University of Zurich

Department of Economics

Working Paper Series

ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 235

Designing Dynamic Research Contests

Jean-Michel Benkert and Igor Letina

Revised version, August 2019

Designing Dynamic Research Contests[‡]

Jean-Michel Benkert and Igor Letina*

August 14, 2019

Abstract

This paper studies the optimal design of dynamic research contests. We introduce interim transfers, which are paid in every period while the contest is ongoing, to an otherwise standard setting. We show that a contest where: (i) the principal can stop the contest in any period, (ii) a constant interim transfer is paid to agents in each period while the contest is ongoing, and (iii) a final prize is paid once the principal stops the contest, is optimal for the principal and implements the first-best.

Keywords: innovation, dynamic contests, research contests, inducement prizes
JEL: O32, D02, L19

[‡]We are grateful for very helpful comments by Eddie Dekel, Christian Ewerhart, Andreas Hefti, Botond Kőszegi, David Levine, Shuo Liu, Marc Möller, Nick Netzer, Georg Nöldeke, Yuval Salant, Armin Schmutzler, Philipp Strack, Curtis R. Taylor and seminar audiences at the Swiss IO Day 2016, EARIE 2016, the Barcelona GSE Summer Forum 2016, the VfS 2016, the Swiss Theory Day 2016, and the Zurich Workshop on Economics 2016.

*Benkert: Department of Economics and UBS International Center of Economics in Society, University of Zurich, jean-michel.benkert@econ.uzh.ch. Letina: Department of Economics, University of Bern, igor.letina@vwi.unibe.ch. Benkert would like to acknowledge the hospitality of Northwestern University where some of this work was carried out and the Swiss National Science Foundation (Doc.Mobility grant P1ZHP1_161810) as well as the UBS International Center of Economics in Society at the University of Zurich for financial support. Letina would like to acknowledge the hospitality of Stanford University where some of this work was carried out and the Swiss National Science Foundation (Doc.Mobility grant P1ZHP1_155283) for financial support.

1 Introduction

Research contests have a long history as mechanisms for inducing innovation. From navigation and food preservation to aviation, research contests have been used to find solutions to some of society's most pressing problems.[1] Recently, the use of research contests by both the private and public sector has been expanding rapidly. Some examples of problems to which research contests have been applied include vaccine technology, antibiotics overuse, space flight, robotics and AI, as well as environment and energy efficiency.[2] Given that the 2010 America Competes Reauthorization Act authorized US Federal agencies to use prizes and contests, it can be expected that the importance of research contests will only grow in the coming years.

The literature generally classifies research contests into fixed-prize contests and innovation races.[3] To win a fixed-prize contest, an agent needs to have the best innovation at the end of the contest, whereas an agent needs to have a specific innovation as quickly as possible to win an innovation race. The advantage of a race is that it proceeds until an appropriate innovation has been developed (reducing the risk of ending the contest prematurely) and that it minimizes wasteful duplication (since the contest stops as soon as the winning criteria have been met). However, for an innovation race to be feasible, verifiability of outcomes is necessary in order to determine when the race has been won. Since the quality of an innovation is rarely verifiable, innovation races in practice commonly rely on some verifiable proxy of quality.[4] This may result in innovations which are not useful to the sponsor. The 2006 Netflix Prize contest, which aimed to improve the algorithms behind movie recommendations, featured a $1,000,000 prize for the first contestant who would improve upon Netflix's own algorithm by at least 10%. The winning algorithm, however, was never employed because the "the additional accuracy gains [...] did not seem to justify the engineering effort needed to bring them into a production environment." Thus, having to rely on a proxy meant that Netflix did not fully benefit from the innovation that won the final prize.[5]

When implementing a fixed-prize contest, the principal in general does not have to use

[1]1714 Longitude Prize, 1795 Napoleon's Food Preservation Prize and 1919 Orteig Prize, respectively.

[2]2012 EU Vaccine Prize; 2015 Better Use of Antibiotics Prize; 1996 Ansari X Prize, 2006 Northrop Grumman Lunar Lander XCHALLENGE and 2007 Google Lunar X-Prize; 2004 DARPA Grand Challenge, 2007 Urban Challenge, 2014 A.I. presented by TED XPRIZE; 1992 Super-Efficient Refrigerator Program, Progressive Insurance Automotive X PRIZE and 2015 NRG Cosia Carbon XPRIZE.

[3]See the discussion in Taylor (1995).

[4]In the case of the 1996 Ansari X Prize, the objectively verifiable proxy was to have two manned space flights in two weeks using the same spacecraft, while the larger objective of the organizer was to develop private space travel. See http://ansari.xprize.org.

[5]Of course, it is entirely possible that Netflix implemented parts of the winning algorithm and also made use of the ideas from algorithms that did not win the contest. See *Netflix Recommendations: Beyond the 5 stars (Part 1)* by (at the time) Netflix's Engineering Director Xavier Amatriain and Lead Researcher Justin Basilico. http://tiny.uzh.ch/F0.

a proxy and can instead evaluate the quality of the innovation directly. However, other problems arise, since a fixed-prize contest requires a finite deadline, at which submissions will be evaluated. If the agents are not given enough time, they may fail to produce a good enough innovation.[6] If the deadline is very late, the implementation of the innovation will be delayed and, in addition, there is a risk of wasteful duplication, as agents invest resources while an adequate innovation might already have been developed.

This paper deals with the optimal design of dynamic research contests. We identify a novel design lever that the contest designer can use in order to increase the efficiency of the contest: interim transfers to the agents while the contest is ongoing. The optimal contest takes a remarkably simple form. The principal decides in each period whether to stop the contest or not. If the contest is not stopped, the principal has to pay out an interim transfer, the value of which was announced before the contest started, and which is divided equally among the agents. If the contest is stopped, the principal has to pay out the final prize, the value of which was also announced in advance. We call this type of contest an *interim-transfer contest*. An interim-transfer contest offers a solution to the trade-off between innovation races and fixed-prize contests. It induces an equilibrium where all agents conduct research until they achieve an innovation quality (which does not have to be verifiable) set out by the principal. The agent that discovers such an innovation immediately reveals it to the principal, who in turn immediately stops the contest. An interim-transfer contest thus combines the best of the two formats: it inherits the race-like structure, thereby eliminating wasteful duplication and the risk of a premature ending, without requiring a verifiable proxy for quality.

Our setting closely follows the seminal work by Taylor (1995). There is a principal who would like to procure an innovation and multiple agents who can choose to conduct research in each period. Research is costly and yields an innovation whose value is a random draw from some distribution. The research decisions and the values of innovations are private information to the agents. If an agent reveals the innovation to the principal, the principal can costlessly and accurately determine its value. However, the value is not verifiable by outside parties. In particular, a contract which conditions on the innovation value is not enforceable by the courts.[7] In order to incentivize the agents to conduct research, the principal announces a contest. Taylor (1995) studies what we call a fixed-prize contest, in which the principal commits to paying a final prize at the end of an exogenously given deadline of T periods.[8] We introduce the interim-transfer contest,

[6] The objective of the 2004 DARPA Grand Challenge was to "accelerate the development of autonomous vehicle technologies that can be applied to military requirements" but none of the competitors managed to fulfill the requirements of the tournament. Eventually, the requirements were matched in the 2005 DARPA Grand Challenge, suggesting that more time was needed to be successful. See the official website on http://archive.darpa.mil/grandchallenge04/.

[7] These assumptions are standard in the literature. See Taylor (1995) or Che and Gale (2003).

[8] The principal only needs to commit to pay out the final prize to one of the agents. Once that

where the principal commits to both the final prize (which is paid out once the contest ends), and to the interim transfer (which is paid out in each period as long as the contest continues). How the interim transfer is allocated among the agents is not important for our result. Therefore, we focus on the simplest allocation rule, where the transfer is divided equally among all agents.[9] We endogenize the deadline T, which becomes a design choice of the principal and allow for an infinite horizon. Furthermore, the principal can end the interim-transfer contest before the deadline is reached, in which case the final prize is paid out. Like in Taylor (1995), the principal determines the number of agents invited to the contest and sets the entry fee or entry subsidy which the agents need to pay in order to participate in the contest. Agents join the contest if their expected payoff is at least zero.

Our first main result shows that we can implement any innovation value as the threshold of a global stopping equilibrium with any number of agents using an appropriate interim-transfer contest. In a global stopping equilibrium, all agents conduct research in every period until some agent discovers an innovation with a value above the threshold, at which point the research effort is stopped by all agents. This result is in sharp contrast to the finding in Taylor (1995), who shows that a fixed-prize contest implements only individual stopping equilibria. With an individual stopping equilibrium, each agent conducts research in every period until she herself discovers an innovation with a value above the threshold, irrespective of what the other agents find. Our result is remarkable for three reasons. First, it eliminates the wasteful duplication of research efforts which occurs in individual stopping equilibria. Second, it eliminates the possibility that the contest will end before an appropriate innovation has been found. Third, the innovation is made available to the principal as soon as it is discovered. In addition, if a natural restriction is placed on the principal's equilibrium strategies, then a global stopping equilibrium is the unique outcome implemented by an appropriately designed interim-transfer contest.

There is a clear economic interpretation underlying our implementation result. On the one hand, the final prize takes care of the agents' incentives. Namely, it induces them to conduct research in every period in order to obtain an innovation of value above the threshold and win the final prize. The presence of the interim transfer complicates matters, because an agent with a value above the threshold may wish to delay the submission of the innovation in order to benefit from interim transfers. Such a delay is risky, though, as another agent may also obtain an innovation above the threshold and win the contest. A sufficiently high final prize ensures that delaying is not profitable and induces truthful reporting of research outcomes in addition to research effort in all periods.[10] On the other

commitment is made, the promise to award the prize to the agent with the best innovation is credible.

[9]More complex allocation rule could, for example, allocate the entire transfer to the agent with the best innovation so far. We provide further discussion of this possibility in the conclusion.

[10]The danger of delaying submission is relevant in the real-world too. In the Netflix Prize contest

hand, the interim transfer takes care of the principal's incentives. In a fixed-prize contest, the principal has no incentive to stop the contest once a value above the threshold has been submitted, because she does not bear the marginal cost of research, but stands to gain by obtaining an even higher value in the next period. The interim transfer is set in such a way that it is weakly greater than the marginal benefit of continuing the contest for one more period when the threshold value has been obtained. Thus, the marginal cost of continuing the contest for one more period outweighs the marginal benefit. Consequently, the principal wants to stop the contest if a value at or above the threshold is obtained, thereby giving rise to a global stopping equilibrium.

While the ability to implement global stopping equilibria using interim-transfer contests is of interest in itself, it does not answer the question whether the principal would actually want to do so. Indeed, one can show that for a given threshold value and number of agents, a global stopping equilibrium has lower expected research costs than an individual stopping equilibrium, but that the latter has a higher expected value of innovation than the former. Hence, it is unclear a priori which contest maximizes the principal's utility. Moreover, it is not clear if some other contest or some other mechanism performs better than either an interim-transfer or a fixed-prize contest. Our second main result addresses this issue by showing that the principal can implement the first-best outcome with an interim-transfer contest. Then, the principal can extract the entire surplus with appropriate entry fees. Thus, interim-transfer contests constitute optimal mechanisms more generally.

