

Letina, Igor; Benkert, Jean-Michel

Conference Paper

Designing Dynamic Research Tournaments

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel -
Session: Competition and Tournaments, No. C04-V3

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Letina, Igor; Benkert, Jean-Michel (2016) : Designing Dynamic Research Tournaments, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel - Session: Competition and Tournaments, No. C04-V3, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg

This Version is available at:

<https://hdl.handle.net/10419/145738>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Designing Dynamic Research Tournaments

February 29, 2016

Abstract

This paper considers the optimal design of dynamic research tournaments when the buyer can set time-dependent prizes. We derive the buyer-optimal tournament and show that it entails an increasing prize schedule. Remarkably, this allows the buyer to implement a global stopping rule. In particular, the optimal tournament attains the first-best. More generally, we show that global stopping rules can be implemented robustly and compare them to individual stopping rules which have been analyzed in the existing literature. We conclude by discussing policy implications of our findings and highlight that global stopping rules combine the best aspects of innovation races and fixed prize tournaments.

Keywords: innovation, contests, dynamic tournaments, global stopping rule, breakthroughs, R&D

JEL: O32, D02, L19

Conference Draft; preliminary and incomplete

1 Introduction

Research tournaments¹ have a long history as a mechanism for inducing innovation. From navigation and food preservation, to aviation,² research tournaments have been used to find solutions to some of society’s most pressing problems. Recently, the use of research tournaments by both private and public sector has been expanding rapidly. As in the past, research tournaments are used in order to foster innovation in some of the most pressing and difficult issues that society is facing. Some examples of problems to which research tournaments have been applied include vaccine technology, antibiotics overuse, space flight, robotics and AI, and environment and energy efficiency.³ Given that the 2010 America Competes Reauthorization Act authorized US Federal agencies to use prizes and tournaments, it can be expected that the importance of research tournaments will only grow in the coming years.

Mistakes in tournament design can waste R&D funds and slow the development of important innovations. While some aspects of tournament design, like the effect of the number of competitors or the allocation of prizes, are well studied,⁴ less is known about the dynamic aspects of tournament design. At the same time, as Lang, Seel, and Strack (2014) point out, “there are surprisingly few multi-period contest models in which each player’s decision problem is dynamic.” This paper deals with the (buyer-)optimal design of dynamic research tournaments in precisely such an environment. We identify a novel design lever that the tournament designer can use in order to increase efficiency of the tournament — namely the fact that the prizes can differ depending on the time when they are awarded.⁵

We build on the model of Taylor (1995), where N sellers choose in each of the T periods whether to invest in research or not. Investment is costly, but in each period a seller does research he obtains an innovation, the quality of which is a random draw from some distribution F . The investment decisions, and the current highest quality of their innovations are private information to the sellers. If a seller reveals the innovation to the buyer, the buyer can costlessly and accurately determine its quality. However, the quality of innovations is not verifiable by outside parties. In particular, a contract which conditions on the innovation quality is not enforceable by courts. These assumptions are

¹Research tournaments are sometimes referred to as innovation contests or inducement prizes.

²1714 Longitude Prize, 1795 Napoleon’s Food Preservation Prize and 1919 Orteig Prize respectively.

³2012 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.

⁴See Che and Gale (2003), Moldovanu and Sela (2001) and Moldovanu and Sela (2006).

⁵To the best of our knowledge dynamic prizes have not been studied in the context of tournament design before. Of course, dynamic payoffs have been used in other context, for example in bandit models like in Green and Taylor (2015).

standard in the literature.⁶ In order to incentivize the sellers to invest in innovation, the buyer commits to a scheme that will, in some enforceable way, pay out a prize to the winning supplier(s). For example in Taylor (1995), the seller commits to paying a prize P at the end of the T periods. We will call this institution a *fixed-prize tournament*.

We deviate from Taylor's model in two important ways. First, we assume that the buyer can commit to a time-dependent prize scheme. This assumption only requires that the contracts can be time-dependent. There are many examples of such contracts — bills have to be paid by a certain date, with penalties for late payment, savings can be deposited for a fixed period of time, and delivery of parcels can be guaranteed within a specified time period. Even in the context of research tournaments we find examples of time-dependent contracts: in the 2006 Netflix prize in addition to the final Grand Prize, a \$50,000 Progress Prize could be awarded *each year* before the tournament ended.

Second, in the main section we assume that the innovations have what we will call a *breakthrough structure*. That means that each innovation falls broadly into two categories — it is either a breakthrough or not. All breakthrough innovations are worth approximately the same to the buyer and all non-breakthrough innovations are worth approximately the same. As the name suggests, breakthrough innovations are worth substantially more. This does not imply that there are only two quality levels — there may be many. What is important is that there is a distinction between the very valuable breakthrough innovations and relatively less valuable other innovations. For example, many tournaments explicitly state what is the goal of the tournament. Then a breakthrough is any innovation that reaches the goal. In the case of the 1996 Ansari X Prize, the publicly stated goal was to have two manned space flights within two weeks using the same spacecraft. The larger objective of the organizer was to spur the development of technology needed for civilian space flight. While the spacecraft that met the goal could be better or worse, the difference to the organizer did probably not matter as much as achieving the publicly stated goal. Thus, we believe that the breakthrough innovation structure fits well with many research tournaments.

Given the breakthrough innovation structure, we show that in the first-best⁷ all suppliers invest and conduct research until at least one of them achieves a breakthrough, at which point all research stops. This is an example of a global stopping rule. However, giving just a single prize at the end of the T periods of the tournament cannot implement a global stopping rule. Indeed, it uniquely implements an individual stopping rule, where each seller performs research until he achieves some threshold value which is determined by the size of the prize (Taylor, 1995). In order to implement a global stopping rule, the sellers must be incentivized to conduct research in every period and to truthfully reveal

⁶See for example Taylor (1995) and Che and Gale (2003).

⁷By first-best we mean maximizing total surplus in the absence of informational barriers.

that they have achieved an innovation, while the buyer has to be incentivized to stop the other firms from conducting further research. We show that by appropriately choosing a dynamic prize schedule, so that the prizes increase in each period the tournament continues, the buyer can be incentivized to stop the tournament as soon as a seller submits a breakthrough. This is done by setting the slope of the increase in prizes such that it is exactly equal to the marginal benefit to the buyer of one more round of research. At the same time, if the intercept of the prize schedule is high enough, the sellers will have an incentive to conduct research in every period and to truthfully reveal their breakthroughs as soon as they have them, because otherwise they would risk losing the tournament entirely. Thus, by appropriately choosing the slope and the intercept of the prize schedule the seller can implement any global stopping rule, and the first-best stopping rule in particular.

Next, we relax the assumption on breakthrough innovations. First we characterize the optimal global stopping rule with the more general innovation structure and show that it satisfies the property that the marginal benefit of one more round of research equals the marginal cost. We then show that any global stopping rule can be implemented with an appropriately designed schedule of prices. Thus, our implementation result is robust to the specification of the innovation structure. Moreover, we show that any global stopping rule can be implemented even if the number of sellers doing research changes from period to period, as long as the sequence is fixed *ex ante*. That is, our implementation result does not depend on the assumption that the set of sellers is constant for the entire duration of the tournament. This is especially important as some tournaments proceed in stages, so that some participants are eliminated in every stage. Finally, we consider the effect of a change in the horizon T of the tournament. We show that, with a global stopping rule, an increase in T always increases the payoff of the buyer. The intuition is that, with a global stopping rule, the tournament is stopped whenever an innovation of high enough quality is realized. Thus, the tournament only continues past the time T if it was beneficial to do so. Whereas in a fixed-prize tournament, an increase in the horizon T can lead to a decrease in buyer payoff. The reason for this is that the longer the time horizon is, the higher the chance of wasteful duplication.

