \lambda V_{0,1} \right) - \lambda V_{0,1} \right\}, \tag{24}$$

which implies that exerting full effort for the laggard is optimal if and only if the following IC constraint holds

$$V_{1,1} - V_{0,1} \geq \frac{c}{\lambda}. \tag{25}$$

Finally, before the arrival of any success, the continuation value of each agent is given by:

$$V_{0,0} = \max_{x_{0,0}} \left\{ x_{0,0}(\lambda V_{1,0} - c)dt + \lambda V_{0,1}dt + (1 - \lambda x_{0,0}dt - \lambda dt)V_{0,0} \right\}$$

$$\Rightarrow 0 = \max_{x_{0,0}} \left\{ x_{0,0} \left(\lambda V_{1,0} - c - \lambda V_{0,0} \right) + \lambda \left(V_{0,1} - V_{0,0} \right) \right\}. \tag{26}$$

From (26), exerting $x_{0,0} = 1$ is incentive compatible for each agent if and only if

$$V_{1,0} - V_{0,0} \geq \frac{c}{\lambda}. \tag{27}$$

We now verify that the proposed flexible-reward schedule in Theorem 1 satisfies all of the above IC constraints and spends the minimum first-best expected reward. Given (25), we can see that $V_{1,1} = c/\lambda$ is the minimum required continuation payoff to incentivize the laggard to put full effort. From (20), we know that $V_{1,1} = \frac{1}{2} \left(R_{2,1} - \frac{c}{\lambda} \right)$. Thus, the principal has to specify a reward $R_{2,1} = 3c/\lambda$ in order to satisfy $V_{1,1} = c/\lambda$. Given these values, it is straightforward to check that the IC constraint

in (21) is satisfied, that is $R_{2,1} - V_{1,1} = \frac{3c}{\lambda} - \frac{c}{\lambda} > \frac{c}{\lambda}$. Also, plugging in the value of $V_{1,1} = c/\lambda$ into (24), we obtain that $V_{0,1} = 0$ and so the IC constraint for the laggard is binding. Next, from (27), we can conclude that $V_{1,0} = c/\lambda$ is the minimum required continuation payoff to motivate the agents to exert effort. Plugging in this value into (22), $R_{2,0} = 2c/\lambda$ is needed to satisfy the HJB. It follows that the IC constraint in (23) is indeed binding for the leader as $R_{2,0} - V_{1,0} = \frac{2c}{\lambda} - \frac{c}{\lambda} = \frac{c}{\lambda}$. Finally, given $V_{1,0} = c/\lambda$ and $V_{0,1} = 0$, we conclude by (26) that $V_{0,0} = 0$ which shows that the last IC constraint in (27) is also binding, that is $V_{1,0} - V_{0,0} = \frac{c}{\lambda} - 0 = \frac{c}{\lambda}$. Therefore, full effort is incentive compatible at all times.

To calculate the expected reward of the contest with flexible reward, note that when both agents have already obtained one success, the expected reward of the contest is given by $R_{2,1} = 3c/\lambda$. When there is a leader and a laggard, the expected reward can be computed as follows:

$$\int_t^\infty \lambda \left(\frac{2c}{\lambda} + \frac{3c}{\lambda} \right) e^{-2\lambda(\tau - t)} d\tau = \frac{5c}{2\lambda}.$$

To interpret the above equation note that if the leader obtains her second success, the reward is $R_{2,0} = 2c/\lambda$ and if the laggard obtains her first success, the state of the game transitions to the case where both agents have already obtained one success and the reward is adjusted upward to $R_{2,1} = 3c/\lambda$. Finally, the ex-ante expected reward of the contest is given by:

$$\int_0^\infty 2\lambda \left(\frac{5c}{2\lambda} \right) e^{-2\lambda t} dt = \frac{5c}{2\lambda}. \blacksquare$$

Proof of Proposition 3: To derive the symmetric pure-strategy Nash equilibrium with full effort, we shall fix the opponent's effort $x_{k,t}^{-i} = 1$ for all k and t and try to find conditions under which agent i best-responds by choosing $x_{k,t}^i = 1$ for all k, t.

Consider the problem faced by an agent who has not yet achieved a success. Dropping the superscript i in (7) by using the symmetry of agents, the equivalent Bellman equation for the agent's problem is as follows:

$$V_{0,t} = \max_{x_{0,t}} \left\{ -cx_{0,t}dt + x_{0,t}\lambda V_{1,t}dt + (1 - x_{0,t}\lambda dt - p_t\lambda dt)V_{0,t+dt} \right\}. \tag{28}$$

Note that $cx_{0,t}dt$ denotes the agent's cost of effort within the time interval $(t, t+dt)$, while $x_{0,t}\lambda dt$ denotes the probability that a success arrives within $(t, t+dt)$, in which case the agent receives a continuation payoff, $V_{1,t}$. On the other hand, the probability that her opponent is in the second stage and achieves a success in this time interval is $p_t\lambda dt$, and in that case, the agent receives a continuation value of zero. With probability $(1 - x_{0,t}\lambda dt - p_t\lambda dt)$, neither the agent achieves a success, nor does her opponent achieve the second success, in which case the contest continues and the agent anticipates to receive her continuation payoff, $V_{0,t+dt}$. Given that we have an infinite

horizon dynamic model with no deadline, from (28) one can verify that the continuation payoff solely depends on the probability p_t rather than time itself. Thus, we can define a stationary Bellman function $V_{k,p}$ for $k \in \{0, 1\}$ that does not depend on time but depends on the current state of p_t. Let p be a state variable that corresponds to the probability that each agent assigns to the fact that her opponent is in the second stage under no information disclosure. Then, we can express each agent's continuation payoff as $V_{k,p}$. Thus, we can rewrite (28) as follows:

$$V_{0,p} = \max_{x_{0,p}} \left\{ -cx_{0,p}dt + x_{0,p}\lambda V_{1,p}dt + (1 - x_{0,p}\lambda dt - p\lambda dt)V_{0,p+dp} \right\}. \tag{29}$$

Using a Taylor expansion (Ito's Lemma), we have

$$V_{0,p+dp} \simeq V_{0,p} + V'_{0,p}dp = V_{0,p} + \lambda(1-p)^2 V'_{0,p}dt,$$

where we have used that $x_{k,t}^{-i} = 1$ and $dp = \lambda(1-p)^2 dt$ according to (5). Substituting this expression into (29), dropping the terms of the order dt^2 (since $dt^2 \simeq 0$), canceling terms and dividing both sides by dt, we obtain the following HJB equation for the agent's problem:

$$0 = \max_{x_{0,p}} \left\{ \underbrace{-cx_{0,p}}_{cost} + \underbrace{x_{0,p}\lambda \left(V_{1,p} - V_{0,p} \right)}_{benefit} - \underbrace{\lambda \left[pV_{0,p} - (1-p)^2 V'_{0,p} \right]}_{externality} \right\}. \tag{30}$$

Note that the first term reflects the agent's flow cost of effort, the second term reflects her flow benefit from effort, and the third term captures the externality imposed by her opponent's effort. Since the HJB in (30) is linear in $x_{0,p}$, it can be concluded that $x_{0,p} = 1$ is optimal if and only if

$$V_{1,p} - V_{0,p} \geq \frac{c}{\lambda}. \tag{31}$$

The above IC constraint implies that an agent with no success finds it optimal to work if the principal rewards the agent with additional utility of at least c/λ upon the arrival of a success.