Our optimality result relies crucially on the assumption that the principal can charge entry fees. The ability of the principal to charge entry fees has been previously used in the literature and it has also been used in practice.[11] However, entry fees in the optimal interim-transfer contest could potentially be very large. If the agents who are supposed to participate in the contest have budget constraints, then implementing an optimal interim-transfer contest may not be possible. This implies that the scope of applications of our optimal contest might be limited and would need to be verified on a case by case basis. Finally, even if the principal could not charge entry fees, our results would still imply that an interim-transfer contest is a welfare-maximizing mechanism.

In the main model, we assume that the principal can choose an infinite deadline for the contest. In an extension, we consider the case where there is an exogenous finite deadline beyond which the contest cannot run or the innovation becomes worthless. With

the second-placed team had the same performance measure as the winners. They lost because they submitted their bid 20 minutes after the eventual winners. See "A \$1 Million Research Bargain for Netflix, and Maybe a Model for Others" in *The New York Times*. https://nyti.ms/2InaHqN.

[11]Entry fees are assumed in Taylor (1995) and are central in papers on auctioning entry into contests, such as Fullerton and McAfee (1999) and Giebe (2014). As an example of their use in practice, Taylor (1995) mentions the \$200,000 entry fee into the FCC's contest aimed to develop the HD TV standard.

a finite deadline, the first-best outcome does not have a structure as simple as with an infinite deadline. Gal, Landsberger, and Levykson (1981) and Morgan (1983) have shown that while a global stopping equilibrium is still optimal, the first-best search intensity generally depends both on the number of periods until the deadline and on the current highest innovation quality. For tractability, we impose assumptions on the research process which guarantee that the first-best search intensity, as long as the contest is ongoing, depends only on time and not on the highest innovation quality. In particular, we assume breakthrough innovation structure, which implies that the principal only cares about making a breakthrough, and that all innovations constituting a breakthrough are worth approximately the same. Under the assumption of a breakthrough innovations structure, the first-best search rule is a global stopping equilibrium where the number of participants in the contest increases as the deadline draws near. Importantly, the optimal number of contestants can be fixed ex ante. We generalize the notion of the interim-transfer contest to allow for an increasing number of contestants over time and show that this generalized interim transfer contest is an optimal mechanism and implements the first best.

The paper is structured as follows. The model is introduced in Section 2. We study the infinite horizon case in Section 3, while finite horizon is considered in Section 4. Discussion of the literature is in Section 5. We conclude in Section 6. All proofs are in the appendix.

2 Model

2.1 Setting

There is a risk-neutral principal who wants to procure an innovation, and a set of identical risk-neutral agents $N = \mathbb{N}$ who can potentially discover the innovation by conducting research. If the principal obtains the innovation in any period $s \in \{1, \dots, T\}$ with $T \leq \infty$, her payoff is $\delta^{s-1}\theta - \sum_{t=1}^{T} \delta^{t-1} w_t$, where θ is the value of the innovation and $w_t = \sum_i w_{ti}$ is the sum of transfers made to all agents in period t. Agent i's payoff is $\sum_{t=1}^{T} \delta^{t-1}(w_{ti} - c_{ti})$, where w_{ti} is the transfers received and c_{ti} is the cost incurred through research activities in period t. We assume that the innovations are of no intrinsic value to the agents and allow for any discount factor $\delta \in (0, 1)$.

An agent can conduct research in any period t at per-period cost $C > 0$. In each period in which the agent performs research, he obtains an innovation of value $\theta \in \Theta$. The innovation value obtained is an independent draw from some distribution F with full and finite support $\Theta = \{\theta^1, \theta^2, \dots, \theta^K\}$. Without loss, assume that $\theta^{k+1} > \theta^k$ and let $\theta^1 = 0$. We interpret the outcome $\theta^1 = 0$ as a failure to develop an innovation which is more valuable than the current technological level. Agents can repeatedly conduct research and have perfect recall, that is, they can access all their own previously obtained innovations at any point in time. Initially, every agent is endowed with a worthless innovation and in each period in which an agent does not conduct research he receives a worthless innovation (i.e., a shirking agent does not advance beyond the current technological level).

The agents' research activity (whether or not they conduct research in any given period) and research outcomes (the value of an innovation obtained in any given period) are private information. If an agent submits an innovation to the principal, the principal can determine the value of the innovation at no cost. However, the value of an innovation is not verifiable by a court. Thus, contracts conditioning on the value of an innovation are not credible.[12] The agents are not budget constrained. This has two implications. The agents can always afford to conduct research and the principal can charge positive entry fees. The outside option of the agents is normalized to zero.

2.2 Interim-Transfer Contests

Taylor (1995) considers what we will refer to as a *fixed-prize contest* (FPC). An FPC is a tuple $\Gamma = \langle E, p, n, T\rangle$, where $E \in \mathbb{R}$ is an entry fee, $n \in \mathbb{N}$ is the (invited) number of participating agents, $T \in \mathbb{N}$ is the finite contest deadline and $p \in \mathbb{R}$ is the final prize, which is awarded to an agent at the end of T periods. In contrast, we introduce an *interim-transfer contest* (ITC) as a tuple $\Gamma = \langle E, p, m, n, T\rangle$. An ITC differs from an FPC in the three ways. First, there could also be an "infinite" deadline $T = \infty$, which means that there is no deadline and the contest can run for as long as the principal wants it to. Second, there is an interim transfer $m \in \mathbb{R}$ which the principal has to pay in every period until the contest stops. Third, the principal can stop the contest in every period, which cannot be done in an FPC. Figure 1 illustrates the timing of an ITC.[13]

In order for an ITC to be credible, courts need to be able to verify (i) that the interim transfer was paid in every period but the last, and (ii) that the final prize was awarded, which can happen no later than at the deadline T. Thus, courts only need to be able to verify time-contingent payments. We view this assumption as uncontroversial, as such contracts are common. For example, utility bills come with payment deadlines and penalties for late payment.

The contest $\Gamma = \langle E, p, m, n, T\rangle$ induces an extensive form game of incomplete information. The set of players is $I = \{0, 1, \ldots, n\}$, where player 0 is the principal and players 1 to n are the agents. The set H is the set of histories and the actions available after any non-terminal history h is denoted $A(h) = \{a : (h, a) \in H\}$. The set of initial histories is

[12] Non-observability and non-verifiability is a typical feature of research activity. As Taylor (1995, p. 873) notes "research inputs are notoriously difficult to monitor" and "courts seldom possess the ability or expertise necessary to evaluate technical research projects".

[13] Of course, one can imagine different contest formats, for example where the allocation of transfers among agents is dependent on interim performance or with time-contingent interim transfers m_t. As it turns out, the optimum can be achieved with a simple ITC.

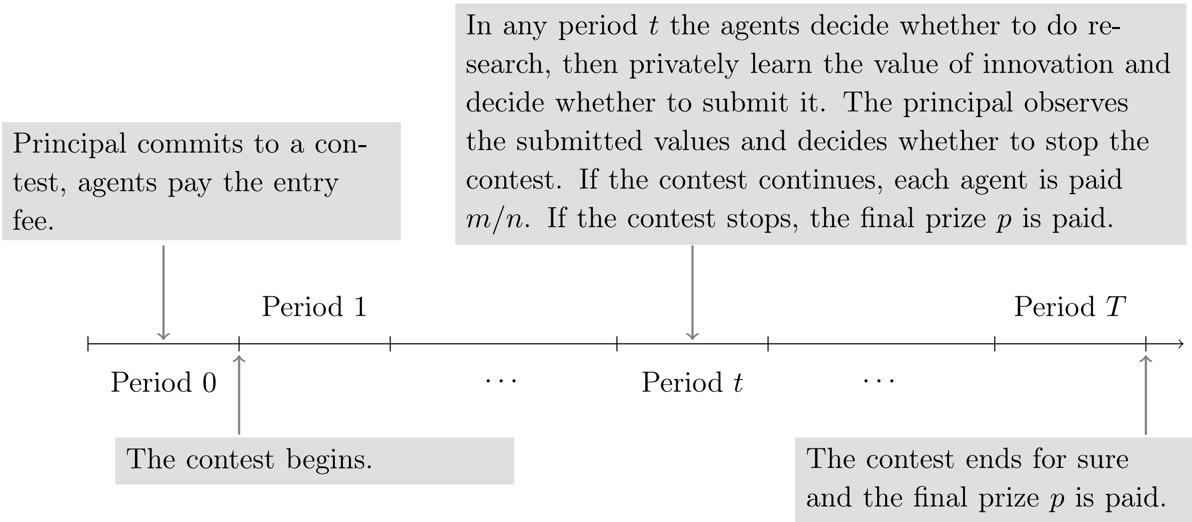

Figure 1: Timeline of an ITC.

the set of the states of the world. The true initial history is $\theta \in \Theta^{NT}$, where each element $\theta_{it} \in \Theta$ (where $i \in \{1, \dots, n\}$ and $t \in \{1, \dots, T\}$) is drawn i.i.d. from the probability distribution F. An agent i who conducts research in period t receives a value equal to θ_{it}. For the agents the payoffs are determined by the research costs they have incurred, the entry fee they pay if they enter the contest, and the transfers they receive. The principal's payoff is determined by the value of the innovation she gets, the entry fees of the participants, and the transfers she makes to the agents. In what follows, we will use the terms doing research and investing (in research) interchangeably. The set of players, their payoff functions, the research technology and the contest structure are common knowledge. The timing of the game is as follows.

Period 0**:**

- The principal announces the contest $\Gamma = \langle E, p, m, n, T \rangle$.

- Agents decide whether to enter or not. All agents who enter the contest pay the entry fee E.

Period $t < T$**:**

- Stage 1: Each agent simultaneously decides whether to perform research at cost C. Agents do not observe the actions taken by their competitors.

- Stage 2: Each agent i who conducted research receives value equal to θ_{it}. All other agents receive value 0.

- Stage 3: Having privately observed the value of their innovation, agents simultaneously decide whether to privately submit their innovation.

- Stage 4: The principal observes the set of submissions (if any) and decides whether to stop the contest. If the contest is stopped, the principal obtains the highest submitted innovation and the agent with the highest submitted quality receives the prize p. If the highest quality is submitted by multiple agents, the winner is chosen randomly among them. If the contest continues the principal pays the interim transfer m, so that each agent receives m/n.

Period T (if $T < \infty$):

- Stages 1-3: As above.

- Stage 4: The contest stops, the principal obtains the highest submitted innovation and the agent with the highest submitted quality receives the prize p. If the highest quality is submitted by multiple agents, the winner is chosen randomly among them. If there have been no submissions, the winner is randomly chosen.