Generally, one can distinguish between research tournaments and innovation races.⁸ To win a tournament a seller needs to have the best innovation at some specific date, whereas a seller needs to have a specific innovation as quickly as possible to win an innovation race. Thus, for an innovation race to be feasible, verifiability is necessary in order to determine whether some proposed innovation is indeed the innovation required to win the race. When innovation races are implemented in practice, a verifiable proxy is commonly used to determine whether an innovation meets the buyer's requirements. In

⁸See the discussion in Taylor (1995).

the Ansari X Prize, the goal of the tournament was to “incentivize the creation of a safe, reliable, reusable, privately-financed manned space ship to demonstrate that private space travel is commercially viable”, yet the sponsors used the objectively verifiable proxy of “build[ing] and launch[ing] a spacecraft capable of carrying three people to 100 kilometers above the Earth’s surface, twice within two weeks”.⁹ The advantage of a race is that it proceeds until an appropriate innovation has been developed and that it minimizes the wasteful duplication. When implementing research tournaments, the buyer in general does not have to use a proxy. However, other problems arise. In contrast to innovation races the sponsor of a research tournament needs to announce an end date at which the submissions will be judged. If the competitors are not given enough time, they may fail to produce a good enough innovation.¹⁰ If the deadline is very late, however, there is a risk of wasteful duplication.

Our dynamic prize tournament offers a solution to these problems. We show that when implementing a global stopping rule the buyer benefits from increasing the duration of the tournament, thereby allowing for very late deadlines which increase the chance of getting a sufficiently good innovation. At the same time, the nature of the global stopping rule ensures that throughout the tournament an innovation race is taking place, thus avoiding the risk of duplication arising in standard research tournaments. Moreover, this innovation race is implemented without requiring a proxy, hence eliminating this source of inefficiency. Effectively, a dynamic prize tournament inherits the best properties of both innovation races and research tournaments. Overall, our results indicate that using dynamic prizes could result in substantially more efficient tournaments.

The structure of the paper is as follows. Section 2 reviews the relevant literature. Section 3 presents the model. Section 4 characterizes the first-best and shows that it can be implemented with a dynamic-prize tournament. Section 5 considers extensions to the main model and in particular relaxes the assumption of breakthrough innovation. Section 6 concludes. All proofs are relegated to the appendix.

2 Related Literature

The seminal paper on dynamic research tournaments is Taylor (1995) on which we build our model.¹¹ He shows that a T -period tournament with N sellers where a fixed prize

⁹See the Ansari X Prize website <http://ansari.xprize.org>.

¹⁰The 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/>.

¹¹Konrad (2009) provides an excellent overview of the literature on contests. See also Siegel (2009) for general results on all-pay auctions.

is awarded at the end of the tournament uniquely implements an individual stopping rule among the sellers. Further, Taylor shows that it is optimal to limit the number of sellers in the tournament and that the buyer can extract the entire ex ante surplus using appropriate entry fees. As Taylor himself notes, however, his contest generally fails to implement the first-best, which entails a global stopping rule instead of an individual stopping rule (Gal, Landsberger, and Levykson, 1981). Moreover, Morgan (1983) shows that holding the number of sellers conducting research fixed across time is generally not optimal either, as the optimal number should vary over time depending on the currently highest quality among the sellers and the number of periods left.¹² Fullerton, Linster, McKee, and Slate (2002) compare Taylor’s fixed-prize tournament to an auction in the same setting. That is, the only change to Taylor’s framework is that at the end of the tournament an auction is used to allocate the prize among the sellers. The authors argue that this lowers the buyer’s informational requirements and they provide experimental evidence that this is more cost-effective when the buyer cannot charge entry fees. It is noteworthy that the buyer continues to employ an individual stopping rule when using an auction and therefore still fails to implement the first-best. Rieck (2010) considers a variation of Taylor’s framework which enables him to study the optimal information revelation. He shows that when the sellers’ research outcomes are publicly revealed there are essentially two thresholds instead of one. If the highest quality among the sellers is above the upper threshold all sellers stop research, if the highest quality is between the two thresholds only the leading seller stops research and if all qualities are below the lower threshold all sellers continue to do research. Depending on the parameters the buyer may be better off with or without information revelation.

Recently a number of papers have used bandit models to study the problem of incentive provision for dynamic research activity. Halac, Kartik, and Liu (forthcoming) consider the optimal design of contests for innovation when it is unclear ex ante whether or not the innovation in question can actually be successfully realized. Thus, in contrast to our setting the reason for research activity is not to get as good an innovation as possible, but rather to determine an innovation’s feasibility. The buyer who designs the contest can choose the prize-sharing scheme and a disclosure policy which determines what information is revealed to the sellers about their respective outcomes. Similarly to our setting, the first-best features a global stopping rule. However, Halac et al. (forthcoming) find that a contest which does not entail a global stopping rule can be optimal in the presence of private effort provision by the sellers. More generally, they show that in a broad class of contests it is optimal to stop the contest only once a certain number of sellers had a success and to share the prize between them. Bimpikis, Ehsani, and Mostagir (2014) study a closely related question to Halac et al. (forthcoming) but allow for partial progress, i.e., in

¹²See Morgan and Manning (1985) for general results on the first-best search rule in this environment.

order for an innovation to be feasible a milestone or breakthrough is necessary before the potential success is realized, and the buyer may set a prize for the breakthrough and the eventual success each. While they do not consider the question of optimal design, they show that the buyer may benefit from not revealing a partial success. Along similar lines Green and Taylor (2015) 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 buyer 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 seller’s incentives in contrast to our paper, where the increase in prizes over time serves to align the buyer’s incentives.

The seminal paper in the literature on optimal design of research tournaments in the static context is Che and Gale (2003). The authors consider a model where the innovation technology is deterministic. Once the innovations are developed, the sellers bid for a price at which the buyer can obtain their innovation. The buyer then chooses the bid which offers her the highest surplus, that is, the highest difference between the value of the information and the bid. The *contest design* consists of the choice of the set from which the sellers can choose their bids. This set of mechanisms turns out to be very flexible and includes the fixed-price tournament (when the set of allowable bids is a singleton) and the auction (when the set is \mathbb{R}^+). The authors show that with symmetric sellers the optimal contest is an auction and the optimal number of sellers is two. When sellers are asymmetric, the optimal contest is still an auction with two sellers, but the optimal auction handicaps the more efficient sellers.¹³ The major difference between Che and Gale (2003) and our paper is that we focus on the dynamic aspects of tournament design. Thus, the question of wasteful duplication of effort, which is the central issue we address with the dynamic prize tournament, does not show up in Che and Gale (2003). Another difference is the choice of innovation technology — when innovation is deterministic, as in Che and Gale (2003), there is no *sampling benefit* from having more than two sellers.¹⁴ Additionally, an auction gives market power to the sellers, and when the innovation technology is sufficiently random, an auction might perform badly as the seller profits from the good realizations.¹⁵ Several other directions have been explored in the static setting. Letina and Schmutzler (2015) consider the optimal contest design when the sellers can choose their approach to innovation and the buyer attempts to give them incentives to diversify their approaches because of the resulting option value. They find that the optimal tournament

¹³Discrimination in contests is also studied in Pérez-Castrillo and Wettstein (forthcoming).