Next, consider the problem faced by an agent who has achieved one success as formulated in (6). Since the continuation payoffs of agents depend on the state variable p rather than time, the principal's problem is also stationary (i.e., independent of t) and hence it is optimal for the principal to choose a reward schedule that depends only on p. In other words, an agent who achieves two successes first is rewarded R_p, where p is her belief about her opponent's progress. As a result, after dropping the superscript i in (6) by using the symmetry of agents, the corresponding Bellman equation for the agent's problem is given by:

$$V_{1,p} = \max_{x_{1,p}} \left\{ -cx_{1,p}dt + x_{1,p}\lambda R_p dt + (1 - x_{1,p}\lambda dt - p\lambda dt)V_{1,p+dp} \right\}$$

which using the previous techniques gives us the following HJB eqution

$$0 = \max_{x_{1,p}} \left\{ \underbrace{-cx_{1,p}}_{cost} + \underbrace{x_{1,p}\lambda\left(R_p - V_{1,p}\right)}_{benefit} - \underbrace{\lambda\left[pV_{1,p} - (1-p)^2 V'_{1,p}\right]}_{externality} \right\}. \tag{32}$$

Since the HJB in (32) is linear in $x_{1,p}$, it can be concluded that $x_{1,p} = 1$ is optimal if and only if

$$R_p - V_{1,p} \geq \frac{c}{\lambda}. \tag{33}$$

We are now ready to prove that full effort is incentive compatible at all times given the proposed flexible-reward schedule in Proposition 3. First, notice that when we fix the opponent's effort $x_{k,t}^{-i} = 1$ for all t and solve (5) with initial condition $p_0^i = 0$, we obtain $p_t = \lambda t/(1 + \lambda t)$ as stated in the proposition. Second, note that if an agent with no success receives a continuation payoff $V_{1,p} = c/\lambda$, $\forall p$, by substituting this value into the integral form of the agent's problem in (7), we obtain $V_{0,p} = 0$. Hence, (31) is always binding. Moreover, if $V_{1,p} = c/\lambda$, the flexible-reward schedule $R_p = (2 + p)c/\lambda$ always satisfies (33). Plugging in $R_p = (2+p)c/\lambda$ into (32), it can be verified that $V_{1,p} = c/\lambda$ for all p is a solution. Finally, plugging in $V_{1,p} = c/\lambda$ into (30), it can be verified that $V_{0,p} = 0$ for all p is a solution. Therefore, the design is incentive compatible at all times.

Next, to prove the optimality of this design, we provide a simple surplus argument. Notice that agents' and principal's combined surplus is the same across any design that implements full effort at all times. Therefore, the principal's surplus is maximized when the agents' surplus is minimized. The minimum continuation payoff for an agent who has not yet achieved a success is $V_{0,p} = 0$. In order to satisfy the IC constraint in (31), the principal has to offer the minimum continuation payoff $V_{1,p} = c/\lambda$ to an agent upon the arrival of a success and by (32), this value uniquely gives $R_p = (2 + p)c/\lambda$. Therefore, this design minimizes the agents' surplus and hence maximizes the principal's surplus.

Finally, we can verify that this design spends the first-best expected reward. To show this, we compute the expected reward that the principal has to pay under this design. Denote by $R_{k,l}$ the principal's expected payout conditional on the first agent having achieved $k \in \{0, 1\}$ successes, and the second agent having achieved $l \in \{0, 1\}$ successes. Let us consider the state of the game when both agents have already achieved one success, then the expected payout is given by:

$$R_{1,1,t} = \int_t^\infty 2\lambda \left(2 + \frac{\lambda\tau}{1 + \lambda\tau}\right) \frac{c}{\lambda} e^{-2\lambda(\tau - t)} d\tau = \frac{3c}{\lambda} - \frac{2c}{\lambda} e^{2(1+\lambda t)} \int_{2(1+\lambda t)}^\infty \frac{e^{-x}}{x} dx,$$

where the first equality can be interpreted as follows: if any agent obtains the second success during interval $(\tau, \tau + dt)$ which happens with probability $2\lambda dt$, the principal has to pay $(2 + p_\tau)c/\lambda$ to the winner, provided that none of the agents have already obtained the second success by time τ which is captured by the term $e^{-2\lambda(\tau - t)}$, and the second equality is obtained by change of variables.

Next, consider the state of the game with a leader and a laggard. Then the expected payout can be computed as follows:

$$R_{1,0,t} = \int_t^\infty \lambda \left[\left(2 + \frac{\lambda\tau}{1+\lambda\tau} \right) \frac{c}{\lambda} + R_{1,1,\tau} \right] e^{-2\lambda(\tau-t)} d\tau$$

$$= \int_t^\infty \left(2 + \frac{\lambda\tau}{1+\lambda\tau} \right) c e^{-2\lambda(\tau-t)} d\tau + \int_t^\infty \lambda R_{1,1,\tau} e^{-2\lambda(\tau-t)} d\tau$$

$$= \frac{1}{2} R_{1,1,t} + \int_t^\infty \lambda \left[\frac{3c}{\lambda} - \frac{2c}{\lambda} e^{2(1+\lambda\tau)} \int_{2(1+\lambda\tau)}^\infty \frac{e^{-x}}{x} dx \right] e^{-2\lambda(\tau-t)} d\tau$$

$$= \frac{3c}{\lambda} - \frac{c}{\lambda} e^{2(1+\lambda t)} \int_{2(1+\lambda t)}^\infty \frac{e^{-x}}{x} dx - 2ce^{2(1+\lambda t)} \int_t^\infty \int_{2(1+\lambda\tau)}^\infty \frac{e^{-x}}{x} dx d\tau$$

$$= \frac{2c}{\lambda} + \frac{c}{\lambda} e^{2(1+\lambda t)} (1+2\lambda t) \int_{2(1+\lambda t)}^\infty \frac{e^{-x}}{x} dx.$$

Finally, starting from time zero, the expected reward of the contest is given by:

$$R_{0,0,0} = \int_0^\infty 2\lambda R_{1,0}(t) e^{-2\lambda t} dt$$

$$= \int_0^\infty \left[4ce^{-2\lambda t} + 2ce^2(1+2\lambda t) \int_{2(1+\lambda t)}^\infty \frac{e^{-x}}{x} dx \right] dt$$

$$= \frac{2c}{\lambda} + \frac{c}{2\lambda} = \frac{5c}{2\lambda}. \blacksquare$$

Proof of Proposition 4: We prove the proposition in multiple steps.

Step 1: *We first verify that the strategy of the agents in the proposition forms a symmetric equilibrium.*

To check this, we fix the strategy of agent $-i$ to the proposed one in the proposition and verify that agent i best-responds by playing the same strategy. First, using condition (33), it is easy to verify that exerting full effort is incentive compatible for an agent with one success for all p since $V_{1,p}$ can not exceed $\overline{R} - c/\lambda$. Given that agent $-i$ exerts effort $x_{0,t}^{-i} = p_r$ for $t \geq t_r$, and $x_{1,t}^{-i} = 1$, by (5) we obtain $\dot{p}_t^i = 0$. As a result, $p_t^i = p_r$ for $t \geq t_r$. Following this observation, note that if agent i with no success, holding a belief p_r, receives a continuation payoff $V_{1,p_r} = c/\lambda$, by substituting this value into the integral form of the agent's problem in (7), we get $V_{0,p_r} = 0$. Hence, the incentive compatibility condition in (31) is binding for all $t \geq t_r$ implying that agent i is indifferent between any level of effort and so exerting $x_{0,p_r} = p_r$ is optimal. Plugging in $\overline{R} = (2+p_r)c/\lambda$ into (32), it can be verified that $V_{1,p} = c/\lambda$ is a solution for all $t \geq t_r$ where $p_t = p_r$. Finally, to prove that exerting full effort is optimal for agent i with no success for all $p < p_r$, we move backward from time t_r associated with belief p_r and prove that if the agent finds it optimal to exert strictly positive effort at any belief p' where $p \leq p' \leq p_r$ (i.e., if $V_{1,p'+dp} - V_{0,p'+dp} \geq c/\lambda$), then we have $V_{1,p} - V_{0,p} \geq c/\lambda$