Throughout the paper, we focus on the following equilibrium candidate. Denote with σ_i the strategy function and with μ_i the beliefs of player i. Denote with θ_i^{max} the highest value available to player i. For agents, this is the highest value they have so far discovered. For the principal, this is the highest value currently submitted. Denote with $\theta^g \in \Theta$ the threshold of the global stopping equilibrium. For agents, the equilibrium strategy is:

- if $A(h) = \{\text{Invest, Not Invest}\} = \{I, NI\}$ then

$$\sigma_i(\theta_i^{max}) = \begin{cases} I & \text{if } \theta_i^{max} < \theta^K \\ NI & \text{else} \end{cases};$$

- if $A(h) = \{\text{Submit, Not Submit}\} = \{S, NS\}$ then

$$\sigma_i(\theta_i^{max}) = \begin{cases} S & \text{if } \theta_i^{max} \geq \theta^g \\ NS & \text{else} \end{cases}.$$

The equilibrium strategy of the principal is:

- if $A(h) = \{\text{Continue, Not Continue}\} = \{Cont, NCont\}$ then

$$\sigma_0(\theta_0^{max}) = \begin{cases} NCont & \text{if } \theta_0^{max} \geq \theta^g \\ Cont & \text{else} \end{cases}.$$

The equilibrium beliefs of agent i over the state θ_{jt} are given either by Bayesian updating or by prior beliefs, i.e., the distribution F. For own elements ($i = j$), the agent learns the exact state if he invests. If he does not invest, or if the chance to invest has not occurred yet, he holds initial beliefs given by distribution F. For the elements of other agents ($i \neq j$), if no deviation has been observed, the agent concludes that everybody has invested up to that point and that nobody has a value higher than θ^g. This implies that each individual θ_{jt} is drawn from the distribution F truncated at θ^g. If the agent observes that the principal deviated, he learns nothing about the realization of the state in that period, and holds the initial beliefs. The principal believes that any agent who has not submitted an innovation has conducted research in every period, yet the draws were always below θ^g (i.e., the distribution is from the truncated F). For an agent who has submitted an innovation, the principal believes that research has been conducted in every period and that the submission is the current highest value. For all states which have not occurred yet, the agents and the principal hold initial beliefs.

3 Infinite Horizon

3.1 Implementation of Global Stopping Equilibria

Stopping behavior provides an easy and intuitive way to describe the agents' equilibrium play in the contest in terms of their research activity. Under an *individual stopping equilibrium*, every agent does research in every period of the contest until an individual threshold value of innovation is reached, irrespective of the innovations discovered by other agents. Such an individual stopping equilibrium consequently entails a risk of duplication of research effort across agents. In contrast, such wasteful duplication of research efforts is avoided under a *global stopping equilibrium*, where every agent does research in every period of the contest until a single agent reaches some global threshold at which point all agents stop doing research.

An FPC uniquely implements individual stopping equilibria, as the chance of winning the final prize provides the agents with an incentive to conduct research until a threshold is reached (Taylor, 1995). However, an FPC cannot implement a global stopping equilibrium because the principal cannot credibly commit to stop the contest after the threshold value has been achieved. The reason is that she does not bear the marginal cost of continued research, while she stands to benefit from any marginal increase in the value of innovation. Our first result shows that an appropriately designed ITC can implement any value of innovation as the threshold of a global stopping equilibrium, thereby eliminating the risk of wasteful duplication of research effort present in an FPC.

Proposition 1 *Any $\theta^g \in \Theta$ can be implemented as the threshold of a global stopping equilibrium with any $n \geq 2$ and for any $T \leq \infty$ by using an ITC with final prize $p \geq \bar{p}$*

and the interim transfer $m = p(1-\delta) + \delta\Delta(\theta^g, n) - \theta^g + \varepsilon$, *where*

$$\Delta(\theta^g, n) = F^n(\theta^g)\theta^g + \sum_{j=g+1}^{K} \left(F^n(\theta^j) - F^n(\theta^{j-1})\right)\theta^j \tag{1}$$

$$\bar{p} = \max\left\{\frac{\delta n\Delta(\theta^g, n) - \theta^g + \varepsilon}{\delta(n-1)(1-F(\theta^{K-1}))^{n-1}}, \frac{2nC}{(1-F(\theta^{K-1}))^n}\right\} \tag{2}$$

$$\varepsilon \in [0, (1-\delta F^n(\theta^{g-1}))(\theta^g - \theta^{g-1})). \tag{3}$$

The proposition shows that for any $\theta^g \in \Theta$, there exists an ITC with low enough E so that in a perfect Bayesian equilibrium n agents enter the contest, conduct research and submit any innovation of value at least θ^g to the principal. The principal stops the contest as soon as an innovation of value at least θ^g has been submitted and so stops any further research effort. The agent who submitted the highest value innovation wins the final prize.[14] Moreover, as long as the deadline is not reached and the principal does not end the contest, the principal pays m/n to each agent.[15]

In order to implement a global stopping equilibrium, the agents need to conduct research until the threshold is reached and then immediately report their innovation to the principal. The principal needs to stop the contest once this happens. The final prize and the interim transfer are the instruments which ensure this behavior, respectively.

The interim transfer m when $\varepsilon = 0$ corresponds exactly to the principal's marginal benefit of continuing the contest one more period when an innovation of value θ^g has been submitted. Thus, through the interim transfer, the principal incurs a constant marginal cost of one more round of research by the agents. Since the marginal benefit of research to the principal is decreasing in θ, the principal strictly prefers to continue the contest whenever the highest innovation value is below θ^g, strictly prefers to stop it whenever it is above, and if $\varepsilon = 0$ she is indifferent at the threshold value. For positive values of ε, the preference to stop is strict also for the threshold value. As a consequence, the principal will credibly stop the contest if at least the threshold value θ^g was reached.

Conversely, the final prize gives the incentivizes the agents to perform research in every period and to report their research outcomes truthfully. Intuitively, as in the FPC, each agent pursues an individual stopping threshold which is determined by the expected probability of winning the final prize. Increasing the final prize increases the individual stopping threshold. If the individual stopping threshold is above the global stopping threshold, the agents will conduct research as long as the contest is ongoing. Similarly,

[14] In case of multiple innovations of equal highest value being reported simultaneously, the winner is randomly chosen.

[15] Instead of paying m/n to each agent, the principal could allocate m in different ways. For example, entire m could be paid to the agent with the best intermediate value. We discuss this possibility further in the conclusion.

the final prize induces the agents to truthfully report their research outcomes. By not reporting an innovation above the threshold, an agent could attempt to obtain the interim transfer in the current period in addition to the final prize in the next period. However, not reporting exposes the agent to the risk that another agent will win in the current period and end the contest. As long as the size of the final prize is sufficiently large relative to the interim transfer, the agent will report truthfully. Thus, incentives of the agents can be satisfied by making the final prize large enough. Since the interim transfer is a function of the final prize, the incentives of both the principal and the agents can be satisfied simultaneously, so that the equilibrium exhibits global stopping. The final prize p and the interim transfer m are set high enough so that the agents' participation constraints are satisfied and n agents want to enter the contest. The condition $p \geq \bar{p}$ is a sufficient condition needed to guarantee that any $\theta^g \in \Theta$ can be implemented as the threshold of a global stopping equilibrium. In particular instances, a global stopping equilibrium could be implemented with final prizes which are lower than $\bar{p}$.

The above intuition also serves as a sketch of the proof. The first step is to show that the principal does not want to deviate from the equilibrium strategy given the interim transfer, while the second step is to show that for a sufficiently high final prize, the agents do not want to deviate either.

An ITC does not have a unique equilibrium. For example, an agent submitting an innovation below the threshold is also an equilibrium of this contest. While a multiplicity of equilibria can be an issue when the principal wants to implement some specific outcome, this alternative equilibrium profile does not change the outcome of the game in any way. The principal simply ignores any submissions below the threshold and the game results in exactly the same payoffs in all states of the world. Our next result shows that this observation holds for any equilibrium in which the principal plays a *threshold strategy*. We say that a principal's strategy σ_0 is a threshold strategy if there is a $\theta^s \in \Theta$ such that

$$\sigma_0(\theta_0^{max}) = \begin{cases} NCont & \text{if } \theta_0^{max} \geq \theta^s \\ Cont & \text{else} \end{cases}.$$

Given this restriction on the equilibrium strategy of the principal (but not on her deviation strategies), an ITC exists which uniquely implements any $\theta^g > 0$ as the threshold of a global stopping equilibrium.

Proposition 2 *Fix any $\theta^g \in \Theta$ such that $\theta^g > \theta^0$ and consider any ITC satisfying the conditions of Proposition 1 with the final prize $p > \bar{\bar{p}}$ and $\varepsilon > 0$, where*

$$\bar{\bar{p}} = \max\left\{\frac{n^2 C + (\delta\Delta(\theta^g, n) - \theta^g + \varepsilon)\left(1 - F(\theta^{K-1})\right)}{(1-\delta)(n-1)\left(1 - F(\theta^{K-1})\right)}, \bar{p}\right\}.$$

Then, any perfect Bayesian equilibrium in which the principal follows a threshold strategy implements a global stopping with θ^g.

In order to guarantee uniqueness of outcomes, we further restrict the ITC from Proposition 1 in two ways. First, we require that $\varepsilon > 0$, which provides strict incentives to the principal to stop the contest at the threshold value. Second, the final prize needs to be higher. The proof shows that if the principal follows a threshold strategy, then in any PBE the agents will do research and submit all innovations which are at or above the threshold. The principal, in turn, will use θ^g as the threshold of her strategy. Thus, any multiplicity of equilibria is a result of either different submission strategies for innovations below the threshold, or different strategies off equilibrium, neither of which is payoff relevant for the principal or the agents.

3.2 Optimal Contests

Our first main result above shows that it is possible to implement global stopping equilibria without relying on verifiability of research outcomes. While this is of interest in itself, the fact that the principal can implement a global stopping equilibrium does not imply that she actually wants to do so. One may suspect that a global stopping equilibrium is preferable to an individual stopping equilibrium, but this need not be true in general. Indeed, comparing an FPC and an ITC with the same threshold value and number of contestants, one can show that the individual stopping equilibrium yields a higher expected value of innovation, but has higher expected costs than the global stopping equilibrium. It is thus not obvious whether the ITC or the FPC performs better. Moreover, it is not clear if there exists some other contest, or indeed some other mechanism, which performs better than either the ITC or the FPC. To shed light on this, we turn to the question of optimality in this section.

The first step is to characterize the first best, that is the outcome without the principal-agent problem. If the principal could directly control the agents, then our framework would correspond to classic search problems (Gal et al., 1981; Morgan, 1983). In these models, the principal in every period decides whether to continue the search or not. If the search is continued, the principal decides with which intensity to search and bears the corresponding search costs. If the search stops, the principal obtains, in that period, utility equal to the highest value found. In each period that the search continues, a value is drawn as a function of the search intensity.

As Benkert, Letina, and Nöldeke (2018) argue, this model can be embedded into a Markov decision process framework.[16] This is done in the following way. Let $s \in \Theta^z$ be a state variable with $\Theta^z = \Theta \cup \{z\}$, where $s \in \Theta$ implies that the highest value found

[16]For an overview of Markov decision process framework, see DeGroot (1970, Chapter 14).

so far is s and $s = z$ implies that the search has been terminated. Let $n \in N^z$ be the action variable, with $N^z = N \cup z$, where $n \in N$ implies continued search with n agents and $n = z$ implies that the search is terminated. The reward function is given by

$$r(s,n) = \begin{cases} -\delta n C, & \text{if } n \neq z \text{ and } s \neq z \\ s, & \text{if } n = z \text{ and } s \neq z \\ 0, & \text{if } s = z. \end{cases}$$

Finally, the transition function $\phi(s'|s,n)$ defines the probability that the state transitions from s to s' given the action n. Once terminated, the search cannot leave the terminal state, so that $\phi(z|z,n) = \phi(z|s,z) = 1$ for all $s \in \Theta^z$ and all $n \in N^z$. Furthermore, the search never enters the terminal state without being stopped, that is $\phi(z|s,n) = 0$ for all $s, n \neq z$. Since there is perfect recall, the state never decreases, that is $\phi(s'|s,n) = 0$ for all $s' < s$ and $s, n \neq z$. For all remaining cases, $\phi(s'|s,n)$ is derived from the probability that s' is drawn from F given n draws. Then the problem of the principal is to choose the horizon $T \leq \infty$ and a sequence of functions $n_0, n_1, \ldots, n_{T-1}$ where $n_t : \Theta^z \to N^z$ which will maximize the expected value of her discounted sum of rewards. Denoting with S_t the random state variable in period t, this can be expressed as

$$\max_{T, n_0, \ldots, n_{T-1}} \frac{r(0, n_0)}{\delta} + \mathbb{E}\left[\sum_{t=1}^{T-1} \delta^{t-1} r(S_t, n_t) + \delta^{T-1} r(S_T, z)\right].$$

While $n_t(\cdot)$ could potentially depend on time and state variables in complex ways, Benkert et al. (2018) show that the solution to the principal's maximization problem takes a very simple form. The optimal horizon T is infinite, and there exist $n^{FB} \in N$ and $\theta^{FB} \in \Theta$ such that

$$n_t^*(\theta) = \begin{cases} n^{FB}, & \text{if } \theta < \theta^{FB} \\ z, & \text{if } \theta \geq \theta^{FB}, \end{cases}$$

for all t. That is, the principal optimally searches with a constant sample size n^{FB} across time until some threshold θ^{FB} has been reached.[17] While we do not have explicit expressions for n^{FB} and θ^{FB}, it is simple to find both for any given parameters of our model.[18] Put differently, a global stopping equilibrium with constant search intensity is an optimal search strategy. This leads to our next result.