¹⁴See for example Terwiesch and Xu (2008) and Letina and Schmutzler (2015).

¹⁵This is the case in Schöttner (2008) who shows that when innovation technology is sufficiently random, a fixed-price tournament can outperform an auction.

is a *bonus tournament* where a winner gets a fixed prize, plus a bonus if he outperforms the second best seller with a high enough margin. Moldovanu and Sela (2001) consider how a total prize sum should be split by a buyer who is maximizing the expected effort. They find that if the cost functions are concave or linear in effort, then the it is optimal to allocate the entire prize sum to the winner. If, on the other hand, cost functions are convex, it may be optimal to offer multiple prizes.

Lang et al. (2014) consider a two-player, fixed-price contest where sellers exert effort over time and breakthroughs arrive according to a Poisson process. The seller with the most breakthroughs wins. The authors consider the effect of changing the time T when the contest ends. They find that the buyer can be better off with a shorter deadline, which is exactly our finding in Proposition 6 for a fixed-prize tournament. However, we also show that this result does not hold for arbitrary contests — in our setting, a buyer is always better off with a longer deadline in a tournament implementing a global stopping rule.

More broadly related are papers examining if buyers should be split into several sub-contests or eliminated over time. Moldovanu and Sela (2006) find that if the buyer is maximizing the expected value of the highest effort it is beneficial to split the sellers into preliminary subcontests and to have the finalists compete against each other. In a different setting,¹⁶ Fu and Lu (2012) find that maximizing expected effort involves eliminating one seller in each round. An additional question that has been examined in the literature on dynamic contests is the question of how the information about the relative performance during the contest should be used. Gershkov and Perry (2009) study tournament with two agents and two stages and ask if a review should be conducted after stage one. They show that, assuming reviews are aggregated optimally, it is always beneficial to conduct a midterm review. In a related setting, Klein and Schmutzler (forthcoming) consider how a prize sum should be allocated for the first and second period performance, and how information about it should be used. They show that for large parameter regions, entire prize sum should be paid to the winner but both first-period and second-period performance should be weighted when determining the winner.

There is relatively little empirical work on dynamic research tournaments. 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 in the number of competitors yields heterogeneous responses. Namely, low ability agents respond weakly, medium ability agents decrease their efforts while high ability agents

¹⁶They use nested contest as in Clark and Riis (1996).

increase their efforts. We refer to the very recent survey Dechenaux, Kovenock, and Sheremeta (2015) for experimental work on tournaments.

3 The Model

There is a risk-neutral buyer who wants to procure an innovation and $N \geq 2$ ex ante identical risk-neutral sellers who can potentially produce the innovation by conducting research. If the buyer obtains the innovation in any period $t \in \{1, \dots, T\}$ with $T < \infty$, her payoff is $\theta - p$, where θ is the quality of the innovation and p is the sum of any transfers. A seller's payoff is $m - c$, where m is the sum of any transfers received and c is the total cost incurred through research activities. The innovations are of no intrinsic value to the sellers.

A seller can conduct research in any period t at per-period cost $C > 0$. In each period in which the seller 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 where $\Theta = \{\theta^1, \theta^2, \dots, \theta^K\}$. Suppose without loss that $\theta^{k+1} > \theta^k$ and we normalize $\theta^1 = 0$. Sellers can repeatedly conduct research and have perfect-recall, that is, they can access all their own previous innovations at any point in time. Initially, every seller is endowed with a worthless innovation and in each period a seller does not perform research he receives a worthless innovation.

We say a research tournament features *breakthrough innovation structure* if there exists some innovation $\theta^b \in \Theta$ such that all innovations below θ^b are worth very little to the buyer, while all innovations at or above θ^b are worth approximately the same. Thus, there are essentially two levels of innovation — breakthroughs and low-value innovations. In Section 5 we will relax the assumptions on Θ . Formally, we will say that the innovation process has a breakthrough innovation structure if the following assumption is satisfied.

Assumption 1 (breakthrough innovation structure) *There exists $\theta^b \in \Theta$ such that (i) $\theta^K - \theta^b < C$ and (ii) $\theta^b \geq \bar{\theta}$, for some threshold value $\bar{\theta}$.*¹⁷

This assumption captures the intuition that (i) all innovations in which a breakthrough was realized are of roughly the same value to the buyer, and (ii) reflects that all breakthroughs are of sufficiently high value. A very simple example of breakthrough innovation

¹⁷More precisely,

$$\theta^b \geq \bar{\theta} = \frac{1}{1 - (1 - F^{N(T-1)}(\theta^{b-1}))(1 - F^N(\theta^{b-1}))} \times \left(\frac{C - (M(N) - M(N-1))}{(1 - F(\theta^{b-1}))F^{N-1}(\theta^{b-1})} - F^{N(T-1)}(\theta^{b-1})\theta^{b-1} + (1 - F^{N(T-1)}(\theta^{b-1}))M(N) - NC \sum_{j=1}^{T-1} F^{Nj}(\theta^{b-1})j \right).$$

structure would be $\Theta = \{0, \varepsilon, B, B + \varepsilon\}$ for sufficiently large B and small ε . Then B and $B + \varepsilon$ are breakthrough innovations.

The sellers' 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 a seller submits an innovation to the buyer, the buyer 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 innovation are not credible.¹⁸ To overcome the hold-up problem, the buyer commits to holding a tournament and to paying a prize p to the winner of the tournament. We assume that the buyer can commit to paying a prize p_t if the tournament ends in period t . This requires that the courts can verify (i) when the tournament was declared over, and (ii) if the correct prize was paid.

In period 0 the buyer announces the tournament $\langle E, \mathbf{p}, n \rangle$ which consists of an entry fee E , a prize schedule $\mathbf{p} = [p_1, p_2, \dots, p_T]$ and the maximal number of participants n . Sellers observe $\langle E, \mathbf{p}, n \rangle$ and decide whether to pay the (possibly negative) entry fee E .¹⁹ If less than two sellers decide to participate, the tournament is canceled. If more than n sellers wish to participate, n are selected randomly. In each period $t \in \{1, \dots, T\}$ the participants of the tournament can simultaneously conduct research and subsequently decide whether or not to submit an innovation to the buyer. At the end of the period the buyer can declare any of the submissions the winner in which case the tournament ends and the seller who submitted the winning innovation receives the prize p_t .²⁰ If no winner is declared in period t or there have been no submissions, the tournament proceeds to period $t + 1$ unless it was already the last period of the tournament. In this case, all sellers submit an innovation and the buyer must declare a winner.

This model of dynamic research tournaments is essentially the same as the one proposed in Taylor (1995). We make only two changes. First, we assume that the buyer can commit to a schedule of period-specific prizes \mathbf{p} . This assumption only requires that the contracts can be time-dependent, which we view as uncontroversial. Second, to derive our main results we impose some structure on the set of innovations Θ . Namely, we assume that innovations have a breakthrough structure as discussed above and that the set Θ is

¹⁸Non-observability and non-verifiability is a typical feature of research activity. As Taylor (1995) notes "research inputs are notoriously difficult to monitor" and "courts seldom possess the ability or expertise necessary to evaluate technical research projects".

¹⁹The assumption that the buyer can charge an entry fee is taken from Taylor (1995). It is essentially an assumption that the sellers are not liquidity constrained. Alternatively, the results would remain unchanged if we assumed that the buyer cannot charge an entry fee, but instead maximizes social welfare. Many innovation tournaments are motivated exactly by social welfare and not by the private profit of the tournament organizer.