implying that exerting full effort is optimal at belief $p - dp$. This can be seen by the following analysis:

$$V_{1,p} - V_{0,p} =$$

$$-cdt + \lambda \overline{R}dt + (1 - \lambda dt - p\lambda dt)V_{1,p+dp} + cdt - \lambda V_{1,p}dt - (1 - \lambda dt - p\lambda dt)V_{0,p+dp} \geq$$

$$\lambda \overline{R}dt + (1 - \lambda dt - p\lambda dt)\frac{c}{\lambda} - \lambda V_{1,p}dt \geq \frac{c}{\lambda},$$

where the last inequality results from the fact that

$$V_{1,p} = \int_t^\infty (\lambda \overline{R} - c)e^{-\int_t^\tau \lambda(1+p_s)ds}d\tau \leq \int_t^\infty (\lambda \overline{R} - c)e^{-\lambda(1+p)(\tau-t)}d\tau$$

$$= \frac{\lambda \overline{R} - c}{\lambda(1+p)} \leq \overline{R} - (1+p)\frac{c}{\lambda},$$

where the first line results from the fact that p_t is weakly increasing and the second line holds if and only if $(2+p)c/\lambda \leq \overline{R}$ which is satisfied for $p \leq p_r$. This verifies the equilibrium. Next, we prove the uniqueness of the symmetric equilibrium.

Step 2: *Let p_r solve $(2+p_r)c/\lambda = \overline{R}$. There is no symmetric equilibrium in which an agent with no success exerts full effort at some $p > p_r$.*

First note that $x_{0,p} < 1$ for some p. This is because as p approaches 1, $V_{1,p}$ approaches $V_{1,1} = \frac{1}{2}(\overline{R} - c/\lambda) < c/\lambda$ given the budget constraint, where we use the fact that an agent with one success exerts full effort at all times. Then, suppose t is the first time at which the belief of an agent with no success reaches its maximum level ($\overline{p}$) in a symmetric equilibrium, and $\overline{p} > p_r$. Let us focus on a region where the belief is strictly increasing and reaches $\overline{p}$ for the first time. If the agent exerts full effort at $\overline{p}$, by (5) p strictly increases which is a contradiction. Therefore, we must have $V_{1,\overline{p}+dp} - V_{0,\overline{p}+dp} \leq c/\lambda$. This condition implies that exerting zero effort is optimal at belief $\overline{p}$. Then, we consider the following two cases:

- $V_{1,\overline{p}+dp} - V_{0,\overline{p}+dp} = c/\lambda$: To find the agent's optimal effort at belief $\overline{p} - dp$, given that $x_{0,\overline{p}} = 0$, we write the following:

$$V_{1,\overline{p}} - V_{0,\overline{p}} =$$

$$(\lambda \overline{R} - c)dt + (1 - \lambda dt - \overline{p}\lambda dt)V_{1,\overline{p}+dp} - (1 - \overline{p}\lambda dt)V_{0,\overline{p}+dp} =$$

$$V_{1,\overline{p}+dp} - V_{0,\overline{p}+dp} + \left[\lambda \overline{R} - c - \lambda V_{1,\overline{p}+dp} - \overline{p}\lambda \left(V_{1,\overline{p}+dp} - V_{0,\overline{p}+dp}\right)\right]dt$$

$$= \frac{c}{\lambda} + (\lambda \overline{R} - c - \lambda V_{1,\overline{p}+dp} - \overline{p}c)dt < \frac{c}{\lambda},$$

where the last inequality results from the fact that $V_{1,\overline{p}+dp} \geq c/\lambda$ and $\overline{R} = (2 + p_r)c/\lambda < (2 + \overline{p})c/\lambda$. Thus, an agent with no success exerts zero effort at belief $\overline{p} - dp$ and by (5) p decreases which is a contradiction.

- $V_{1,\overline{p}+dp} - V_{0,\overline{p}+dp} < c/\lambda$: From the previous case, we know that an agent with no success must exert full effort at belief $\overline{p} - dp$ which requires $V_{1,\overline{p}} - V_{0,\overline{p}} \geq c/\lambda$. By continuity of $V_{1,p} - V_{0,p}$, we conclude that $V_{1,\overline{p}} - V_{0,\overline{p}} = c/\lambda$. Doing the same analysis as before, we find that the agent finds it optimal to put zero effort at belief $\overline{p} - 2dp$ which violates the assumption that the agent's belief is strictly increasing in this region.

Step 3: *Let p_r solve $(2+p_r)c/\lambda = \overline{R}$. There is no symmetric equilibrium in which an agent with no success does not exert full effort at some $p < p_r$.*

Suppose τ is the first time that the agent with no success does not exert full effort. Let $t > \tau$ be the first time at which the belief of an agent with no success in the equilibrium reaches its minimum level $(\underline{p})$ and $\underline{p} < p_r$. Let us focus on a region where the agent's belief is strictly decreasing and reaches $\underline{p}$ for the first time. If the agent exerts zero effort at belief $\underline{p}$, by (5) p strictly decreases which is a contradiction. Therefore, we must have $V_{1,\underline{p}+dp} - V_{0,\underline{p}+dp} \geq c/\lambda$. This condition implies that exerting full effort is optimal at belief $\underline{p}$. Then we consider two cases:

- $V_{1,\underline{p}+dp} - V_{0,\underline{p}+dp} = c/\lambda$: To find the agent's optimal effort at belief $\underline{p} - dp$, given that $x_{0,\underline{p}} = 1$, we can write the following:

$$V_{1,\underline{p}} - V_{0,\underline{p}} =$$

$$(\lambda\overline{R} - c)dt + (1 - \lambda dt - \underline{p}\lambda dt)V_{1,\underline{p}+dp} - (\lambda V_{1,\underline{p}+dp} - c)dt - (1 - \lambda dt - \underline{p}\lambda dt)V_{0,\underline{p}+dp} =$$

$$V_{1,\underline{p}+dp} - V_{0,\underline{p}+dp} + \left[\lambda\overline{R} - \lambda V_{1,\underline{p}+dp} - (1+\underline{p})\lambda\left(V_{1,\underline{p}+dp} - V_{0,\underline{p}+dp}\right)\right]dt =$$

$$\frac{c}{\lambda} + \left[\lambda\overline{R} - \lambda V_{1,\underline{p}+dp} - (1+\underline{p})c\right]dt > \frac{c}{\lambda},$$

where the last inequality results from the fact that

$$V_{1,p} = \int_t^\infty (\lambda\overline{R} - c)e^{-\int_t^\tau \lambda(1+p_s)ds}d\tau \leq \int_t^\infty (\lambda\overline{R} - c)e^{-\lambda(1+\underline{p})(\tau-t)}d\tau$$

$$= \frac{\lambda\overline{R} - c}{\lambda(1+\underline{p})} \leq \overline{R} - (1+\underline{p})\frac{c}{\lambda},$$

where the first line results from the fact that $\underline{p}$ is the minimum belief in the equilibrium and the second line holds since $(2 + \underline{p})c/\lambda \leq \overline{R}$. Therefore, the agent exerts full effort at belief $\underline{p} - dp$ and by (5) p strictly increases which is a contradiction.