[17] This result is related to the optimality of Gittins index policies in multi-armed bandit problems. See Bergemann and Välimäki (2001).

[18] Since the per-period search costs are equal to $-nC$, increasing the sample size leads to higher costs. The benefit is that an innovation above the threshold will be discovered sooner, which is desirable due to discounting. The optimal sample size n^{FB} balances the speed of discovery and the cost of research.

Proposition 3 *If $n^{FB} \geq 2$ the first-best outcome can be implemented using an ITC.*

The result follows in three steps. First, by Proposition 1 in Benkert et al. (2018), the first best is equivalent to a global stopping equilibrium with (i) θ^{FB} threshold, (ii) infinite horizon, and (iii) n^{FB} agents. Second, by Proposition 1 such an outcome can be implemented by an interim-transfer contest. Third, choosing the entry fee E such that the agents' participation constraint is binding allows the principal to fully extract the first-best surplus. Thus, the ITC not only outperforms the FPC, rather, it is the optimal mechanism more generally.[19]

As noted in the previous section, we need $n^{FB} \geq 2$ in order to induce any effort in any contest. However, a slight twist to the ITC allows us to implement the first-best outcome when $n^{FB} = 1$, too. More precisely, to implement a global stopping equilibrium with only one agent doing research, the principal announces an ITC between a "real agent" and a "fictitious agent". The interim transfers are always paid to the fictitious agent, while only the real agent can receive the final prize, which occurs only when he submits an innovation value above the threshold. This way the interim transfer still ensures that the principal will adhere to the global stopping equilibrium, the agent will exert effort in every period in order to obtain the final prize by reaching at least the threshold value, and the fictitious agent has no incentive to exert any research effort.

Another advantage of holding an ITC over an FPC is that the equilibrium of an ITC is more robust to information disclosure by the agents. Suppose that each agent could credibly disclose the quality of its innovation to its rivals, for example by holding a public demonstration of his current technology. In an FPC, an agent with a high enough innovation quality might have an incentive to disclose his innovation in order to discourage the competitors from investing. This, in turn, could lower the expected quality that the principal obtains from an FPC.[20] The equilibrium of the ITC would not be affected by the ability of agents to disclose their innovation quality. If an agent disclosed a quality below the winning threshold, this would not affect the payoffs of his competitors in any way, since winning requires an innovation above the threshold. Hence, no agent benefits from disclosing such an innovation to his competitors. Any innovation above the threshold immediately causes the principal to stop the contest, so that again the equilibrium is not affected by the possibility of information disclosure.[21]

Taylor (1995) notes that the first best could be achieved if the principal, instead of holding one multi-period contest, held a series of one-period contests. However, if

[19]Of course, this requires that the principal can charge entry fees and that the agents are not budget constrained. As mentioned in the introduction, entry fees have been used both in the literature and in practice. Nevertheless, the requirement may limit the scope of applications.

[20]See Rieck (2010). Disclosure of intermediate innovations is also studied in Gordon (2011) and Gill (2008).

[21]We would like to thank an anonymous referee for pointing this out.

inspecting the agents' submissions is costly, running a sequence of one-period contests and inspecting submissions after every period may be prohibitively costly. This points to another advantage of the ITC. Namely, the principal only has to inspect submissions once and only from those agents who have developed an innovation of high enough quality. This argument relies on the implicit assumption that the agents can evaluate their innovations at no cost. If it was costly for the agents to evaluate the innovations, they could submit without evaluating, thereby shifting the evaluation costs to the principal. The principal can address this concern in several ways. First, the principal can change how the interim transfer m is allocated among agents to provide incentives to agents to actually evaluate their innovations. This can be done by splitting the interim transfer m (in the case the contest continues) only among those agents who have not submitted.[22] Second, if the principal can observe whether or not the agents have evaluated their submitted innovation, she could refuse to consider submissions which have not been evaluated. For example, while evaluating the value of pharmaceutical innovations is certainly costly, the principal can easily observe whether or not the new drug went through the clinical trials. Finally, even if the principal had to bear the cost of evaluating all innovations in each period, she would not be worse off than in a series of one-period contests.[23]

4 Finite Horizon

As noted above, the first best takes a particularly simple form in the case of an infinite horizon. With a finite horizon, the first best is more complicated. It is characterized by a function $n^{FB}(\theta, t)$, which specifies the number of agents which optimally do research in period t when the current highest value is θ. Gal et al. (1981) and Morgan (1983) have shown that there exists a global stopping value θ^g such that $n^{FB}(\theta, t) = 0$ for $\theta \geq \theta^g$ and that the number $n^{FB}(\theta, t)$ is decreasing in θ and increasing in t. Thus, a global stopping threshold is optimal even with a finite horizon, but the optimal number of agents will generally change non-monotonically over time.[24] Without additional assumptions it is difficult to say more about the optimal innovation contest. To make progress, we assume that all innovations are of similar value, which we will call a breakthrough.

[22]Since m can be made arbitrarily large, there always exists an ITC where m/n is larger than the agents' evaluation cost. We discuss additional ways in which the principal can change the allocation rule of the interim transfer in the conclusion.

[23]Since in Taylor (1995) the agents can observe the value of their innovation at no cost, it is not obvious how the equilibrium of an FPC would change if evaluating innovations was costly for the agents.

[24]To understand this non-monotonicity, suppose that the current best innovation available to the principal is $\theta^1 = 0$. As the deadline T approaches, the risk that the final innovation is worth nothing increases. To mitigate this risk, the principal *increases* the number of firms investing in every period. However, suppose that in the period $T - 1$ an innovation with relatively high value $\theta^k > 0$ is discovered. Since now the principal will obtain at least θ^k, she is less willing to risk costly duplication of effort, so in the final period she *reduces* the number of firms investing.

Assumption 1 *The set of values of innovation is given by* $\Theta = \{0, \theta^b, \theta^{b+1}, \ldots, \theta^K\}$ *where* $\theta^K - \theta^b < C$*.*

This assumption fits well with environments where the goal of the principal is a scientific discovery. For example, determining that tuberculosis is caused by bacteria, that DNA has a double-helix structure, or that Fermat's Last Theorem is true are examples of breakthroughs where the main value comes from the insight. A clearer image of the double-helix or a more elegant proof of Fermat's Last Theorem would certainly be valuable, but not nearly as much as the original insight. Environments which do not fit well tend to be more applied innovations, where the values are significantly dispersed. For example, a pharmaceutical innovation can be effective at treating a disease, but it can be costly to produce or have significant side effects. Improvements along either of those dimensions can significantly increase the value of that innovation to the principal.

Given Assumption 1, the first best takes a simple form. Once any breakthrough has been achieved, further research is not worthwhile and the principal stops the contest. As long as no breakthrough has been achieved, the principal will weakly increase the research effort in every period. The ITC we introduced in Section 2 cannot implement the first best, because the number of participants in an ITC is constant over time. However, we can easily generalize the notion of an ITC to accommodate the increasing number of researchers over time. Let a *generalized interim-transfer contest* (gITC) be a tuple $\Gamma = \langle \mathbf{E}, p, \mathbf{m}, \mathbf{n}, T \rangle$, where $\mathbf{n} = [n_1, n_2, \ldots, n_T]$ is a vector specifying the number of participants in each period, $\mathbf{E} = [E_1, \ldots, E_T]$ is the vector of entry fees to be paid by participants entering in period t, and $\mathbf{m} = [m_1, \ldots, m_{T-1}]$ is the vector of interim transfers, where transfer m_t is paid in period t if the contest continues to the next period.

The game induced by a gITC is analogous to the one induced by a regular ITC, but with two differences. First, in any period t with $n_t > n_{t-1}$, $n_t - n_{t-1}$ additional agents are randomly chosen from the set N and invited to participate in the contest and to pay the entry fee E_t.[25] Agents may only enter the contest in the period in which they are invited. Second, the interim transfer is not the same in each period, but can vary from period to period. Together with Assumption 1, a gITC implements the first-best outcome.

Proposition 4 *Given Assumption 1 and a finite deadline* T*, the first-best outcome can be implemented using a generalized ITC with sufficiently high* p *and interim transfers*

$$m_t = p(1-\delta) + \delta(\Delta(\theta^b, n^{FB}_{t+1}) + (n^{FB}_{t+1} - n^{FB}_t)E_{t+1}) - \theta^b \quad \textit{for } t = 1, \ldots, T-1.$$

[25] Note that in each period at most a finite number of researchers is invited. Since the horizon is finite and the contest terminates in period T, the highest possible number of participants in the contest is n_T. The assumption that the set of available agents N is infinite is made for simplicity, our results would hold if N was finite but sufficiently large. An additional implicit assumption that we make is that there are no recruiting costs. If recruiting agents was costly, the first best would remain qualitatively unchanged (with a weakly increasing number of agents over time), even if the exact vector $\mathbf{n}^{FB}$ would change. Furthermore, gITC would still implement the first-best outcome for the principal.

The proof mirrors that of Proposition 1 and the intuition is unchanged, that is, the interim transfers ensure the principal's adherence to the global stopping equilibrium and the final prize incentivizes the agents to conduct research until they achieve a breakthrough.

5 Related Literature

The seminal paper on dynamic research contests is Taylor (1995) on which we build our model.[26] He shows that a T-period FPC with N agents uniquely implements an individual stopping equilibrium among the agents. Further, Taylor shows that it is optimal to limit the number of agents in the contest and that the principal can extract the entire ex-ante surplus using appropriate entry fees. Rieck (2010) considers a variation of Taylor's framework, which enables him to study the role of information revelation. Depending on the parameters, the principal may be better off with or without information revelation and firms may voluntarily reveal information. As we mentioned before, an ITC is more robust to information disclosure than an FPC.