²⁰In principle, the buyer could commit to not end the tournament early. We assume that this is not possible in order to reduce the dimension of the buyer's strategy space. The assumption is without loss, as setting very high p_t essentially commits the buyer to not end the tournament in period t .

a discrete set.²¹

The tournament $\langle E, \mathbf{p}, n \rangle$ induces a T -period dynamic game of incomplete information with the set of players being the buyer and the sellers who participate in tournament. The set of players, their payoff functions, the research technology and the tournament structure are common knowledge. The seller's research activity and outcomes are private information. The timing of the game is as follows.

Period 0:

- All N invited sellers decide whether to enter or not. If they enter they pay the entry fee E .

Period $t < T$:

- Stage 1: Each seller simultaneously decides whether to perform research at cost C . Sellers do not observe the actions taken by their competitors.
- Stage 2: Each seller i who conducted research receives quality which is a random draw from F . All other sellers receive quality 0.
- Stage 3: Having privately observed the value of their innovation, sellers simultaneously decide whether to privately submit their best innovation.
- Stage 4: The buyer observes the set of submissions. If it is empty the tournament continues. If not, the buyer decides whether to declare a winner or not. If a winner is declared the tournament stops, the buyer obtains the winning innovation and the seller who submitted the winning innovation receives the prize p_t .

Period T :

- Stages 1-3: As above.
- Stage 4: The tournament stops and the buyer has to declare a winner whose submissions the buyer then obtains in exchange for the prize p_T . If no seller submitted, the prize is randomly allocated.

4 Optimal Tournament

In this section we characterize the optimal research tournament. We first show that given our assumptions on the breakthrough innovation structure, the first-best is equivalent to a

²¹The discreteness assumption is made to avoid the technicalities of defining beliefs in a dynamic game with continuous spaces. As already noted we will relax the breakthrough structure in Section 5.

global stopping rule with all N firms conducting research — that is, all firms do research in every period until at least one firm achieves a breakthrough. As soon as a breakthrough is achieved, all firms stop doing research. Next, we show that using an appropriately designed dynamic prize schedule, where the prize that the buyer awards is increasing in t , any global stopping rule can be implemented. Finally, combining these two results, we show that in the optimal tournament the buyer implements the first-best outcome. The optimal tournament features a dynamic prize structure which gives the buyer the incentives to stop the tournament as soon as at least one seller has had a breakthrough, and to the sellers the incentives to do research in every period and to truthfully report as soon as they have a breakthrough.

Proposition 1 *Consider a breakthrough innovation structure such that Assumption 1 holds. Then, in the first-best, all N firms conduct research in every period until a breakthrough is achieved after which research is stopped completely.*

In general, the first-best is characterized by a function $n(\theta, t)$, which specifies the number of sellers which optimally do research as a function of time and the current highest quality θ . Gal et al. (1981) and Morgan (1983) have shown that in the first-best there exists a global stopping quality θ^g such that $n(\theta, t) = 0$ for $\theta \geq \theta^g$ and that the first-best number $n(\theta, t)$ is decreasing in θ and increasing in t . In particular, the optimal number of sellers doing research will generally change non-monotonically over time.²² However, given our assumption about the breakthrough innovation structure, the first-best plan takes a very simple form. Since the breakthrough is very valuable, it is optimal to have all firms perform research as long as no breakthrough has been achieved. However, as all innovations with a successful breakthrough have roughly the same value, continuing the research effort is not worth its cost once a breakthrough was achieved and it is optimal to stop all research activities.

Taylor (1995) shows that any tournament with a fixed prize paid out at the end of the tournament uniquely implements an individual stopping rule, where each firm does research until it has reached its individual threshold level, irrespective of the qualities discovered by other sellers. Such an individual stopping rule consequently entails a fairly large risk of wasteful duplication across sellers. Intuitively, a global stopping rule seems better than any individual stopping rule, as it reduces the amount of wasteful research by stopping research once any seller has discovered an innovation of high enough quality. However, it is not clear if a global stopping rule can be implemented with a tournament.

²²Given the finite time-horizon there is the basic trade-off between increasing the chance of getting a high quality innovation by having many sellers do research in a given period and risking wasteful duplication. As the deadline approaches the buyer becomes more willing to risk duplication for a given quality level as there are less research opportunities in the future. Having a relatively high quality innovation early on reduces the pressure to get a better innovation in the future and the buyer is therefore less willing to risk duplication by having many sellers do research simultaneously.

Our next proposition shows that if it is possible to commit to a dynamic prize schedule, then any global stopping rule can be implemented.

Proposition 2 *Any global stopping rule $\theta^g \in \Theta$ with N sellers can be implemented using a tournament $\langle E, \mathbf{p}, N \rangle$ with a prize schedule $\mathbf{p} = [p_1, p_2, \dots, p_T]$ such that*

$$p_t = p_1 + (t - 1)\Delta(\theta^g, N),$$

where

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

for $t \in \{2, \dots, T\}$ and $p_1 \geq \bar{p}$, where \bar{p} is some cutoff value and E is sufficiently low to induce entry of all N sellers.

The proposition shows that for any $\theta^g \in \Theta$, there exists a tournament $\langle E, \mathbf{p}, N \rangle$ with a prize schedule \mathbf{p} , and an entry fee E , such that in equilibrium all N sellers enter. Furthermore, each seller performs research until it reaches quality θ^g . Once this occurs, the seller reports the discovered innovation to the buyer, who immediately stops the tournament, declares a winner and pays out the prize corresponding to the period.²³ In a fixed prize tournament as in Taylor (1995), the buyer cannot credibly commit to stop the tournament after the threshold quality has been achieved because she does not bear the marginal cost of continued research, while she stands to benefit from any marginal increase in quality. The dynamic prize schedule in Proposition 2 increases by $\Delta(\theta^g, N)$ in every period, which is exactly the marginal benefit of one additional round of research when the current highest quality is θ^g and N sellers will be conducting research in the next period. Thus, the marginal cost of an additional round equals its marginal benefit, which makes this commitment credible, as the buyer is indifferent between continuing and stopping the tournament when the current highest quality is θ^g . Since the marginal benefit of research is decreasing in θ , the buyer strictly prefers to continue the tournament whenever the highest quality is below θ^g and strictly prefers to stop it whenever it is above. Thus, it is the slope of the prize schedule that provides incentives to the buyer necessary to implement a global stopping rule.

On the other hand, the intercept of the prize schedule gives the incentives to the sellers to perform research in every period and to report their outcomes truthfully. Intuitively, as in the fixed prize tournament, each seller pursues an individual stopping rule which is determined by the expected prize. Increasing the expected prize increases the individual

²³In case of multiple breakthroughs of equal quality being reported simultaneously, the buyer randomly declares a winner among these innovations.

stopping threshold. If the individual stopping threshold is above the global stopping threshold, the sellers will conduct research as long as the tournament is ongoing. Similarly, it is the size of the prizes that induces the sellers to truthfully report their research outcomes. By not reporting an innovation above the threshold, a seller could win a higher prize in the future. However, not reporting exposes the seller to the risk that another seller will win in the current period and end the tournament. As long as the increase of the prize in the next period is sufficiently small relative to the current-period prize, the seller will report truthfully. Thus, incentives of the firms can be satisfied by making the average prize large enough by shifting up the prize schedule sufficiently. Since such a shift in the prize schedule does not affect its slope, which is the sole determinant of the buyer's incentives across periods, both buyer and seller incentives can be satisfied simultaneously, such that a global stopping rule results in equilibrium. Finally, the entry fee E can be chosen such that N sellers indeed want to participate in such a tournament by making it an entry subsidy if necessary.