- $V_{1,\underline{p}+dp} - V_{0,\underline{p}+dp} > c/\lambda$: From the previous case, we know that an agent with no success must exert zero effort at belief $\underline{p} - dp$ which requires $V_{1,\underline{p}} - V_{0,\underline{p}} \leq c/\lambda$. By continuity of $V_{1,p} - V_{0,p}$, we conclude that $V_{1,\underline{p}} - V_{0,\underline{p}} = c/\lambda$. Doing the same analysis as before, we find that the agent finds it optimal to put full effort at belief $\underline{p} - 2dp$ which violates the assumption that the agent's belief is strictly decreasing in this region.

From steps 2 and 3, we conclude that the symmetric equilibrium in the proposition is unique. ∎

Proof of Proposition 5:   We build on the proofs of Propositions 4 and B1. As before, let us fix the strategy of agent $-i$ to the proposed one in the proposition and verify the best response of agent i. Consider the very last instant of the first cycle at which the belief of agent i reaches p_r. The IC condition (31) implies that full effort is optimal for an agent with no success if and only if $V_{1,p} - V_{0,p} \geq c/\lambda$. We can rewrite this condition at time t_r associated with belief p_r as follows:

$$p_r V_{1,1} + (1 - p_r)V_{1,quit} - (1 - p_r)V_{0,0} \geq c/\lambda. \tag{34}$$

The above condition can be interpreted as follows: if agent i obtains her first success at t_r, her expected continuation payoff is given by $p_r V_{1,1} + (1 - p_r)V_{1,quit}$ anticipating that the principal discloses full information at the end of the cycle. Therefore, with probability p_r her opponent has already made partial progress which in that case they keep working until the end and the continuation payoff is $V_{1,1}$, or her opponent quits if she has not obtained any success and the continuation payoff is $V_{1,quit}$. On the other hand, if agent i does not succeed at t_r, she quits if her opponent has progressed to the second stage. Otherwise, the contest and the beliefs reset and a new cycle begins with a continuation payoff of $V_{0,0}$.

Given that an agent with one success always puts full effort until the end, we know $V_{1,1} = \frac{1}{2}(\overline{R} - c/\lambda)$ and $V_{1,quit} = \overline{R} - c/\lambda$. Moreover, the upper bound for $V_{0,0}$ is given by $V_{0,0}^F$, where F stands for full information, which is the continuation payoff if full information is provided during each cycle. To see this, suppose that full information is provided during each cycle. We consider two cases: i) if agent i obtains the first success, her opponent immediately quits. This leads to a higher continuation payoff than the case of silent period where the opponent keeps working until the end of the cycle; ii) if agent i's opponent obtains the first success, agent i's best response is to quit. However, in a silent period, agent i earns a negative ex-post payoff. Therefore, the upper bound for $V_{0,0}$ is given by $V_{0,0}^F = \frac{1}{2}(\overline{R} - 2c/\lambda)$. Plugging in these values into (34), it is easy to verify that the condition is binding implying that full effort is optimal.

Finally, to prove that exerting full effort is optimal for agent i with no success during each cycle where $p < p_r$, we can show that if the agent finds it optimal to exert strictly positive effort at any belief p' where $p \leq p' \leq p_r$ (i.e., if $V_{1,p'+dp} - V_{0,p'+dp} \geq c/\lambda$), then we have $V_{1,p} - V_{0,p} \geq c/\lambda$ implying that exerting full effort is optimal at belief $p - dp$. This can be seen by the following analysis:

$$V_{1,p} - V_{0,p} =$$

$$\int_t^{t_r} (\lambda\overline{R} - c)e^{-\int_t^{\tau} \lambda(1+p_s)ds}d\tau + \left[p_r V_{1,1} + (1-p_r)V_{1,quit}\right]e^{-\int_t^{t_r} \lambda(1+p_s)ds}$$

$$- \int_t^{t_r} (\lambda V_{1,\tau} - c)e^{-\int_t^{\tau} \lambda(1+p_s)ds}d\tau - (1-p_r)V_{0,0}e^{-\int_t^{t_r} \lambda(1+p_s)ds}$$

$$= \int_t^{t_r} \lambda \left(\overline{R} - V_{1,\tau}\right) e^{-\int_t^\tau \lambda(1+p_s)ds} d\tau + [p_r V_{1,1} + (1-p_r)V_{1,quit} - (1-p_r)V_{0,0}] e^{-\int_t^{t_r} \lambda(1+p_s)ds} \geq \frac{c}{\lambda},$$

where the last inequality can be verified after plugging in the values of $V_{1,p}$, $V_{1,1}$, $V_{1,quit}$, and $V_{0,0}$ into the above expression, computing the above integral and some tedious algebra. ∎

Proof of Proposition 6: Let us fix the strategy of agent $-i$ to the proposed one in the proposition and verify that agent i best-responds by playing the same strategy. We already show that an agent with one success finds it optimal to put full effort if and only if

$$\overline{R} - V_{1,p} \geq \frac{c}{\lambda}, \tag{35}$$

which always holds as $V_{1,p} \leq \overline{R} - c/\lambda$. Using p as the state variable, consider the Bellman equation for the maximization problem of agent i with no success as follows:

$$V_{0,p} = \max_{x_{0,p}} \left\{-cx_{0,p}dt + x_{0,p}\lambda V_{1,p}dt + (1 - x_{0,p}\lambda dt - p\lambda dt - p\gamma_p dt)V_{0,p+dp}\right\}.$$

Using the same techniques as before, we can derive the following HJB equation:

$$0 = \max_{x_{0,p}} \left\{-cx_{0,p} + x_{0,p}\lambda \left(V_{1,p} - V_{0,p}\right) - p\lambda V_{0,p} - p\gamma_p V_{0,p} + (1-p)(\lambda - p\lambda - p\gamma_p)V'_{0,p}\right\}.$$

Therefore, the IC constraint for an agent with no success implies that $x_{0,p} = 1$, if and only if

$$V_{1,p} - V_{0,p} \geq c/\lambda, \tag{36}$$

which is similar to (31). To derive the expected continuation payoff of agent i, holding a belief p, upon the arrival of the first success, we can write:

$$V_{1,p} = pV_{1,1} + (1-p)V_{1,0}. \tag{37}$$

where $V_{1,1} = \frac{1}{2}(\overline{R} - c/\lambda)$ is the expected continuation payoff if the opponent has already progressed to the second stage, and $V_{1,0}$ is the expected continuation payoff if the opponent has not progressed to the second stage. Given the probabilistic encouragement rate of information disclosure $\gamma_r = \lambda(1-p_r)/p_r = 1/t_r$, by (8) we obtain $p_t = p_r$ remains constant for $t \geq t_r$. Therefore, at any threshold belief $p = p_r$, we have:

$$V_{1,0} = \int_t^\infty \left[(\lambda\overline{R} - c) + \lambda\frac{1}{2}\left(\overline{R} - \frac{c}{\lambda}\right) + \gamma_r\left(\overline{R} - \frac{c}{\lambda}\right)\right] e^{-(2\lambda+\gamma_r)(\tau-t)}d\tau,$$

given that during interval $(\tau, \tau+dt)$, the leader puts full effort and earns $(\lambda\overline{R} - c)dt$, or the laggard may achieve her first success (given her full effort strategy in the equilibrium) and the continuation payoff is $\frac{1}{2}(\overline{R} - c/\lambda)$, or partial progress may be disclosed and in that case the leader gets $(\overline{R} - c/\lambda)$. Taking the above integral, we obtain:

$$V_{1,0} = \frac{3\lambda + 2\gamma_r}{2(2\lambda + \gamma_r)}\left(\overline{R} - \frac{c}{\lambda}\right). \tag{38}$$