Recently, a number of papers have used bandit models to study the problem of incentive provision for dynamic research activity. In bandit models it is unclear ex ante if the innovation in question can actually be successfully realized, so in contrast to our setting, these models focus on learning over time. Halac, Kartik, and Liu (2017) consider the optimal design of contests for innovation where the principal chooses the prize-sharing scheme and a disclosure policy which determines what information is revealed to the agents about their respective outcomes. Similarly to our setting, the first-best features a global stopping equilibrium. However, they find that a contest which does not entail a global stopping equilibrium can be optimal in the presence of learning. In particular, in a broad class of contests it is optimal to stop the contest only once a certain number of agents had a success and to share the prize between them.[27] Along similar lines Green and Taylor (2016) consider the role of breakthroughs in a single-agent contracting environment. In contrast to our framework, the research outcome can be contracted upon and the problem the principal faces is how to optimally induce effort over time using a first deadline for the breakthrough, a second deadline for the final outcome and a monetary transfer. In their paper, the monetary transfer is decreasing over time, which induces the agent to aim for an early success. Thus, the slope of the prize schedule is used to affect the agent's incentives. In contrast, in our paper, the final prize aligns the agents' incentives, while the interim transfers serve to align the principal's incentives.

Also related is the literature on optimal design of research contests in the static setting,

[26]Konrad (2009) provides an excellent overview of the literature on contests. See also Siegel (2009) for general results on all-pay auctions.

[27]For a related model featuring partial progress see Bimpikis, Ehsani, and Mostagir (2019).

where the seminal contribution is Che and Gale (2003). They show that with symmetric agents, the optimal contest is a scoring auction and the optimal number of agents is two. When agents are asymmetric, the optimal contest is still an auction with two agents, but the optimal auction handicaps the more efficient agents.[28] The innovation in their setting is deterministic, so there is no sampling benefit from having more than two agents.[29] Several other directions have been explored in the static setting. Letina and Schmutzler (2019) consider the optimal contest design when the agents can choose their approach to innovation and the principal attempts to give them incentives to diversify their approaches because of the resulting option value. They find that the optimal contest is a bonus tournament, where a winner gets a fixed prize, plus a bonus if he outperforms the second best agent with a high enough margin. Olszewski and Siegel (2019) provide a novel approach to optimal contest design with many agents. In a very general setting, they characterize the optimal prize structure and find that with convex costs and risk-averse agents, multiple prizes are optimal.

Lang, Seel, and Strack (2014) is related to our result about the optimal contest length T. They consider a two-player FPC where agents exert effort over time and breakthroughs arrive according to a Poisson process. The agent with the most breakthroughs wins. They find that the principal can be better off with a shorter deadline. Seel (2018) characterizes the optimal deadline in a two-player FPC where the player with highest effort wins. He finds that a short deadline is optimal. In our paper, if the principal was limited to a FPC, shorter deadlines would also be optimal. However, since an ITC is stopped as soon as the threshold is reached, infinite deadlines are optimal.

Related to our implementation result is Kruse and Strack (2015). They study a dynamic principal-agent problem, where the agent observes realizations of a stochastic process over time. They show that for any threshold value, the principal can induce the agent to stop the game as soon as the process is above the threshold by committing to an appropriate schedule of transfers which depend only on the period when the game is stopped. In our paper, the stochastic process comes from the research done by the agents and the goal of the contest is to incentivize the agents to engage in research.

There is a growing empirical literature on dynamic research contests.[30] Using data on software contests Boudreau, Lacetera, and Lakhani (2011) find that increasing the number of participants reduces average effort but increases the chance of getting a very high quality innovation. Also using data on software contests Boudreau, Lakhani, and Menietti (2016) find that the results derived in Moldovanu and Sela (2001) generally perform quite well. In particular, they find that the response of participants to an increase

[28] Discrimination in contests is also studied in Pérez-Castrillo and Wettstein (2016).

[29] See for example Terwiesch and Xu (2008), Schöttner (2008) and Letina (2016).

[30] For a survey of experimental work on contests see Dechenaux, Kovenock, and Sheremeta (2015).

in the number of competitors yields heterogeneous responses. Namely, low ability agents respond weakly, medium ability agents decrease their efforts while high ability agents increase their efforts. Bhattacharya (2018) estimates a multistage research contest based on the U.S. Department of Defense SBIR program. Lemus and Marshall (2019) estimate a dynamic contest model using data from Kaggle.com.[31] Using their model, they can evaluate counterfactual contest designs. One of their main findings is that having a public leaderboard during the contest improves the outcomes.

6 Conclusion

The goal of the present paper is to improve our understanding of the optimal design of research contests, which have recently seen a rapid expansion in practice. It does so by constructing a theoretical benchmark which enables the principal to implement the first-best outcome. However, our model relies on several strong assumptions which may limit the extent to which our result can be applied in practice. We have already mentioned the importance of entry fees and the assumption that the agents are not budget constrained. In addition, we model research as an independent draw in each period. This simplifies analysis, but in reality the research is likely to be cumulative, with researchers building on their previous discoveries. Finally, to implement the optimal contest, the principal needs detailed information about the environment, including the level of research costs and the distribution from which innovations are drawn. In addition, this information is not updated over time. Analyzing the effects of these assumptions is an exciting avenue for future research.

The main result of this paper is that an interim-transfer contest constitutes the optimal mechanism. As was already mentioned, the purpose of interim transfers is to provide incentives to the principal for stopping the contest once an innovation of sufficient quality has been discovered. This incentive is present whenever the principal commits to paying m in every period of the contest. The exact rule for allocating m in a particular innovation contest can then be chosen in response to the specific setting of that innovation contest. For example, if maintaining the competitive balance among contestants is an issue in a given contest, then splitting the transfer m equally among all agents (as is done in our ITC) alleviates that concern.[32] Alternatively, the principal may want to learn about the innovations as soon as they are discovered, even if the innovations are below the threshold at which the contest is stopped. This could be either for marketing purposes or because the principal can start benefiting from innovations while the contest is ongoing. In this case, the principal needs to provide incentives to the agents to also report innovations

[31] They host online prediction competitions, which are closely related to innovation contests.

[32] For an example of this see Möller (2012).

below the threshold. This can be done by awarding the entire transfer m to the agent who currently has the best innovation as a progress prize.[33]

In a similar vein, an alternative contest which does not entail interim transfers but instead features a dynamic prize schedule does almost as well as the interim-transfer contest. In this alternative contest, the principal commits to a sequence of prizes $p_1, ..., p_T$. If the principal ends the contest in period t, the prize p_t has to be paid out. The difference between two prizes p_t and p_{t+1} plays a similar role to our interim transfer. This alternative contest with a dynamic prize schedule can also implement global stopping equilibria, but only with finite deadlines. Thus, the first-best outcome can be approximated arbitrarily well, but not fully achieved.

In addition to innovation contests, interim-transfer contests can be used to study promotion tournaments within firms. In a promotion tournament, firms commit to paying a salary (i.e., an interim transfer) to a group of workers, until one worker is promoted (which is the final prize). Our model suggests that the payment of salaries and the timing of the promotion are linked. Of course, in a promotion tournament the firm may be interested in the total effort provided by all workers or in selecting the best worker among the contestants. As this example illustrates, the insights from our paper can potentially extend to settings beyond innovation contests.

[33]Awarding the transfer m to the agent with the current highest innovation implies that the principal would have to potentially inspect submissions in every period.

Appendix

A Proofs

A.1 Proof of Proposition 1

The proof proceeds as follows. First, three lemmas show that no profitable one-shot deviation exists after any history of the game. Then, we show that no one-shot deviation implies that no profitable deviation exists.

Lemma 1 *There exists no profitable one-shot deviation for the principal.*

Proof. First, observe that the principal cannot profitably deviate when allocating the prize. Thus, we only need to consider deviations in instances when the principal chooses whether to stop or continue the contest. If $T < \infty$ there is a final period. In this final period the contest has to end and thus the principal has no action to take. Next, consider any period $t < T$. Suppose an innovation of value $\theta^k \geq \theta^g$ has been submitted to the principal. Stopping yields $\theta^k - p$, whereas continuing yields $-m + \delta(\Delta(\theta^k, n) - p)$. Thus, stopping is optimal whenever

$$m \geq p(1-\delta) + \delta\Delta(\theta^k, n) - \theta^k.$$

Recall that $m = p(1-\delta) + \delta\Delta(\theta^g, n) - \theta^g + \varepsilon$. Simple algebra shows that $\Delta(\theta, n) - \theta$ is strictly decreasing in θ. Thus, the principal will stop the contest if a value $\theta^k \geq \theta^g$ has been submitted.

Suppose now a value $\theta^k < \theta^g$ has been submitted. We will show in three steps that stopping is not optimal. Steps 1 and 2 cover the case when T is finite, while Step 3 deals with the infinite horizon case.

Step 1. Denote with $U_0(\sigma|\theta^k, t)$ the expected utility to the principal of having the highest value θ^k in period t and given a strategy σ. We will prove by induction that $U_0(\sigma|\theta^k, t) > \theta^k - p$, which shows that no profitable one-shot deviation occurs. For the base step, we show that $U_0(\sigma|\theta^k, T-1) > \theta^k - p$. We can write

$$\begin{aligned} U_0(\sigma|\theta^k, T-1) &= -m + \delta\left(F^n(\theta^k)\theta^k + \sum_{j=k+1}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j\right) - \delta p \\ &= \theta^g - p - \delta(\Delta(\theta^g, n) - \Delta(\theta^k, n)) - \varepsilon. \end{aligned} \tag{4}$$

Thus, $U_0(\sigma|\theta^k, T-1) > \theta^k - p$ if and only if

$$\theta^g - \theta^k - \varepsilon - \delta(\Delta(\theta^g, n) - \Delta(\theta^k, n)) > 0.$$

If $g > k$, then simple algebra gives

$$\Delta(\theta^g, n) - \Delta(\theta^k, n) = \sum_{j=k}^{g-1} F^n(\theta^j)(\theta^{j+1} - \theta^j). \tag{5}$$

Using this in the inequality above, $U_0(\sigma|\theta^k, T-1) > \theta^k - p$ if and only if

$$\theta^g - \theta^k - \delta \sum_{j=k}^{g-1} F^n(\theta^j)(\theta^{j+1} - \theta^j) - \varepsilon > 0$$
$$\sum_{j=k}^{g-1} (1 - \delta F^n(\theta^j))(\theta^{j+1} - \theta^j) - \varepsilon > 0,$$

which is satisfied whenever

$$(1 - \delta F^n(\theta^{g-1}))(\theta^g - \theta^{g-1}) > \varepsilon.$$

Step 2. To prove the inductive step, we show that if $U_0(\sigma|\theta^k, t+1) > \theta^k - p$ then also $U_0(\sigma|\theta^k, t) > \theta^k - p$. We can write

$$U_0(\sigma|\theta^k, t) = -m + \delta \left(F^n(\theta^k) U_0(\sigma|\theta^k, t+1) + \sum_{j=k+1}^{K} (F^n(\theta^j) - F^n(\theta^{j-1})) U_0(\sigma|\theta^j, t+1) \right).$$

Since $U_0(\sigma|\theta^j, t+1) = \theta^j - p$ for all $j \geq g$, and by the previous step $U_0(\sigma|\theta^j, t) > \theta^j - p$ for all $k \leq j < g$, then we can write

$$\begin{aligned} U_0(\sigma|\theta^k, t) &\geq -m + \delta \left(F^n(\theta^k)(\theta^k - p) + \sum_{j=k+1}^{K} (F^n(\theta^j) - F^n(\theta^{j-1}))(\theta^j - p) \right) \\ &= -m - \delta p + \delta \Delta(\theta^k, n) \\ &= \theta^g - p - \delta(\Delta(\theta^g, n) - \Delta(\theta^k, n)) - \varepsilon. \end{aligned}$$

Observe that the last expression is identical as equation (4) and the proof that $U_0(\sigma|\theta^k, t) > \theta^k - p$ proceeds analogously.