Note that the Proposition 2 does not rely on the breakthrough innovation structure. That is, it holds for any vector of feasible innovation qualities Θ . In fact, the only strengthening of the assumptions from Taylor (1995) required to implement a global stopping rule instead of an individual stopping rule is that the buyer can commit to a dynamic prize schedule instead of only a fixed prize at the end of the tournament. We view this assumption as uncontroversial as time-dependent contracts are abundant.

Proposition 2 establishes that N participants would wish to take part in such a tournament and behave such that a global stopping rule is implemented. However, we still need to show if the buyer would wish to announce this tournament in the first place. The next proposition answers this question in the case of a breakthrough innovation structure.

Proposition 3 *Suppose that Assumption 1 is satisfied. Then the optimal tournament is a dynamic prize tournament and it implements the first-best.*

The proof of this result is straightforward. We know from Proposition 1 that under Assumption 1 the first-best is a global stopping rule in which all N sellers conduct research in every period until a breakthrough was had and then stop. Thus, the first-best is characterized by a constant number of sellers doing research until a global threshold is reached. However, this is precisely what Proposition 2 tells us we can achieve using an appropriate dynamic prize tournament. Moreover, with the tournament implementing the first-best they buyer would obviously want to announce it.

Taylor (1995) notes that the first-best could be achieved if the buyer, instead of holding one multi-period tournament, held a series of one-shot tournaments. In this case the buyer could choose the optimal number of sellers in each period given the current highest quality and compensate them for the one-period effort. However, if inspecting the sell-

ers' submissions is costly, running a sequence of one-period tournaments and inspecting submissions after every period may be very costly. This points to another advantage of the dynamic prize tournament. Namely, the buyer only has to inspect submissions once and only from the sellers who have developed an innovation of high enough quality. As mentioned in the introduction, we commonly distinguish between innovation races and research tournaments. Interestingly, our dynamic prize tournament acts like an innovation race as long as the tournament runs, as the first seller who makes a breakthrough wins and eventually turns into a classic research tournament with the best innovation available in the last period winning. Moreover, this innovation race is implemented without requiring verifiability, a sharp contrast to regular innovation races.

5 Robustness

In this section we show that we can implement global stopping rules more generally using dynamic prize tournaments and we compare their performance to fixed prize tournaments. Throughout this section we do not impose any structure on $\Theta = \{\theta^1, \theta^2, \dots, \theta^K\}$ except $\theta^{k+1} > \theta^k$ and the normalization $\theta^1 = 0$.

We stated in Proposition 1 that it is optimal to stop research once a breakthrough has been achieved. The next result characterizes the threshold at which research should stop in the absence of the breakthrough structure, i.e., it gives the optimal stopping rule for N sellers.

Proposition 4 *The optimal global stopping rule with N sellers is given by the smallest $\theta^k \in \Theta$ such that*

$$F^N(\theta^k) \theta^k + \sum_{j=k+1}^K (F^N(\theta^j) - F^N(\theta^{j-1})) \theta^j - \theta^k \leq NC. \quad (1)$$

When the number of sellers is not fixed the global stopping level is determined by the point at which the marginal gain of doing another round of research with exactly one seller equals its marginal cost. Essentially, there is a quality level at which another round of research with even only a single seller is not worth the cost, irrespective of how many opportunities to do research are yet to arrive. If the buyer could reduce the number of sellers further, another round of research might be optimal. However, as she has hit the lower bound, there is no room to reduce the number further and the research optimally stops. Inspecting equation (1) we see that this intuition is also present in the case where the number of sellers is fixed across time. In contrast to the case with a flexible number of sellers, the buyer compares the marginal gain of doing another round of research to its marginal cost with N sellers instead of only one. But given that the number of sellers is

fixed, there is no possibility to continue doing research with less sellers and, hence, she has hit the lower bound and research optimally stops.

In practice we observe that the number of participants in a tournament may change over time. Typically, participants are being eliminated as the end of the tournament draws closer.²⁴ Moreover, we noted earlier that the first-best may include both elimination and addition of participants over time. We showed in Proposition 2 that we can implement any global stopping rule with a fixed number of participants using a dynamic prize schedule. However, this result can be generalized considerably. The following result shows that a dynamic prize schedule allows us to implement any global stopping rule with any, arbitrarily changing number of participants. Formally, a tournament is then described by $\langle \mathbf{E}, \mathbf{p}, \mathbf{N} \rangle$, where $\mathbf{N} = [N_1, \dots, N_T]$ is the fixed number of participants in each in period. Note that the numbers N_t are fixed ex ante and do not depend on the actual research outcomes of the sellers. However, the identity of the sellers who continue can depend on their research outcomes. Given this environment, we obtain the following result.

Proposition 5 *Any global stopping rule $\theta^g \in \Theta$ with a sequence of $\mathbf{N} = [N_1, \dots, N_T]$ sellers participating in each period can be implemented using a tournament $\langle \mathbf{E}, \mathbf{p}, \mathbf{N} \rangle$ with some increasing prize schedule $\mathbf{p} = [p_1, \dots, p_T]$ and some sequence of entry fees $\mathbf{E} = [E_1, \dots, E_T]$.*

The intuition for the result is analogous to the result in Proposition 2. Namely, the slope of the prize schedule takes care of the buyer's incentives while the size of the prizes incentivizes the sellers to conduct research, truthfully reveal their innovations and to stay in the tournament for as long as possible. The difference is that the slope of the prize schedule is no longer constant but changing over time. Recall that the prizes increase from period to period to equate the marginal benefit of an additional round of research to the buyer with its marginal cost when the threshold is reached. Since the marginal benefit changes when the number of sellers is changed, the prize schedule has to increase less strongly following a round of elimination and more strongly following an addition of new sellers. Moreover, it is easier to incentivize the sellers to conduct research and report their innovations when there is a round of elimination ahead, because the threat of elimination makes truthful reporting more attractive and increases the rewards of research, as this increases the chances of not being eliminated and therefore retaining a chance at getting the prize. Similarly, the prospect of increasing the number of sellers in the next period increases the incentives to report truthfully and conduct research, as more competition in the future makes trying to win in the current period more attractive.

²⁴For example, the 2015 NRG COSIA Carbon XPRIZE consists of three rounds. Only up to 15 participants will proceed to the second round and only up to 5 five will proceed to the third round. See <http://carbon.xprize.org/about/overview>.

We noted earlier that the finite time-horizon of the tournament induces a fundamental trade-off between increasing the chance of getting a high quality innovation and risking wasteful duplication when increasing the number of sellers. What if the buyer instead had the possibility to relax the time constraint, i.e., what if the buyer could choose a later deadline $T' > T$? Intuitively, a longer time-horizon should be beneficial to the buyer, as it should allow the buyer to lessen the risk of duplication while keeping the chance of getting a high innovation. It turns out that this is not necessarily the case as our final result in this section shows.

Proposition 6 *A buyer who implements a global stopping rule is strictly better off by increasing the duration of the tournament. A buyer who implements an individual stopping rule may be worse off by increasing the duration of the tournament.*

There are two opposing effects at play when the length of the tournament is increased. First, the buyer benefits because it increases the expected quality of the innovation she will eventually obtain. Second, it increases the risk of wasteful research in the form of duplication. The beneficial effect is clearly present in both kinds of tournaments. However, in contrast to the fixed prize tournament, there is no change in the amount of duplication when a global stopping rule is implemented, as it ensures that research stops conditional on reaching the threshold. This effect is not present with an individual stopping rule. Thus, depending on which effect dominates, increasing the duration of the tournament may be harmful for the buyer under an individual stopping rule, whereas the buyer is unambiguously better off under a global stopping rule.