Plugging in this value into (37), we obtain

$$V_{1,p_r} = p_r V_{1,1} + (1 - p_r) V_{1,0} = p_r \frac{1}{2}\left(\overline{R} - \frac{c}{\lambda}\right) + (1 - p_r)\frac{3\lambda + 2\gamma_r}{2(2\lambda + \gamma_r)}\left(\overline{R} - \frac{c}{\lambda}\right) = \frac{c}{\lambda}, \qquad (39)$$

where the last equality results from substituting $\gamma_r = \lambda(1 - p_r)/p_r$ and $\overline{R} = (2 + p_r)c/\lambda$. Hence, $V_{1,p_r} = c/\lambda$ for $t \geq t_r$. Using (9), we obtain $V_{0,p_r} = 0$ for $t \geq t_r$ and hence (36) is binding. Finally, the exact same argument in the proof of Proposition 4, step 1 can be provided to prove that exerting full effort is optimal for agent i with no success for all $p < p_r$, by showing that if the agent finds it optimal to exert strictly positive effort at any belief p' where $p \leq p' \leq p_r$, then exerting full effort is optimal at belief $p - dp$. Therefore, an agent with no success puts full effort until she succeeds, or the game ends, or partial progress is disclosed. ∎

Proof of Theorem 2: We prove the theorem in multiple steps. To gain insights for why probabilistic encouragement design reduces the expected lead time of the contest, we prove the result for a more general class of contests in which the principal commits to disclose information about any partial progress at constant rate $\lambda(x_0 - p_r)/p_r$ after t_r so that in equilibrium an agent with no success reduces her effort to $x_0 \geq p_r$ for all $t \geq t_r$. Notice that no information disclosure is a special case with $x_0 = p_r$ and $\gamma_t = 0$, and probabilistic encouragement design is a special case with $x_0 = 1$ and $\gamma_r = \lambda(1 - p_r)/p_r$ for all $t \geq t_r$.

    Step 1: *We calculate the expected lead time of the contest under probabilistic encouragement design.*

    Denote by $T_{k,l,t}$ the expected lead time of the contest when one agent has obtained k successes and the other one has obtained l successes from any time t onward. Let us consider the state of the game when both agents have already obtained one success. Then the expected arrival time for the second success is given by:

$$T_{1,1,t} = \int_t^\infty 2\lambda(\tau - t)e^{-2\lambda(\tau - t)}d\tau = \frac{1}{2\lambda}.$$

Here, information disclosure does not affect the outcome since both agents exert full effort until the end. Next, consider the state of the game with a leader (an agent with one success) and a laggard (an agent with no success). Then, the expected lead time of the contest from any time $t \geq t_r$ can be expressed as follows:

$$T_{1,0,t \geq t_r} = \int_t^\infty \left[\lambda(\tau - t) + x_0\lambda\left(\tau - t + \frac{1}{2\lambda}\right) + \frac{\lambda(x_0 - p_r)}{p_r}\left(\tau - t + T_{1,quit,\tau}\right)\right]e^{-\left(\lambda + x_0\lambda + \frac{\lambda(x_0 - p_r)}{p_r}\right)(\tau - t)}d\tau$$

$$= \frac{2 + p_r}{2\lambda(1 + p_r)}, \qquad (40)$$

where $T_{1,quit,\tau}$ is the expected arrival time for the second success once the principal discloses that the leader has made partial progress and the laggard quits, namely,

$$T_{1,quit,t} = \int_t^\infty \lambda(\tau - t)e^{-\lambda(\tau - t)}d\tau = \frac{1}{\lambda}.$$

(40) can be interpreted as follows: conditional on reaching to any instant τ, the leader exerts full effort and if she succeeds the contest ends at $\tau - t$, or the laggard who is putting x_0 effort may achieve her first success and in that case the contest's expected lead time is $\tau - t + 1/2\lambda$, or information may be disclosed by the principal and in that case the laggard quits and the contest ends by the leader at $\tau - t + 1/\lambda$ in expectation. Interestingly, $T_{1,0,t \geq t_r}$ is independent of x_0. Next, for any $t < t_r$, the expected lead time is given by:

$$T_{1,0,t<t_r} = \int_t^{t_r} \left[\lambda(\tau - t) + \lambda\left(\tau - t + \frac{1}{2\lambda}\right)\right]e^{-2\lambda(\tau - t)}d\tau + \left(t_r - t + \frac{2 + p_r}{2\lambda(1 + p_r)}\right)e^{-2\lambda(t_r - t)}$$

$$= \frac{3}{4\lambda} + \frac{1 - p_r}{4\lambda(1 + p_r)}e^{-2\lambda(t_r - t)}, \tag{41}$$

where we use the fact that no information is disclosed by the principal before t_r. Finally, the ex-ante expected lead time of the contest for any $t \geq t_r$ can be expressed as follows:

$$T_{0,0,t \geq t_r} = \int_t^\infty 2x_0\lambda\left(\tau - t + T_{1,0,\tau \geq t_r}\right)e^{-2x_0\lambda(\tau - t)}d\tau$$

$$= \int_t^\infty 2x_0\lambda\left[\tau - t + \frac{2 + p_r}{2\lambda(1 + p_r)}\right]e^{-2x_0\lambda(\tau - t)}d\tau = \frac{1 + p_r + x_0(2 + p_r)}{2x_0\lambda(1 + p_r)}, \tag{42}$$

where we use that an agent with no success exerts effort x_0 after t_r, and for any $t < t_r$ is given by:

$$T_{0,0,t<t_r} = \int_t^{t_r} 2\lambda\left(\tau - t + T_{1,0,\tau<t_r}\right)e^{-2\lambda(\tau - t)}d\tau + \left(t_r - t + T_{0,0,t_r}\right)e^{-2\lambda(t_r - t)}$$

$$= \int_t^{t_r} 2\lambda\left[\tau - t + \frac{3}{4\lambda} + \frac{1 - p_r}{4\lambda(1 + p_r)}e^{-2\lambda(t_r - \tau)}\right]e^{-2\lambda(\tau - t)}d\tau + \left[t_r - t + \frac{1 + p_r + x_0(2 + p_r)}{2x_0\lambda(1 + p_r)}\right]e^{-2\lambda(t_r - t)}$$

$$= \left[\frac{2(1 + p_r) - x_0(1 + 3p_r) + 2x_0\lambda(1 - p_r)(t_r - t)}{4x_0\lambda(1 + p_r)}\right]e^{-2\lambda(t_r - t)} + \frac{5}{4\lambda}, \tag{43}$$

given that both agents exert full effort before t_r.

Under probabilistic encouragement design, we have $x_0 = 1$ after t_r. Also, $p_r = \lambda t_r/(1 + \lambda t_r)$. Plugging in these values into (43), the expected lead time of the contest under probabilistic encouragement design is given by:

$$T_{0,0,0} = \left[\frac{1 - p_r + 2\lambda(1 - p_r)t_r}{4\lambda(1 + p_r)}\right]e^{-2\lambda t_r} + \frac{5}{4\lambda} = \frac{1}{4\lambda}e^{-2\lambda t_r} + \frac{5}{4\lambda}. \tag{44}$$

Step 2: *We prove that probabilistic encouragement design dominates no information disclosure.*

This immediately follows from the previous step. We already show that $T_{1,0,t}$ is independent of x_0. This means the expected lead time of the contest from any time t onward once the first success

is obtained is the same across any design with constant information disclosure of rate $\lambda(x_0 - p_r)/p_r$ that stimulates constant effort x_0 after t_r in the equilibrium. However, according to (43), $T_{0,0,0}$ is decreasing in x_0 and probabilistic encouragement design ensures that $x_0 = 1$ as long as both agents have zero success which results in the minimum expected lead time within this class of contests. Notice that no information disclosure or any disclosure with a rate lower than γ_r fails to encourage full effort and hence is dominated by the probabilistic encouragement design. Finally, we can compute the expected lead time of the contest under no information disclosure by plugging in $x_0 = p_r$ into (43).