Step 3. In the infinite horizon case, the expected utility when the principal follows σ, given any $\theta^k < \theta^g$ is equal to the value of search. Thus we can write

$$U_0(\sigma|\theta^k) = -m + \delta \left(F^n(\theta^{g-1}) U_0(\sigma|\theta^k) + \sum_{j=g}^{K} (F^n(\theta^j) - F^n(\theta^{j-1}))(\theta^j - p) \right)$$

$$(1 - \delta F^n(\theta^{g-1})) U_0(\sigma|\theta^k) = -m - \delta(1 - F^n(\theta^{g-1}))p + \delta \sum_{j=g}^{K} (F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j$$

$$(1 - \delta F^n(\theta^{g-1})) U_0(\sigma|\theta^k) = \theta^g (1 - \delta F^n(\theta^{g-1})) - p(1 - \delta F^n(\theta^{g-1})) - \varepsilon$$

$$U_0(\sigma|\theta^k) = \theta^g - p - \frac{\varepsilon}{1 - \delta F^n(\theta^{g-1})}$$

Thus, $U_0(\sigma|\theta^k) > \theta^k - p$ for all $k < g$. ∎

Lemma 2 *There is no profitable one-shot deviation at the submission stage for the agent.*

Proof. Observe that submitting an innovation that has value below θ^g is never profitable. Thus we only need to consider the decision of an agent who has an innovation of value $\theta^k \geq \theta^g$. Suppose the state of the world is such that another agent has a value $\theta \geq \theta^k$. Then, submitting yields a weakly higher payoff, as it could mean the agent wins the prize, whereas not submitting yields zero as the contest ends for sure. Finally, suppose the state of the world is such that no other agent has a value $\theta \geq \theta^k$.

We need to consider two cases: when $\theta^k = \theta^K$ and when $\theta^k < \theta^K$. Suppose first that $\theta^k = \theta^K$. The payoff of following the equilibrium strategy and submitting is p. One-shot deviation is to not submit, then not do research, and then submit. The payoff of this deviation is

$$\frac{m}{n} + \delta \mathcal{P}_{t+1}(\sigma'|\theta^K, t)p$$

where $\mathcal{P}_{t+1}(\sigma'|\theta^K, t)$ is the probability that the agent wins the contest in period $t+1$ given that he has the quality θ^K in period t and follows the deviation strategy σ'. The deviation will not be profitable if $p \geq m/n + \delta \mathcal{P}_{t+1}(\sigma'|\theta^K, t)p$. It is sufficient to show that $p \geq m + \delta \mathcal{P}_{t+1}(\sigma'|\theta^K, t)p$. Substituting $m = (1-\delta)p + \delta\Delta(\theta^g, n) - \theta^g + \varepsilon$ and rearranging, this is equivalent to

$$p \geq \frac{\delta\Delta(\theta^g, n) - \theta^g + \varepsilon}{\delta(1 - \mathcal{P}_{t+1}(\sigma'|\theta^K, t))}.$$

The probability of the event that all the opponents obtain θ^K and one of them wins is given by $(1 - F(\theta^{K-1}))^{n-1}(n-1)/n$. Since $\mathcal{P}_{t+1}(\sigma'|\theta^K, t) < 1 - (1 - F(\theta^{K-1}))^{n-1}(n-1)/n$, it is sufficient to show that

$$p \geq \frac{n(\delta\Delta(\theta^g, n) - \theta^g + \varepsilon)}{\delta(n-1)(1 - F(\theta^{K-1}))^{n-1}}.$$

which always holds since $p \geq \bar{p}$.

In the case $\theta^k < \theta^K$ the one-shot deviation is to not submit, invest, and then submit. However, observe that the quality in the next period cannot be greater than θ^K, which implies that the deviation payoff is less than $m/n + \delta(\mathcal{P}_{t+1}(\sigma|\theta^K, t+1)p - C)$. As this is less than in the previous case, this deviation is also not profitable. ■

Lemma 3 *There exists no profitable one-shot deviation at the research stage for the agent.*

Proof. Suppose that the highest quality agent i has in period t is θ^k. In what follows, we show that for $p \geq \bar{p}$ investing is optimal for all $\theta^k < \theta^K$ and it is not optimal for $\theta^k = \theta^K$. Let σ' be a strategy profile that coincides with the equilibrium candidate σ with the exception of the agent i's action in the investment stage in period t. Thus, it is a one-shot deviation. First note that a deviation in the case $\theta^k = \theta^K$ would imply investing when the agent has the highest feasible quality. This is trivially never optimal, as the agent incurs research costs without an increase in quality. Thus, focus on the case $\theta^k < \theta^K$ where a deviation is to not invest.

Denote the expected utility of agent i following the strategy σ from period t in which his highest quality is θ^k with $U_i(\sigma|\theta^k, t)$. A one-shot deviation is not profitable if

$$U_i(\sigma|\theta^k, t) - U_i(\sigma'|\theta^k, t) \geq 0. \tag{6}$$

As before, let $\mathcal{P}_s(\sigma|\theta^k, t)$ be the probability that the agent i wins the contest in period $s \geq t$, following the strategy σ from period t in which the highest quality was θ^k.

First consider the case $\theta^g \leq \theta^k < \theta^K$. In this case, the game will end with certainty in period t and the LHS of inequality (6) reads

$$U_i(\sigma|\theta^k, t) - U_i(\sigma'|\theta^k, t) = -C + p(\mathcal{P}_t(\sigma|\theta^k, t) - \mathcal{P}_t(\sigma'|\theta^k, t)).$$

Due to perfect recall, in any state of the world in which the agent i wins following the strategy σ', he also wins following the strategy σ. Following strategy σ', agent i has a zero probability of winning if all the opponents have θ^K, while that probability is positive if he follows σ. The event that all the opponents draw θ^K and agent i following the strategy σ wins, happens with probability of at least $(1 - F(\theta^{K-1}))^n/n$. Thus, to show that the inequality (6) holds, it is sufficient to show that

$$-C + p\frac{(1 - F(\theta^{K-1}))^n}{n} \geq 0 \tag{7}$$

$$p \geq \frac{nC}{(1 - F(\theta^{K-1}))^n}$$

which always holds since $p \geq \bar{p}$.

The only remaining case is $\theta^k < \theta^g$, which we now consider. In this case, the agent i could not have observed a deviation by the principal, so he believes all of his opponents have values below θ^g. First suppose that $t = T$ so that the contest ends with certainty in period t. Then the analysis above applies. Next, suppose that $t < T$. Since σ and σ' coincide after t, that we can write the expected utilities in the following way:

$$\begin{aligned} U_i(\sigma'|\theta^k, t) &= F^{n-1}(\theta^{g-1}) \left[\frac{m}{n} + \delta U_i(\sigma|\theta^k, t+1)\right] \\ U_i(\sigma|\theta^k, t) &= p\mathcal{P}_t(\sigma|\theta^k, t) + F(\theta^k)F^{n-1}(\theta^{g-1}) \left[\frac{m}{n} + \delta U_i(\sigma|\theta^k, t+1)\right] \\ &\quad + \sum_{j=k+1}^{g-1} \left(F(\theta^j) - F(\theta^{j-1})\right) F^{n-1}(\theta^{g-1}) \left[\frac{m}{n} + \delta U_i(\sigma|\theta^j, t+1)\right] - C. \end{aligned}$$

Since $U_i(\sigma|\theta^j, t+1) \geq U_i(\sigma|\theta^k, t+1)$ for all $\theta^j > \theta^k$, we can write

$$U_i(\sigma|\theta^k, t) \geq p\mathcal{P}_t(\sigma|\theta^k, t) + F(\theta^{g-1})F^{n-1}(\theta^{g-1}) \left[\frac{m}{n} + \delta U_i(\sigma|\theta^k, t+1)\right] - C.$$

Furthermore, since $\mathcal{P}_t(\sigma|\theta^k, t) \geq (1 - F(\theta^{g-1}))F^{n-1}(\theta^{g-1}) + (1 - F(\theta^{K-1}))^n/n$ and $p \geq m/n + \delta U_i(\sigma|\theta^k, t+1)$ by Lemma 2, we can write

$$U_i(\sigma|\theta^k, t) \geq \frac{(1 - F(\theta^{K-1}))^n}{n} p + F^{n-1}(\theta^{g-1}) \left[\frac{m}{n} + \delta U_i(\sigma|\theta^k, t+1)\right] - C.$$

Then,

$$U_i(\sigma|\theta^k, t) - U_i(\sigma'|\theta^k, t) \geq \frac{(1 - F(\theta^{K-1}))^n}{n} p - C \geq 0$$

which holds by the same argument as for Inequality (7). ■

We conclude the proof by showing that since no one-shot deviation exists, then no profitable deviation exists at all. First, observe that if T is finite, then the result follows by Theorem 1 of Hendon, Jacobsen, and Sloth (1996). If T is infinite, then the game is continuous at infinity and the result follows by Corollary 2 of Hendon et al. (1996).

A.2 Proof of Proposition 2

The claim follows from the lemmas below. Throughout the proof, we will again use $\mathcal{P}(\sigma|\theta_i^{max}, t)$ to denote the probability that the agent i wins the contest in period t, given strategy σ and his current highest value θ_i^{max}. Similarly, $\mathcal{P}^c(\sigma|\theta_i^{max}, t)$ will denote the probability that the contest continues to the next period. We will use $U_i(\sigma|\theta_i^{max}, t)$ to denote the expected payoff of player i (agent or the principal).

Lemma 4 *In any PBE, $\sigma_0(\theta^K) = NCont$.*

Proof. Suppose not. We show that $NCont$ is a profitable one-shot deviation. We can write

$$U_0(NCont, \sigma|\theta^K) = \theta^K - p.$$

The principal's expected utility after $Cont$, given a strategy that terminates the contest in $\bar{t} \leq \infty$ periods, is

$$U_0(Cont, \sigma|\theta^K) = \delta^{\bar{t}}(\theta^K - p) - \sum_{t=0}^{\bar{t}-1} \delta^t m.$$

Then,

$$\begin{aligned} U_0(NCont, \sigma|\theta^K) - U_0(Cont, \sigma|\theta^K) &= \left(1 - \delta^{\bar{t}}\right)\left(\theta^K - p\right) + \sum_{t=0}^{\bar{t}-1} \delta^t m \\ &= \frac{1-\delta^{\bar{t}}}{1-\delta}\left((1-\delta)\theta^K + \delta\Delta(\theta^g, n) - \theta^g + \varepsilon\right) \\ &\geq \frac{1-\delta^{\bar{t}}}{1-\delta}\left(\theta^K - \theta^g + \varepsilon\right) > 0. \end{aligned}$$

■

Lemma 5 *At any submission stage and for any $\theta_i^{max} \in \Theta$*

$$\mathcal{P}(S, \sigma|\theta_i^{max}, t) + \mathcal{P}^c(S, \sigma|\theta_i^{max}, t) \geq \mathcal{P}(NS, \sigma|\theta_i^{max}, t) + \mathcal{P}^c(NS, \sigma|\theta_i^{max}, t).$$