Moreover, suppose that Assumption 1 was satisfied and that the buyer could choose $T \in \mathbb{N}$, the ending of the tournament. By Proposition 3 we know that the optimal tournament is a dynamic prize tournament with a global stopping rule. From Proposition 6, we know that the buyer would optimally set the highest feasible T .

6 Conclusion

The goal of the present paper is to increase our understanding of the optimal design of research tournaments, which have recently seen a rapid expansion in practice. Research tournaments are inherently dynamic by the nature of research itself and by virtue of the tournament taking place over a longer period of time. It turns out that taking into account the dynamic nature of research improves tournament design. Indeed we show that using a prize schedule with increasing prizes over time instead of a fixed prize yields a strict improvement for the buyer in the face of a breakthrough structure. In particular, it gives rise to the optimal tournament and allows the buyer to implement the first-best. More generally, we show that a dynamic prize tournament with an increasing prize schedule

allows us to implement any global stopping rule. The great appeal of this finding lies in the fact that the first-best features a global stopping rule. Moreover, the channels through which this is achieved are strikingly simple: the slope of the prize schedule and its intercept align the buyer's and the sellers' incentives, respectively. Hence, an intuitive and seemingly simple departure from a fixed prize tournament delivers the ability to implement a central feature of the first-best solution.

Our results have important implications for the design of research tournaments. As mentioned before, we can usually distinguish between innovation races and research tournaments. Recall that in an innovation race the winner is the first seller to achieve a pre-determined quality, whereas in a research tournament the winner is the seller with the best quality at some pre-determined date. Naturally, both innovation races and research tournaments have their advantages and disadvantages. An innovation race avoids wasteful duplication as research stops once the goal has been reached. Moreover, it does not end until the goal is reached. However, it requires a verifiable outcome, so often an imperfect proxy of innovation has to be employed. On the other hand, a research tournament does not have to rely on proxies but there is a risk of wasteful duplication and premature ending. Quite remarkably, a dynamic prize tournament with a global stopping rule allows us to implement a hybrid of the two. Namely, throughout the duration of the tournament an innovation race is taking place, as the tournament ends as soon as an innovation above the stopping threshold is realized. However, there is an end date at which the race ends and turns into a classic research tournament in which the best innovation wins without requiring verifiability. Hence, a dynamic prize tournament allows us to get the best out of the innovation race (long horizon with no wasteful duplication) and the research tournament (no proxy required).

We derive our results under assumptions that are barely stronger than those in the seminal work by Taylor (1995). Yet, even within this rich framework there are further avenues to pursue. Comparative static results along the lines of Proposition 6 which considered extending the length of the tournament would allow us to increase our understanding of the optimal mechanism in the absence of the breakthrough innovation structure. Further, one can study the scope for beneficial information revelation between the sellers, both from the sellers' and the buyer's perspective. Moreover, although quite ambitious, deriving the optimal tournament in the absence of the breakthrough structure would be valuable and an interesting direction for future work.

References

- BIMPIKIS, K., S. EHSANI, AND M. MOSTAGIR (2014): “Designing Dynamic Contests,” Mimeo.
- 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,” *The RAND Journal of Economics*, 47, 140–165.
- CHE, Y.-K. AND I. GALE (2003): “Optimal design of research contests,” *The American Economic Review*, 93, 646–671.
- CLARK, D. J. AND C. RIIS (1996): “A multi-winner nested rent-seeking contest,” *Public Choice*, 87, 177–184.
- 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.
- FU, Q. AND J. LU (2012): “The optimal multi-stage contest,” *Economic Theory*, 51, 351–382.
- FULLERTON, R. L., B. G. LINSTER, M. MCKEE, AND S. SLATE (2002): “Using Auctions to Reward Tournament Winners: Theory and Experimental Investigations,” *The RAND Journal of Economics*, 33, 62–84.
- GAL, S., M. LANDSBERGER, AND B. LEVYKSON (1981): “A Compound Strategy for Search in the Labor Market,” *International Economic Review*, 22, 597–608.
- GERSHKOV, A. AND M. PERRY (2009): “Tournaments with midterm reviews,” *Games and Economic Behavior*, 66, 162–190.
- GREEN, B. AND C. R. TAYLOR (2015): “On Breakthroughs, Deadlines, and the Nature of Progress: Contracting for Multistage Projects,” Mimeo.
- HALAC, M., N. KARTIK, AND Q. LIU (forthcoming): “Contests for Experimentations,” *Journal of Political Economy*.
- HENDON, E., H. J. JACOBSEN, AND B. SLOTH (1996): “The One-Shot-Deviation Principle for Sequential Rationality,” *Games and Economic Behavior*, 12, 274–282.

- KLEIN, A. H. AND A. SCHMUTZLER (forthcoming): “Optimal effort incentives in dynamic tournaments,” *Games and Economic Behavior*.
- KONRAD, K. A. (2009): *Strategy and Dynamics in Contests*, Oxford University Press.
- LANG, M., C. SEEL, AND P. STRACK (2014): “Deadlines in stochastic contests,” *Journal of Mathematical Economics*, 52, 134–142.
- LETINA, I. AND A. SCHMUTZLER (2015): “Designing innovation contests for diversity,” *University of Zurich, Department of Economics, Working Paper No. 200*.
- MOLDOVANU, B. AND A. SELA (2001): “The optimal allocation of prizes in contests,” *American Economic Review*, 542–558.
- (2006): “Contest architecture,” *Journal of Economic Theory*, 126, 70–96.
- MORGAN, P. B. (1983): “Search and Optimal Sample Size,” *The Review of Economic Studies*, 50, 659–675.
- MORGAN, P. B. AND R. MANNING (1985): “Optimal Search,” *Econometrica*, 53, 923–944.
- PÉREZ-CASTRILLO, D. AND D. WETTSTEIN (forthcoming): “Discrimination in a model of contests with incomplete information about ability,” *International Economic Review*.
- 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.
- SIEGEL, R. (2009): “All-Pay Contests,” *Econometrica*, 77, 71–92.
- TAYLOR, C. R. (1995): “Digging for Golden Carrots: An Analysis of Research Tournaments,” *The American Economic Review*, 872–890.
- TERWIESCH, C. AND Y. XU (2008): “Innovation contests, open innovation, and multi-agent problem solving,” *Management science*, 54, 1529–1543.

7 Appendix

Proof of Proposition 1. The proposition is established in two steps. First, we provide conditions for the first-best solution to be such that all N firms conduct research in every period until a breakthrough is achieved after which search is stopped completely. Second, we show that if Assumption 1 holds, then the conditions identified in the first step hold as well.

Let $n(\theta, t)$ be a function which specifies the number of sellers which in the first-best do research given period t and the current highest quality is θ . Gal et al. (1981) and Morgan (1983) have shown that $n(\theta, t)$ is decreasing in θ and increasing in t . Denote with $n_\theta(\theta^k, t) = n(\theta^{k+1}, t) - n(\theta^k, t)$ and $n_t(\theta^k, t) = n(\theta^k, t+1) - n(\theta^k, t)$ the change in the optimal number of sellers due to an increase in θ or t respectively.