Step 3: *We prove that probabilistic encouragement design dominates full information disclosure.*

This step is easy to verify. Note that under full information, the laggard quits upon the arrival of the first success at any time t. Therefore, the expected lead time in this case is given by:

$$T_{0,0,0}^F = \int_0^\infty 2\lambda \left(t + \frac{1}{\lambda}\right) e^{-2\lambda t} dt = \frac{3}{2\lambda},$$

where F stands for full information. However, under probabilistic encouragement design, the principal delays the stopping time of the laggard by t_r periods of time on average if success arrives after time t_r (given that $\gamma_r = 1/t_r$) and by $2t_r - t$ periods of time on average if success arrives at any time $t < t_r$. It is easy to see that $T_{0,0,0}^F < (5 + e^{-2\lambda t_r})/(4\lambda)$.

Step 4: *We prove that probabilistic encouragement design dominates cyclic information disclosure.*

During the first cycle in a design with cyclic information disclosure, if the first success arrives at time $t < t_r$, both agents put full effort during the cycle and the laggard quits at time t_r at the end of the cycle. Therefore, we can write:

$$T_{1,0,t<t_r}^C = \int_t^{t_r} \left[\lambda(\tau - t) + \lambda\left(\tau - t + \frac{1}{2\lambda}\right)\right] e^{-2\lambda(\tau - t)} d\tau + \left(t_r - t + \frac{1}{\lambda}\right) e^{-2\lambda(t_r - t)} = \frac{3}{4\lambda} + \frac{1}{4\lambda} e^{-2\lambda(t_r - t)},$$

where C stands for cyclic information disclosure. Given this, the ex-ante expected lead time of the contest under cyclic information disclosure is given by:

$$T_{0,0,0}^C = \int_0^{t_r} 2\lambda \left(t + \frac{3}{4\lambda} + \frac{1}{4\lambda} e^{-2\lambda(t_r - t)}\right) e^{-2\lambda t} dt + (t_r + T_{0,0,t_r}) e^{-2\lambda t_r}$$

$$\Rightarrow T_{0,0,0}^C = \frac{t_r\ e^{-2\lambda t_r}}{2\left(1 - e^{-2\lambda t_r}\right)} + \frac{5}{4\lambda}, \tag{45}$$

where we use the fact that $T_{0,0,0}^C = T_{0,0,t_r}^C$ as the game resets at time t_r. However, under probabilistic encouragement design, information is disclosed at least t_r periods on average after the success is arrived. It is easy to check that $T_{0,0,0}^C < (5 + e^{-2\lambda t_r})/(4\lambda)$. Thus, probabilistic encouragement design dominates cyclic disclosure.   ■

Proof of Theorem 3:   Part (i). As before, we fix the strategy of agent $-i$ to the proposed one in the theorem and verify that agent i best-responds by playing the same strategy. Condition (31) implies that full effort is optimal for an agent with no success if and only if $V_{1,p} - V_{0,p} \geq c/\lambda$. We use this condition to pin down t_d. Consider the very last instant during the initial silent period of length t_d at which an agent with no success finds it optimal to work. Then, we must have:

$$V_{1,p_d}^d = p_d V_{1,1}^d + (1 - p_d) V_{1,0}^d = \frac{c}{\lambda}, \tag{46}$$

where $p_d = \lambda t_d / (1 + \lambda t_d)$ is the belief of agent i at time t_d anticipating the equilibrium behavior of agent $-i$. We use the superscript d to refer to a design with delay. Given that the principal discloses any progress with a delay t_d, agent i's belief remains constant for $t \geq t_d$. Therefore, at any threshold belief $p = p_d$, we can write:

$$V_{1,0}^d = \int_t^{t+t_d} \left[(\lambda \overline{R} - c) + \lambda \frac{1}{2} \left(\overline{R} - \frac{c}{\lambda} \right) \right] e^{-2\lambda(\tau - t)} d\tau + e^{-2\lambda t_d} \left(\overline{R} - \frac{c}{\lambda} \right),$$

given that during interval $(\tau, \tau + dt)$, the leader puts full effort and earns $(\lambda \overline{R} - c)dt$, or the laggard may achieve her first success (given her full effort strategy in the equilibrium) and the continuation payoff is $\frac{1}{2}(\overline{R} - c/\lambda)$. If neither the leader achieves the second success, nor does the laggard achieve her first success, the principal discloses progress t_d periods of time after its arrival and in that case the laggard quits and the leader gets $(\overline{R} - c/\lambda)$. Taking the above integral, we obtain:

$$V_{1,0}^d = \left(\frac{3}{4} + \frac{1}{4} e^{-2\lambda t_d} \right) \left(\overline{R} - \frac{c}{\lambda} \right). \tag{47}$$

Notice from the above expression that $V_{1,0}$ depends on t_d and not p. Substituting the above value and $V_{1,1}^d = \frac{1}{2}(\overline{R} - c/\lambda)$ into (46) and simplifying the equation, we find that t_d must solve (10). Finally, to prove that exerting full effort is optimal for agent i with no success for all $p < p_d$, we move backward from time t_d associated with belief p_d and prove that if the agent finds it optimal to exert strictly positive effort at any belief p' where $p \leq p' \leq p_d$ (i.e., if $V_{1,p'+dp} - V_{0,p'+dp} \geq c/\lambda$), then we have $V_{1,p} - V_{0,p} \geq c/\lambda$ implying that exerting full effort is optimal at belief $p - dp$. This can be seen from the following:

$$V_{1,p}^d - V_{0,p}^d = -cdt + \lambda \overline{R} dt + (1 - \lambda dt - p\lambda dt) V_{1,p+dp}^d + cdt - \lambda V_{1,p}^d dt - (1 - \lambda dt - p\lambda dt) V_{0,p+dp}^d$$

$$\geq \lambda \overline{R} dt + (1 - \lambda dt - p\lambda dt) \frac{c}{\lambda} - \lambda V_{1,p}^d dt \geq \frac{c}{\lambda}.$$

To show the last inequality, we need to show that

$$V_{1,p}^d \leq \overline{R} - (1 + p) \frac{c}{\lambda}.$$

We prove this in two steps. First, we prove that

$$V_{1,p}^d \leq \frac{1}{1+p}\left(\overline{R} - \frac{c}{\lambda}\right)$$

$$\Leftrightarrow V_{1,p}^d = pV_{1,1}^d + (1-p)V_{1,0}^d = p\frac{1}{2}\left(\overline{R} - \frac{c}{\lambda}\right) + (1-p)\left(\frac{3}{4} + \frac{1}{4}e^{-2\lambda t_d}\right)\left(\overline{R} - \frac{c}{\lambda}\right) \leq \frac{1}{1+p}\left(\overline{R} - \frac{c}{\lambda}\right)$$

$$\Leftrightarrow p\frac{1}{2} + (1-p)\left(\frac{3}{4} + \frac{1}{4}e^{-2\lambda t_d}\right) \leq \frac{1}{1+p} \Leftrightarrow 1 + e^{-2\lambda t_d} \leq \frac{2}{1+p}.$$

To show the last inequality, we use that $p = \lambda t/(1 + \lambda t)$ and $t \leq t_d$. Thus, it is enough to show that

$$1 + e^{-2\lambda t_d} \leq 1 + e^{-2\lambda t} \leq \frac{2(1 + \lambda t)}{1 + 2\lambda t} \Leftrightarrow 1 - (1 + 2\lambda t)e^{-2\lambda t} \geq 0,$$

and the last inequality holds given that the left-hand-side is increasing in t and at $t = 0$, it is binding. In the second step, we prove that

$$\frac{1}{1+p}\left(\overline{R} - \frac{c}{\lambda}\right) \leq \overline{R} - (1+p)\frac{c}{\lambda}$$

which holds if and only if $(2+p)c/\lambda \leq \overline{R}$ which is satisfied for $p \leq p_d$. This verifies the equilibrium.