Proof. If $t = T$ or $\theta^s = \theta^0$, then the contest ends for sure in the current period. Since the probability of winning in case the contest ends is minimized after NS, submitting can only weakly increase the probability of winning. Next, consider $t < T$ or $\theta^s > \theta^0$. Then $\mathcal{P}(NS, \sigma|\theta_i^{max}, t) = 0$ and following NS by agent i, the contest continues only if all agents either do not submit or submit a value below θ^s. If $\theta_i^{max} < \theta^s$ then the contest continues in all the same states of the world as with NS. Hence $\mathcal{P}^c(S, \sigma|\theta_i^{max}, t) = \mathcal{P}^c(NS, \sigma|\theta_i^{max}, t)$ and $\mathcal{P}(S, \sigma|\theta_i^{max}, t) = \mathcal{P}(NS, \sigma|\theta_i^{max}, t) = 0$. If $\theta_i^{max} \geq \theta^s$ then the agent wins the contest in all the states of the world in which the contest would continue after NS. Hence, $\mathcal{P}(S, \sigma|\theta_i^{max}, t) \geq \mathcal{P}^c(NS, \sigma|\theta_i^{max}, t)$ and since $\mathcal{P}(NS, \sigma|\theta_i^{max}, t) = 0$ the conclusion follows.
■

Lemma 6 *In any ITC, for any σ, θ_i^{max} and t,*

$$p - (m/n + \delta U_i(\sigma|\theta_i^{max}, t+1)) \geq \frac{(n-1)(1-\delta)p - \delta\Delta(\theta^g, n) + \theta^g - \varepsilon}{n} > 0.$$

Proof. If the contest terminates in $\bar{t}$ periods, where $1 \leq \bar{t} \leq \infty$, then

$$\sum_{t=0}^{\bar{t}-1} \delta^t \frac{m}{n} + \delta^{\bar{t}} p \geq \frac{m}{n} + \delta U_i(\sigma|\theta_i^{max}, t+1)$$

which implies

$$\begin{aligned} p - (m/n + \delta U_i(\sigma|\theta_i^{max}, t+1)) &\geq p - \left(\sum_{t=0}^{\bar{t}-1} \delta^t \frac{m}{n} + \delta^{\bar{t}} p \right) \\ &= (1-\delta^{\bar{t}}) \left(\frac{(1-\delta)(n-1)p - \delta\Delta(\theta^g, n) + \theta^g - \varepsilon}{(1-\delta)n} \right) \\ &\geq \frac{(1-\delta)(n-1)p - \delta\Delta(\theta^g, n) + \theta^g - \varepsilon}{n}. \end{aligned}$$

Finally, note that

$$p > \bar{\bar{p}} \geq \frac{\delta\Delta(\theta^g, n) - \theta^g + \varepsilon}{(1-\delta)(n-1)},$$

which implies

$$\frac{(1-\delta)(n-1)p - \delta\Delta(\theta^g, n) + \theta^g - \varepsilon}{n} > 0.$$

■

Lemma 7 *If σ is a PBE and σ_0 is a threshold strategy, then for each agent i who has not observed a deviation by the principal $\sigma_i(\theta_i^{max}) = S$ for all $\theta_i^{max} \geq \theta^s > \theta^0$.*

Proof. Suppose not. Then there exists some history, some agent i and some $\theta_i^{max} \geq \theta^s > \theta^0$, such that $\sigma_i(\theta_i^{max}) = NS$ in a PBE. We will show that S is a profitable deviation. If $\theta_i^{max} \geq \theta^s$ then

$$U_i(S, \sigma|\theta_i^{max}, t) = \mathcal{P}(S, \sigma|\theta_i^{max}, t)p$$

and

$$U_i(NS, \sigma|\theta_i^{max}, t) = \mathcal{P}(NS, \sigma|\theta_i^{max}, t)p + \mathcal{P}^c(NS, \sigma|\theta_i^{max}, t)\left(\frac{m}{n} + \delta U_i(\sigma|\theta_i^{max}, t+1)\right).$$

Since the agent has not observed a deviation by the principal, then his beliefs are formed by Bayesian updating. Given that $\theta_i^{max} > \theta^0$, then $\mathcal{P}(S, \sigma|\theta_i^{max}, t) > \mathcal{P}(NS, \sigma|\theta_i^{max}, t)$. Next, by Lemma 5, $\mathcal{P}(S, \sigma|\theta_i^{max}, t) \geq \mathcal{P}(NS, \sigma|\theta_i^{max}, t) + \mathcal{P}^c(NS, \sigma|\theta_i^{max}, t)$ and by Lemma 6, $p > (m/n + \delta U_i(\sigma|\theta_i^{max}, t+1))$. Thus $U_i(S, \sigma|\theta_i^{max}, t) > U_i(NS, \sigma|\theta_i^{max}, t)$. ■

Lemma 8 *If σ is a PBE and σ_0 is a threshold strategy, then $\sigma_i(\theta_i^{max}) = I$ for all $\theta_i^{max} < \theta^s$ and for each agent i.*

Proof. Suppose not. Then there exists some history, some agent i and some $\theta_i^{max} < \theta^s$, such that $\sigma_i(\theta_i^{max}) = NI$ in a PBE. We will show that I is a profitable one-shot deviation. If $t < T$, we can write

$$U_i(I, \sigma|\theta_i^{max}, t) = -C + \mathcal{P}(I, \sigma|\theta_i^{max}, t)p + \mathcal{P}^c(I, \sigma|\theta_i^{max}, t)\left(\frac{m}{n} + \delta \mathbb{E}_{\theta_i^{max}} U_i(\sigma|\theta_i^{max}, t+1)\right)$$

and

$$U_i(NI, \sigma|\theta_i^{max}, t) = \mathcal{P}^c(NI, \sigma|\theta_i^{max}, t)\left(\frac{m}{n} + \delta U_i(\sigma|\theta_i^{max}, t+1)\right).$$

In every state of the world in which the agent i loses the contest in period t following the action I, he also loses after the action NI. Thus

$$\mathcal{P}(I, \sigma|\theta_i^{max}, t) + \mathcal{P}^c(I, \sigma|\theta_i^{max}, t) \geq \mathcal{P}^c(NI, \sigma|\theta_i^{max}, t).$$

By Lemma 4, the principal stops the contest if $\theta_0^{max} = \theta^K$. Thus,

$$\mathcal{P}(I, \sigma|\theta_i^{max}, t) \geq \frac{1 - F(\theta^{K-1})}{n}.$$

Furthermore, since U_i is non-decreasing in θ_i^{max} then

$$\begin{aligned} U_i(I, \sigma|\theta_i^{max}, t) - U_i(NI, \sigma|\theta_i^{max}, t) &\geq -C + \frac{1 - F(\theta^{K-1})}{n}\left(p - \frac{m}{n} - \delta U_i(\sigma|\theta_i^{max}, t+1)\right) \\ &\geq \frac{1 - F(\theta^{K-1})}{n}\left(\frac{(n-1)(1-\delta)p - \delta\Delta(\theta^g, n) + \theta^g - \varepsilon}{n}\right) - C > 0 \end{aligned}$$

where the second inequality follows from Lemma 6 and the third inequality follows from

$$p > \bar{\bar{p}} \geq \frac{n^2 C + (\delta\Delta(\theta^g, n) - \theta^g + \varepsilon)\left(1 - F(\theta^{K-1})\right)}{(1-\delta)(n-1)\left(1 - F(\theta^{K-1})\right)}.$$

If $t = T$, then the contest ends for sure. The proof is analogous to the proof of Lemma 9 below, and is therefore omitted. ■

Lemma 9 *If σ is a PBE and σ_0 is a threshold strategy with $\theta^s < \theta^K$, then $\sigma_i(\theta^s) = I$ for each agent i.*

Proof. Suppose not. Then there exists some history and some agent i, such that $\sigma_i(\theta^s) = NI$ in a PBE. We will show that I is a profitable one-shot deviation. Since $\theta_i^{max} \geq \theta^s$ and the principal follows a threshold strategy, the contest ends in period t for sure. Then, we can write

$$U_i(I,\sigma|\theta_i^{max},t) = -C + \mathcal{P}(I,\sigma|\theta_i^{max},t)p$$
$$U_i(NI,\sigma|\theta_i^{max},t) = \mathcal{P}(NI,\sigma|\theta_i^{max},t)p.$$

Observe that in every state of the world in which agent i wins following NI, he also wins following I. We now show that, whatever the beliefs of agent i, and for any strategy profile of his opponents and the principal which are compatible with PBE, his probability of winning increases by at least $(1-F(\theta^{K-1}))^n/(2n)$ after action I.

At any state of the world, one of the three following cases holds.

- Case 1: $\max_{j\neq i}\theta_j^{max} > \theta_i^{max}$. In this case, agent i loses the contest for sure after NI, while his winning probability is at least $(1-F(\theta^{K-1}))/n$ after I.

- Case 2: $\max_{j\neq i}\theta_j^{max} < \theta_i^{max}$. In this case, all opponents have a value strictly below θ^s, so by Lemma 8 they all invest in the current period. Thus, agent i's increase in the probability of winning is at least $(1-F(\theta^{K-1}))^n/n$.

- Case 3: $\max_{j\neq i}\theta_j^{max} = \theta_i^{max}$. In this case, agent i loses with probability at least $1/2$. If he invested, his probability of winning would increase by at least $(1-F(\theta^{K-1}))/(2n)$.

Thus, in any of the three possible cases, agent i's probability of winning increases by at least $(1-F(\theta^{K-1}))^n/(2n)$ after action I. Then we can write

$$U_i(I,\sigma|\theta_i^{max},t) - U_i(NI,\sigma|\theta_i^{max},t) \geq -C + p\frac{(1-F(\theta^{K-1}))^n}{2n} > 0,$$

where the last inequality holds from

$$p > \bar{\bar{p}} \geq \bar{p} \geq \frac{2nc}{(1-F(\theta^{K-1}))^n}.$$

■

Lemma 10 *If σ is a PBE and σ_0 is a threshold strategy, then $\theta^s = \theta^g$.*

Proof. From Proposition 1, we know that $\theta^s = \theta^g$ is an equilibrium. Here we show that $\theta^s \neq \theta^g$ is never an equilibrium. Suppose not. Then there exists a PBE, such that σ_0 is a threshold strategy and either $\theta^s > \theta^g$ or $\theta^s < \theta^g$. The proof follows in three steps. Steps 1 and 2 show that there exists a profitable one-shot deviation when $\theta^s > \theta^g$. Step 3 shows the same when $\theta^s < \theta^g$.

Step 1. Suppose that T is infinite and $\theta^s > \theta^g$. We will show that there exists a profitable one-shot deviation.

Suppose that $\theta_0^{max} = \theta^{s-1}$. The payoff from deviating to $NCont$ is $\theta^{s-1} - p \geq \theta^g - p$. The principal's expected utility from following σ_0 is

$$U_0(Cont, \sigma|\theta^{s-1}) = -m + \delta\left(F^n(\theta^{s-1})U_0(Cont, \sigma|\theta^{s-1}) + \sum_{j=s}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))(\theta^j - p)\right)$$

$$U_0(Cont, \sigma|\theta^{s-1}) = \frac{-m - \delta(1 - F^n(\theta^{s-1}))p + \delta\left(\sum_{j=s}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j\right)}{1 - \delta F^n(\theta^{s-1})}$$

The one-shot deviation to $NCont$ is profitable if

$$\theta^g - p > \frac{-m - \delta(1 - F^n(\theta^{s-1}))p + \delta\left(\sum_{j=s}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j\right)}{1 - \delta F^n(\theta^{s-1})}$$

which is equivalent to

$$\begin{aligned}
0 &> -\delta\Delta(\theta^g, n) + \delta F^n(\theta^{s-1})\theta^g - \varepsilon + \delta\left(\sum_{j=s}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j\right) \\
&\geq (F^n(\theta^{s-1}) - F^n(\theta^g))\theta^g - \sum_{j=g+1}^{s-1}\left(F^n(\theta^j) - F^n(\theta^{j-1})\right)\theta^j \\
&\geq (F^n(\theta^{s-1}) - F^n(\theta^g))\theta^g - \left(F^n(\theta^{s-1}) - F^n(\theta^g)\right)\theta^g
\end{aligned}$$

which is always satisfied.