Step 1: We begin by showing that $n(\theta^{b-1}, 1) = N$ is a sufficient condition for $n(\theta, t) = N$ for all $\theta < \theta^b$ and all $t \geq 1$. Recall that $n_t(\theta, t) \geq 0$. Thus, we have $n(\theta^{b-1}, t) = N$ for all $t \geq 1$. Further, recall that $n_\theta(\theta, t) \leq 0$. Consequently, $n(\theta^{b-1}, 1) = N$ implies $n(\theta^s, 1) = N$ for all $s \in \{1, \dots, b-1\}$. Analogously, $n(\theta^{b-1}, T) = N$ implies $n(\theta^s, T) = N$ for all $s \in \{1, \dots, b-1\}$. Taking this together we have $n(\theta, t) = N$ for all $\theta < \theta^b$ and all $t \geq 1$.

We next show that $n(\theta^b, T) = 0$ is a sufficient condition for $n(\theta, t) = 0$ for all $\theta \geq \theta^b$ and all $t \geq 1$. It follows from $n_t(\theta, t) \geq 0$ that $n(\theta^b, t) = 0$ for all $t \geq 1$. Further, because $n_\theta(\theta, t) \leq 0$ we must have $n(\theta, t) = 0$ for all $\theta \geq \theta^b$ and all $t \geq 1$.

Step 2: Next, we show that Assumption 1 implies that the conditions identified in Step 1 hold. Note that $n(\theta^b, T) = 0$ is equivalent to

$$F(\theta^b)\theta^b + \sum_{j=b+1}^K (F(\theta^j) - F(\theta^{j-1})) \theta^j - \theta^b < C.$$

The left-hand side of the inequality is the expected benefit of conducting research with one seller in the last period given that the an innovation of quality θ^b is already available. The right-hand side is the cost of doing research with a single seller. We can rearrange the inequality to obtain

$$\begin{aligned} (F(\theta^b) - 1) \theta^b + (1 - F(\theta^b)) \theta^b + \sum_{j=b+1}^K (F(\theta^j) - F(\theta^{j-1})) (\theta^j - \theta^b) &< C \\ \sum_{j=b+1}^K (F(\theta^j) - F(\theta^{j-1})) (\theta^j - \theta^b) &< C \end{aligned}$$

This inequality is satisfied for any F if

$$(\theta^K - \theta^b) \leq C$$

which holds by Assumption 1(i).

We will now show that if Assumption 1(ii) holds, then $n(\theta^{b-1}, 1) \geq N$. Note that to demonstrate that $n(\theta^{b-1}, 1) \geq N$ it is enough to show that having N sellers conduct research is better than $N - 1$. First, we cannot have more than N sellers, and second, Morgan (1983, Proposition 2) shows that the expected benefit of conducting research within a period is concave in the number of sellers. Thus, if it is better to have N than $N - 1$ sellers, it is also better than having $N - s$ for $s \geq 1$ sellers. Let $V(\theta, t)$ denote the value of having quality θ in period t given an optimal continuation in subsequent periods. Then, we can write the expected payoff of having N sellers conduct research in period 1 given that we have quality θ^{b-1} as

$$F^N(\theta^{b-1})V(\theta^{b-1}, 2) + \theta^b(1 - F^N(\theta^{b-1})) + M(N) - NC$$

where $M(N) = \sum_{j=b+1}^K (F^N(\theta^j) - F^N(\theta^{j-1})) (\theta^j - \theta^b)$. The expected payoff of having $N - 1$ sellers conduct research is

$$F^{N-1}(\theta^{b-1})V(\theta^{b-1}, 2) + \theta^b(1 - F^{N-1}(\theta^{b-1})) + M(N - 1) - (N - 1)C.$$

Thus, having N firms is better if the difference between the two inequalities is weakly positive, which reads

$$(1 - F(\theta^{b-1}))F^{N-1}(\theta^{b-1}) [\theta^b - V(\theta^{b-1}, 2)] + (M(N) - M(N - 1)) - C \geq 0.$$

Consider now the value θ^b such that the inequality holds with equality, i.e., such that we are indifferent between N and $N - 1$ sellers. This implies that we would rather have N than $N - s$ sellers for $s \geq 2$ by the within-period concavity in the number of sellers. Moreover, since $n_t(\theta, t) \geq 0$ the θ^b which induces indifference in period 1 is such that for any period $t \geq 2$ having N firms is weakly better than having $N - 1$ firms, too. Therefore, the optimal continuation is to always employ N firms. Hence,

$$\begin{aligned} V(\theta^{b-1}, 2) = & F^{N(T-1)}(\theta^{b-1})\theta^{b-1} + (1 - F^{N(T-1)}(\theta^{b-1})) (\theta^b(1 - F^N(\theta^{b-1})) + M(N)) \\ & - NC \sum_{j=1}^{T-1} F^{Nj}(\theta^{b-1})j, \end{aligned}$$

because either we continue to have a quality of θ^{b-1} until the end, or at some point we

have a breakthrough.

Therefore,

$$\theta^b \geq \bar{\theta} = \frac{1}{1 - (1 - F^{N(T-1)}(\theta^{b-1}))(1 - F^N(\theta^{b-1}))} \times \left(\frac{C - (M(N) - M(N-1))}{(1 - F(\theta^{b-1}))F^{N-1}(\theta^{b-1})} - F^{N(T-1)}(\theta^{b-1})\theta^{b-1} + (1 - F^{N(T-1)}(\theta^{b-1}))M(N) - NC \sum_{j=1}^{T-1} F^{Nj}(\theta^{b-1})j \right)$$

is a sufficient condition on θ^b . ■

Proof of Proposition 2. The result follows as a special case (such that $N_t = N_{t+1}$ for all t) of Proposition 5. ■

Proof of Proposition 3. The result follows because the first-best is a global stopping rule with a fixed number of sellers (Proposition 1) and a dynamic fixed prize tournament can implement any global stopping rule with a fixed number of sellers (Proposition 2). ■

Proof of Proposition 4. Fix the number of sellers to N . Then any round of research yields a draw from the distribution $F^N = G$. Hence, we can reformulate the problem to one of either one seller with distribution G doing research to no research taking place at all. Then, the proof in Gal et al. (1981, p. 605) with setting $K = 1$ goes through directly proving the result. ■

Proof of Proposition 5. The tournament $\langle \mathbf{E}, \mathbf{p}, \mathbf{N} \rangle$ induces a game of incomplete information with the set of players being the buyer and the set of sellers \mathcal{N} which is given by all the sellers that are active in the tournament at some point. Abusing notation, let $|\mathcal{N}| = N$. In what follows we prove that there exists a sequential equilibrium in this game in which the sellers will conduct research in every period until they reach an innovation of quality at least θ^g , which they submit to the buyer, who will then stop the tournament. Thus, the equilibrium induces a global stopping rule. We begin by formally describing the game of incomplete information and characterize the equilibrium candidate. We then prove that this equilibrium candidate is indeed a sequential equilibrium using the one-shot deviation principle by Hendon, Jacobsen, and Sloth (1996).