Part (ii). To prove that $t_d < t_r$, first notice that the left-hand-side of (46) is strictly decreasing in t_d. To see this, note that

$$\frac{\partial V_{1,p_d}^d}{\partial t_d} = \frac{\partial p_d}{\partial t_d}V_{1,1}^d - \frac{\partial p_d}{\partial t_d}V_{1,0}^d + (1-p_d)\frac{\partial V_{1,0}^d}{\partial t_d} < 0,$$

where the above inequality holds since $\partial p_d/\partial t_d > 0$, $V_{1,1}^d < V_{1,0}^d$ and $\partial V_{1,0}^d/\partial t_d < 0$ according to (47). Following this observation, suppose that $t_d = t_r$. Then, we can show that $V_{1,p_r}^d < c/\lambda$ implying that $t_d < t_r$. To see this, recall from (39) that under probabilistic encouragement design we have:

$$V_{1,p_r}^r = p_r V_{1,1}^r + (1 - p_r)V_{1,0}^r = \frac{c}{\lambda},$$

where superscript r refers to the probabilistic encouragement design. Therefore, to prove that $V_{1,p_r}^d < c/\lambda$, it is enough to show that

$$V_{1,0}^d < V_{1,0}^r \Leftrightarrow \left(\frac{3}{4} + \frac{1}{4}e^{-2\lambda t_r}\right)\left(\overline{R} - \frac{c}{\lambda}\right) < \frac{3\lambda + 2\gamma_r}{2(2\lambda + \gamma_r)}\left(\overline{R} - \frac{c}{\lambda}\right).$$

Using the fact that $\gamma_r = 1/t_r$ and further simplifying the above inequality, we need to show:

$$1 - (1 + 2\lambda t_r)e^{-2\lambda t_r} > 0,$$

which always holds for $t_r > 0$. This completes the proof of part (ii).

Part (iii). We first calculate the expected lead time of the contest when the principal discloses partial progress with a delay of length t_d. Denote by $T_{k,l,t}^d$ the expected lead time of the contest

when one agent has obtained k successes and the other one has obtained l successes from any time t onward under a design with delay. Consider any time t when the first success arrives. Then, the expected lead time of the contest from t can be expressed as follows:

$$T_{1,0,t}^d = \int_t^{t+t_d} \left[\lambda(\tau - t) + \lambda\left(\tau - t + \frac{1}{2\lambda}\right) \right] e^{-2\lambda(\tau - t)} d\tau + \left(t_d + \frac{1}{\lambda}\right) e^{-2\lambda t_d} = \frac{3}{4\lambda} + \frac{1}{4\lambda} e^{-2\lambda t_d}, \quad (48)$$

where we use the fact that the laggard quits once the principal discloses progress t_d periods after its arrival. Using the above expression, the ex-ante expected lead time of the contest is:

$$T_{0,0,0}^d = \int_0^\infty 2\lambda\left(t + \frac{3}{4\lambda} + \frac{1}{4\lambda} e^{-2\lambda t_d}\right) e^{-2\lambda t} dt = \frac{1}{4\lambda} e^{-2\lambda t_d} + \frac{5}{4\lambda}. \quad (49)$$

Recall that the expected lead time of the contest under probabilistic encouragement design is given by $\frac{1}{4\lambda} e^{-2\lambda t_r} + \frac{5}{4\lambda}$. Part (iii) of the theorem follows from Part (ii) where we show that $t_d < t_r$. ∎

B. Additional Results

PROPOSITION B1. *When the principal is budget-constrained and commits to full information disclosure, there exists a unique symmetric equilibrium in which both agents exert full effort until the first success arrives. After that, the laggard quits and the leader puts full effort until the end.*

Proof of Proposition B1: Incentive compatibility conditions (13) and (15) show that an agent with one success finds it optimal to put full effort until the end. Given this observation, we have $V_{1,1} = \frac{1}{2}(\overline{R} - c/\lambda) < c/\lambda$. Immediately, from IC condition (17), it can be concluded that the laggard quits. Using this observation, we obtain $V_{1,0} = \overline{R} - c/\lambda$ and $V_{0,0} = \frac{1}{2}(\overline{R} - 2c/\lambda) \geq 0$. Therefore, IC condition (19) is satisfied as $V_{1,0} - V_{0,0} = \frac{1}{2}\overline{R} > c/\lambda$. ∎

PROPOSITION B2. *When the principal is budget-constrained, and commits to probabilistic encouragement design with a flexible reward according to $R_t = (2 + p_t)c/\lambda$ where $p_t = \lambda t/(1 + \lambda t)$ if $t < t_r$, and $R_t = \overline{R}$ for $t \geq t_r$, an agent who has not achieved a success exerts full effort until she obtains her first success, or her opponent obtains her second success, or the principal discloses the opponent's partial progress. An agent who has achieved one success exerts full effort until the end.*

Proof of Proposition B2: In the proof of Proposition 6, we already show that $V_{1,p_r} = c/\lambda$ for $t \geq t_r$. Using (9), we obtain $V_{0,p_r} = 0$ for $t \geq t_r$ and hence (36) is binding. Also for $t < t_r$ and with a flexible reward $R_t = (2 + p_t)c/\lambda$, it is easy to verify from (32) that $V_{1,p} = c/\lambda$. Therefore, (36) is binding implying that an agent with no success finds it optimal to exert full effort. Finally, the reward structure is such that (35) is slack for $t < t_r$. Therefore, an agent with one success puts full effort at all times.

Notice that agents' and principal's combined surplus is the same across any design with the same equilibrium and the same expected lead time for the principal. Therefore, the principal's surplus is maximized when the agents' surplus is minimized. Since this design gives $V_{0,p} = 0$ and $V_{1,p} = c/\lambda$ to the agents before and after obtaining a success, we conclude that our probabilistic encouragement design with flexible reward minimizes the expected reward of the contest. ∎

B.1. Discounting

In this section, we extend Proposition 1 and Theorem 1 to the case where the principal and the agents discount future payoffs at rate $r > 0$. Generalizations of other results follow in a similar fashion and are available upon request from the authors.