Step 2. Suppose that T is finite and $\theta^s > \theta^g$. We will show that there exists a profitable one-shot deviation when $t = T - 1$ and $\theta_0^{max} = \theta^g$.

In this case, the payoff from deviating to $NCont$ is $\theta^g - p$. Now, the contest ends in the next period, so that the principal's expected utility from following σ_0 is

$$U_0(Cont, \sigma|\theta^g, t) = -m - \delta p + \delta\left(F^n(\theta^g)\theta^g + \sum_{j=g+1}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j\right).$$

The one-shot deviation to $NCont$ is profitable if

$$\theta^g - p > -m - \delta p + \delta\left(F^n(\theta^g)\theta^g + \sum_{j=g+1}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j\right)$$

which is equivalent to

$$-\delta F^n(\theta^g)\theta^g - \delta\sum_{j=g+1}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j + \delta\Delta(\theta^g, n) + \varepsilon > 0$$

$$-\delta\Delta(\theta^g, n) + \delta\Delta(\theta^g, n) + \varepsilon > 0.$$

Step 3. Suppose that $\theta^s < \theta^g$. We will show that there exists a profitable one-shot deviation when $\theta_0^{max} = \theta^s$.

Following the strategy σ_0 yields $\theta^s - p$. Consider a one-shot deviation to *Cont*. By Lemma 9, all agents will do research, so the principal's expected utility is

$$U_0(Cont, \sigma|\theta^s, t) = -m - \delta p + \delta\left(F^n(\theta^s)\theta^s + \sum_{j=s+1}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j\right)$$

The deviation is profitable if

$$-m - \delta p + \delta\left(F^n(\theta^s)\theta^s + \sum_{j=s+1}^{K}(F^n(\theta^j) - F^n(\theta^{j-1}))\theta^j\right) > \theta^s - p$$

$$\theta^g - \theta^s - \delta(\Delta(\theta^g, n) - \Delta(\theta^s, n)) - \varepsilon > 0$$

Using equation (5), a sufficient condition for the above inequality to hold is

$$\sum_{j=s}^{g-1}(\theta^{j+1} - \theta^j) - \delta\sum_{j=s}^{g-1}F^n(\theta^j)\left(\theta^{j+1} - \theta^j\right) > \varepsilon$$

$$\sum_{j=s}^{g-1}(1 - \delta F^n(\theta^j))\left(\theta^{j+1} - \theta^j\right) > \varepsilon$$

which is always satisfied. ■

A.3 Proof of Proposition 3

The first-best problem corresponds to the optimal search problem in Benkert et al. (2018). By Proposition 1 of Benkert et al. (2018) the first-best is to draw n^{FB} observations in each period until the value of at least θ^{FB} has been discovered. By Proposition 1, we know that there exists an ITC which can implement the global stopping threshold θ_N^g with n_N^{FB} and $T = \infty$, thus generating the first-best surplus. Then, by setting E appropriately, the principal can extract the entire expected surplus and achieve the first-best outcome.

A.4 Proof of Proposition 4

The game induced by a gITC and the equilibrium candidate are analogous to those described above. The proof proceed in the similar fashion. We first characterize the optimal sequence of agents $\mathbf{n}^{FB}$. Next, we show that no profitable one-shot deviation exists. Finally, since the game is finite, Theorem 1 of Hendon et al. (1996) implies that no profitable deviation exists at all.

Lemma 11 *Given Assumption 1,* $n^{FB}(0,t) \leq n^{FB}(0,t+1)$ *and* $n^{FB}(\theta,t) = 0$ *for all* $\theta \geq \theta^b$ *and* $t \leq T$*.*

Proof. It is straightforward that for any quality level $\theta^j \geq \theta^b$ the principal will stop doing research because $\theta^K - \theta^b < C$ and thus the cost of doing more research strictly outweighs the potential benefit. Thus, whenever the principal continues searching, she has a current highest quality of innovation of 0. Hence, the problem is as if the principal had no recall. Proposition 3 in Gal et al. (1981), which can be adapted to the current setting with discounting and a discrete set of innovation levels, then implies that the principal will want to employ an increasing number of agents as the deadline draws nearer. ∎

Lemma 12 *There exists no profitable one-shot deviation for the principal.*

Proof. In this final period the principal has no decision to make. Thus we only need to consider deviations in periods $t < T$. Suppose an innovation of value $\theta^k \geq \theta^b$ has been submitted to the principal. Stopping yields $\theta^k - p$, whereas continuing yields $-m_t + \delta(\Delta(\theta^k, n_{t+1}) - p + \mathcal{E}_{t+1})$, where $\mathcal{E}_{t+1} = (n_{t+1} - n_t)E_{t+1}$ is the sum of entry fees received by the principal in period $t+1$. Thus, stopping is optimal whenever

$$m_t \geq p(1-\delta) + \delta(\Delta(\theta^k, n_{t+1}) + \mathcal{E}_{t+1}) - \theta^k.$$

Recall that $m_t = p(1-\delta) + \delta(\Delta(\theta^b, n_{t+1}) + \mathcal{E}_{t+1}) - \theta^g$. Since $\Delta(\theta, n_{t+1}) - \theta$ is strictly decreasing in θ, the principal will stop the contest whenever a value $\theta^k \geq \theta^b$ has been submitted.

Suppose now a value $\theta^k < \theta^b$ has been submitted. Assumption 1 then implies $\theta^k = 0$. The payoff of stopping is $-p$, and continuing (since it constitutes the first best) always has a positive payoff. Hence, stopping is never optimal. ∎

Lemma 13 *There exists no profitable one-shot deviation at either the submission or the research stage for the agent.*

Proof. The proofs are analogous to the proofs of Lemma 2 and Lemma 3. ∎

Lemma 14 *There exists no profitable one-shot deviation at for agents who have not yet entered the contest.*

Proof. Whenever agents are invited to join the contest, they are chosen randomly from the set N. Thus, each agent has a probability 0 of being invited to the contest and therefore no incentive to conduct any research before being invited to participate. ∎

Finally, since T is finite and no profitable one-shot deviation exist, Theorem 1 of Hendon et al. (1996) implies that no profitable deviation exists at all.

References

Benkert, J.-M., I. Letina, and G. Nöldeke (2018): "Optimal Search from Multiple Distributions with Infinite Horizon," *Economics Letters*, 164, 15–18.

Bergemann, D. and J. Välimäki (2001): "Stationary Multi-Choice Bandit Problems," *Journal of Economic Dynamics and Control*, 25, 1585–1594.

Bhattacharya, V. (2018): "An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program," *Mimeo*.

Bimpikis, K., S. Ehsani, and M. Mostagir (2019): "Designing Dynamic Contests," *Operations Research*, 67, 339–356.

Boudreau, K. J., N. Lacetera, and K. R. Lakhani (2011): "Incentives and Problem Uncertainty in Innovation Contests: An Empirical Analysis," *Management Science*, 57, 843–863.

Boudreau, K. J., K. R. Lakhani, and M. Menietti (2016): "Performance Responses to Competition Across Skill Levels in Rank-Order Tournaments: Field Evidence and Implications for Tournament Design," *RAND Journal of Economics*, 47, 140–165.

Che, Y.-K. and I. Gale (2003): "Optimal Design of Research Contests," *American Economic Review*, 93, 646–671.

Dechenaux, E., D. Kovenock, and R. M. Sheremeta (2015): "A Survey of Experimental Research on Contests, All-Pay Auctions and Tournaments," *Experimental Economics*, 18, 609–669.

DeGroot, M. H. (1970): *Optimal Statistical Decisions*, Hoboken, New Jersey: John Wiley & Sons.

Fullerton, R. L. and R. P. McAfee (1999): "Auctioning Entry into Tournaments," *Journal of Political Economy*, 107, 573–605.

Gal, S., M. Landsberger, and B. Levykson (1981): "A Compound Strategy for Search in the Labor Market," *International Economic Review*, 22, 597–608.

Giebe, T. (2014): "Innovation Contests with Entry Auction," *Journal of Mathematical Economics*, 55, 165–176.

Gill, D. (2008): "Strategic disclosure of intermediate research results," *Journal of Economics & Management Strategy*, 17, 733–758.

Gordon, S. (2011): "Announcing Intermediate Breakthroughs in an R&D Race," *Mimeo*.

Green, B. and C. R. Taylor (2016): "Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects," *American Economic Review*, 106, 3660–99.

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

Hendon, E., H. J. Jacobsen, and B. Sloth (1996): "The One-Shot-Deviation Principle for Sequential Rationality," *Games and Economic Behavior*, 12, 274–282.

Konrad, K. A. (2009): *Strategy and Dynamics in Contests*, Oxford University Press.

Kruse, T. and P. Strack (2015): "Optimal Stopping with Private Information," *Journal of Economic Theory*, 159, 702–727.

Lang, M., C. Seel, and P. Strack (2014): "Deadlines in Stochastic Contests," *Journal of Mathematical Economics*, 52, 134–142.

Lemus, J. and G. Marshall (2019): "Dynamic Tournament Design: An Application to Prediction Contests," *Mimeo.*

Letina, I. (2016): "The Road not Taken: Competition and the R&D Portfolio," *RAND Journal of Economics*, 47, 433–460.

Letina, I. and A. Schmutzler (2019): "Inducing Variety: A Theory of Innovation Contests," *International Economic Review*, forthcoming.

Moldovanu, B. and A. Sela (2001): "The Optimal Allocation of Prizes in Contests," *American Economic Review*, 91, 542–558.

Möller, M. (2012): "Incentives Versus Competitive Balance," *Economics Letters*, 117, 505–508.

Morgan, P. B. (1983): "Search and Optimal Sample Size," *The Review of Economic Studies*, 50, 659–675.

Olszewski, W. and R. Siegel (2019): "Performance-Maximizing Contests," *Mimeo.*

Pérez-Castrillo, D. and D. Wettstein (2016): "Discrimination in a Model of Contests with Incomplete Information about ability," *International Economic Review*, 57, 881–914.

Rieck, T. (2010): "Information Disclosure in Innovation Contests," Bonn Econ Discussion Papers 16/2010, Bonn Graduate School of Economics.

Schöttner, A. (2008): "Fixed-Prize Tournaments Versus First-Price Auctions in Innovation Contests," *Economic Theory*, 35, 57–71.

Seel, C. (2018): "Contests with Endogenous Deadlines," *Journal of Economics & Management Strategy*, 27, 119–133.

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

Taylor, C. R. (1995): "Digging for Golden Carrots: An Analysis of Research Tournaments," *American Economic Review*, 85, 872–890.

Terwiesch, C. and Y. Xu (2008): "Innovation Contests, Open Innovation, and Multiagent Problem Solving," *Management Science*, 54, 1529–1543.