B.1.1. First-Best Contract with Observable Effort and Discounting

PROPOSITION B3. *There exists an individually rational "first-best" contract that induces full effort at all times with the minimum required compensation of $\frac{2c}{2\lambda+r} + \frac{4\lambda c(3\lambda+r)}{(2\lambda+r)^3}$ to agents.*

Proof of Proposition B3: Each agent incurs a flow cost of c while working during the contract. Given this, consider the state when both agents have already achieved one success. Then, each agent's expected cost in such a contract from any time t is given by:

$$\int_t^\infty 2\lambda \left(\int_t^\tau ce^{-r(s-t)}ds \right) e^{-2\lambda(\tau-t)}d\tau = \frac{c}{2\lambda+r}.$$

Next, consider the state of the game with a leader and a laggard. Then, each agent's expected cost from any time t can be expressed as follows:

$$\int_t^\infty \left[\lambda \int_t^\tau ce^{-r(s-t)}ds + \lambda \left(\int_t^\tau ce^{-r(s-t)}ds + \frac{c}{2\lambda+r}e^{-r(\tau-t)} \right) \right] e^{-2\lambda(\tau-t)}d\tau = \frac{c(3\lambda+r)}{(2\lambda+r)^2}.$$

Finally, each agent's ex-ante expected cost is given by:

$$\int_0^\infty 2\lambda \left(\int_0^t ce^{-rs}ds + \frac{c(3\lambda+r)}{(2\lambda+r)^2}e^{-rt} \right) e^{-2\lambda t}dt = \frac{c}{2\lambda+r} + \frac{2\lambda c(3\lambda+r)}{(2\lambda+r)^3}.$$

Multiplying the above value by 2 gives us the first-best cost of the principal. ∎

B.1.2. Full Information Disclosure with Flexible Reward and Discounting

THEOREM B1. *Under full information disclosure, a flexible-reward contest with $R_{2,0} = \frac{c(2\lambda+r)}{\lambda^2}$ and $R_{2,1} = \frac{c(3\lambda+r)}{\lambda^2}$ induces both agents to exert full effort at all times with the first-best expected reward of $\frac{2c}{2\lambda+r} + \frac{4\lambda c(3\lambda+r)}{(2\lambda+r)^3}$.*

Proof of Theorem B1: Consider a flexible-reward contest with $R_{2,0} = 2c/\lambda$ and $R_{2,1} = 3c/\lambda$ where the principal commits to disclose any success upon its arrival. Let us fix agent $-i$'s effort $x_{k,l,t}^{-i} = 1$ for all k, l, and t and find conditions under which agent i optimally chooses $x_{k,l,t}^i = 1$ for all k, l, and t. Consider the state of the game where both agents have already achieved one success. The Bellman equation and the corresponding HJB for agent i's problem can be expressed as follows:

$$V_{1,1} = \max_{x_{1,1}} \left\{ x_{1,1}(\lambda R_{2,1} - c)dt + (1 - \lambda x_{1,1}dt - \lambda dt - rdt)V_{1,1} \right\}$$

$$\Rightarrow 0 = \max_{x_{1,1}} \left\{ x_{1,1} \left(\lambda R_{2,1} - c - \lambda V_{1,1} \right) - (\lambda + r)V_{1,1} \right\}, \tag{50}$$

where we use the fact that the winner receives $R_{2,1}$ in this state of the game. From (50), we can derive that $x_{1,1} = 1$ is optimal if and only if $R_{2,1} - V_{1,1} \geq c/\lambda$. Next, consider the state of the game with a leader and a laggard. The Bellman equation and the corresponding HJB for the leader's problem (which we assume to be agent i) can be written as:

$$V_{1,0} = \max_{x_{1,0}} \left\{ x_{1,0}(\lambda R_{2,0} - c)dt + \lambda V_{1,1}dt + (1 - \lambda x_{1,0}dt - \lambda dt - rdt)V_{1,0} \right\}$$

$$\Rightarrow 0 = \max_{x_{1,0}} \left\{ x_{1,0}\left(\lambda R_{2,0} - c - \lambda V_{1,0}\right) + \lambda\left(V_{1,1} - V_{1,0}\right) - rV_{1,0} \right\}, \tag{51}$$

where we use the fact that the winner receives $R_{2,0}$ in this state of the game. From (51), we can derive the IC constraint for the leader which tells us that $x_{1,0} = 1$ is incentive compatible if and only if $R_{2,0} - V_{1,0} \geq c/\lambda$. Similarly, we can express the Bellman equation and the corresponding HJB for the laggard's problem (assuming to be agent i) as follows:

$$V_{0,1} = \max_{x_{0,1}} \left\{ x_{0,1}(\lambda V_{1,1} - c)dt + (1 - \lambda x_{0,1}dt - \lambda dt - rdt)V_{0,1} \right\}$$

$$\Rightarrow 0 = \max_{x_{0,1}} \left\{ x_{0,1}\left(\lambda V_{1,1} - c - \lambda V_{0,1}\right) - (\lambda + r)V_{0,1} \right\}, \tag{52}$$

which implies that exerting full effort for the laggard is optimal if and only if $V_{1,1} - V_{0,1} \geq c/\lambda$. Finally, before the arrival of any success, the continuation value of each agent is given by:

$$V_{0,0} = \max_{x_{0,0}} \left\{ x_{0,0}(\lambda V_{1,0} - c)dt + \lambda V_{0,1}dt + (1 - \lambda x_{0,0}dt - \lambda dt - rdt)V_{0,0} \right\}$$

$$\Rightarrow 0 = \max_{x_{0,0}} \left\{ x_{0,0}\left(\lambda V_{1,0} - c - \lambda V_{0,0}\right) + \lambda\left(V_{0,1} - V_{0,0}\right) - rV_{0,0} \right\}. \tag{53}$$

From (53), exerting $x_{0,0} = 1$ is incentive compatible for each agent if and only if $V_{1,0} - V_{0,0} \geq \dfrac{c}{\lambda}$.

We now verify that the proposed flexible-reward schedule in Theorem B1 satisfies all of the above IC constraints and spends the minimum first-best expected reward. We know that $V_{1,1} = c/\lambda$ is the minimum required continuation payoff to incentivize the laggard to put full effort. From (50), we know that $V_{1,1} = \frac{\lambda R_{2,1} - c}{2\lambda + r}$. Thus, the principal has to specify a reward $R_{2,1} = \frac{c(3\lambda + r)}{\lambda^2}$ in order to satisfy $V_{1,1} = c/\lambda$. Given these values, the IC constraint $R_{2,1} - V_{1,1} \geq c/\lambda$ is satisfied. Also, plugging in the value of $V_{1,1} = c/\lambda$ into (52), we obtain that $V_{0,1} = 0$ and so the IC constraint for the laggard is binding. Next, we know that $V_{1,0} = c/\lambda$ is the minimum required continuation payoff to motivate the agents to exert effort from the beginning. Plugging in this value into (51), $R_{2,0} = \frac{c(2\lambda + r)}{\lambda^2}$ is needed to satisfy the HJB. It follows that the IC constraint for the leader is satisfied as $R_{2,0} - V_{1,0} \geq \frac{c}{\lambda}$. Finally, given $V_{1,0} = c/\lambda$ and $V_{0,1} = 0$, we conclude by (53) that $V_{0,0} = 0$ which shows that the last IC constraint $V_{1,0} - V_{0,0} = \frac{c}{\lambda} - 0 = \frac{c}{\lambda}$ is binding. Therefore, full effort is incentive compatible at all times.

To calculate the expected reward of the contest with flexible reward, note that when both agents have already obtained one success, the expected reward of the contest with discounting is given by

$$\int_t^\infty 2\lambda \left[\frac{c(3\lambda+r)}{\lambda^2} e^{-r(\tau-t)} \right] e^{-2\lambda(\tau-t)} d\tau = \frac{2c(3\lambda+r)}{\lambda(2\lambda+r)}.$$

When there is a leader and a laggard, the expected reward can be computed as follows:

$$\int_t^\infty \lambda \left[\frac{c(2\lambda+r)}{\lambda^2} + \frac{2c(3\lambda+r)}{\lambda(2\lambda+r)} \right] e^{-(2\lambda+r)(\tau-t)} d\tau = \frac{c}{\lambda} + \frac{2c(3\lambda+r)}{(2\lambda+r)^2}.$$

Finally, the ex-ante expected reward of the contest is given by:

$$\int_0^\infty 2\lambda \left[\frac{c}{\lambda} + \frac{2c(3\lambda+r)}{(2\lambda+r)^2} \right] e^{-(2\lambda+r)t} dt = \frac{2c}{2\lambda+r} + \frac{4\lambda c(3\lambda+r)}{(2\lambda+r)^3}.$$

∎