The Game

The tournament $\langle \mathbf{E}, \mathbf{p}, \mathbf{N} \rangle$ induces the following extensive form game of incomplete information: $\mathcal{G} = \langle I, H, \alpha, F, (\mathcal{I}_i)_{i \in I}, (u_i)_{i \in I} \rangle$. The set of players is $I = \{B, S_1, \dots, S_N\}$. The set H is the set of histories, where the set of terminal histories is denoted Z and the actions available after the non-terminal history h is denoted $A(h) = \{a : (h, a) \in H\}$. Note that sellers who have been eliminated cannot report an innovation anymore and that sellers who have not yet been added to the tournament cannot do research. The function α assigns to each non-terminal history a member of I , i.e., α is the player function. The set of initial histories is the finite set of the states of the world Θ^{NT} . The true initial history is $\theta \in \Theta$, 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 . A seller i who conducts research in period t receives quality equal to θ_{it} . For each player $i \in I$ a partition \mathcal{I}_i of $\{h \in H : \alpha(h) = i\}$ with the property that $A(h) = A(h')$, whenever h and h' are in the same member of the partition. The function $u_i : Z \rightarrow \mathbb{R}$ maps for each player i the payoff at each of the terminal nodes. For the sellers the payoffs are determined by the research costs they have incurred, the entry fee they pay (or receive) if they enter the contest, and the prize they receive. The buyer's payoff is determined by the quality of the innovation she gets, the entry fees of the participants, and the prize she pays to the winning seller. In what follows we will use the terms doing research and investing (in research) interchangeably.

Timing

Period 0:

- All N invited sellers decide whether to enter or not. If they enter they pay the entry fee E_t according to what period they are supposed to start.

Period $t < T$:

- Stage 1: Each seller simultaneously decides whether to perform research at cost C . Sellers do not observe the actions taken by their competitors.
- Stage 2: Each seller i who conducted research receives quality equal to θ_{it} . All other sellers receive quality 0.
- Stage 3: Having privately observed the value of their innovation, sellers simultaneously decide whether to privately submit their best innovation.
- Stage 4: The buyer observes the set of submissions. If there have been no submissions, the buyer cannot end the tournament. If there have been submissions, the

buyer decides whether to declare a winner or not. If a winner is declared the tournament stops, the buyer obtains the winning innovation and the seller who submitted the winning innovation receives the prize p_t . If the tournament continues and the set of sellers has to be reduced to N_{t+1} , the buyer selects N_{t+1} sellers depending on their submission (if necessary randomly) to continue. If the tournament continues and the set of sellers has to be increased to N_{t+1} , $N_{t+1} - N_t$ sellers who paid E_{t+1} entry fee join in the next period.

Period T :

- Stages 1-3: As above.
- Stage 4: The tournament stops and the buyer has to declare a winner whose submissions the buyer then obtains in exchange for the prize p_T . If no seller submitted, the prize is randomly allocated.

Equilibrium Candidate

Denote with $\theta^i|h$ the highest quality available to player i at history h . For sellers, this is the highest quality they have so far discovered. For buyers, this is the highest quality currently submitted. The equilibrium candidate (σ, μ) is defined as follows:

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

$$\sigma^i(h) = \begin{cases} I & \text{if } \theta^i|h < \hat{\theta}(h) \\ NI & \text{else} \end{cases}, \quad (2)$$

where $\hat{\theta} : H \rightarrow \Theta$. If $A(h) = \{\text{Submit, Not Submit}\} = \{S, NS\}$ then

$$\sigma^i(h) = \begin{cases} S & \text{if } \theta^i|h \geq \theta^g \\ NS & \text{else} \end{cases}. \quad (3)$$

As we already noted, a seller who has not been added to the tournament cannot invest and a seller who has been eliminated cannot submit any innovation.

Equilibrium beliefs of seller i are as follows. Denote a history in the period t' as $h_{t'}$. Let the last period when the buyer i has not observed a deviation by the seller be $t^e|h_{t'}$. This means that in period t^e , seller i did not submit and that in all periods $t^e + 1, \dots, t' - 1$ the seller i submitted a quality over θ^g but the buyer did not end the tournament. The beliefs that the element of state of the world $\tilde{\theta}$ in some period t and for a player j , where the true state of the world is θ_{jt} are given by the following cases.

- Own elements of the state of the world ($i = j$):

$$\mu_{jt}^i(\theta^k|h_{t'}) = \begin{cases} 1 & \text{if } t \leq t', a_{it}|h_{t'} = I \text{ and } \theta^k = \theta_{jt} \\ 0 & \text{if } t \leq t', a_{it}|h_{t'} = I \text{ and } \theta^k \neq \theta_{jt} . \\ F(\theta^k) - F(\theta^{k-1}) & \text{else} \end{cases} \quad (4)$$

- Others' elements of the state of the world ($i \neq j$):

$$\mu_{jt}^i(\theta^k|h_{t'}) = \begin{cases} \frac{F(\theta^k) - F(\theta^{k-1})}{F(\theta^g - 1)} & \text{if } t \leq t^e|h_{t'} \text{ and } \theta^k < \theta^g \\ 0 & \text{if } t \leq t^e|h_{t'} \text{ and } \theta^k \geq \theta^g . \\ F(\theta^k) - F(\theta^{k-1}) & \text{else} \end{cases} \quad (5)$$

For own elements, the seller learns exact state if he invests, if he does not, or if the chance to invest has not occurred yet, he holds initial beliefs. For the others' elements, once a period starts after a no deviation from the buyer, the seller concludes that everybody has invested up to that point and that nobody has a quality higher than θ^g . This implies that each individual θ_{jt} is drawn from the truncated distribution. If the seller observes that the buyer deviated, he learns nothing about the realization of the state in that period, hence he should hold the initial beliefs. For all the states which have not been revealed yet, the seller holds initial beliefs.

Buyer Consider any information set at which the buyer is moving, i.e., there has been at least one submission. The buyer stops the game if and only if there has been a submission of quality at least θ^g . The buyer's beliefs need to be specified only when there has been at least one submission, as the buyer doesn't move if there hasn't been any. If there has been a submission, the beliefs are as follows. For any seller who has not submitted an innovation, the buyer believes that research has been conducted in every period, yet the draws were always below θ^g . For a firm which submitted an innovation, the buyer believes that research has been conducted in every period and that the submission is the currently highest quality.

There is no profitable one-shot deviation for the buyer

First, observe that the buyer cannot stop the tournament in any period if no firm submits an innovation. Thus we only need to consider cases where at least one firm has submitted an innovation. Further, whenever the buyer declares a winner, she will always choose the highest submission. Moreover, in period T the tournament ends in any case and consequently the buyer simply declares the highest quality innovation the winner.

Period $T - 1$

Let the expected highest quality after one additional round of research by M firms, given that the current highest quality is $\theta^k \in \Theta$, be given by a function $R : \Theta \times \mathbb{N} \rightarrow [\theta^1, \theta^K]$. That is

$$R(\theta^k, M) = F(\theta^k)^M \theta^k + \sum_{j=k+1}^K (F(\theta^j)^M - F(\theta^{j-1})^M) \theta^j.$$

Suppose the sellers play the candidate equilibrium strategy. Consider the incentives of the buyer in the period $T - 1$, when the highest quality is $\theta^k < \theta^g$. Stopping the tournament results in the payoff $\theta^k - p_{T-1}$ while continuing the tournament yields her $R(\theta^k, N_T) - p_T$. The buyer continues the tournament if and only if

$$R(\theta^k, N_T) - \theta^k > p_T - p_{T-1}.$$

Next, we show that the LHS of the above inequality is strictly decreasing in θ^k .

$$\begin{aligned} & R(\theta^{k+1}, N_T) - \theta^{k+1} - (R(\theta^k, N_T) - \theta^k) \\ &= F(\theta^{k+1})^{N_T} \theta^{k+1} + \sum_{j=k+2}^K (F(\theta^j)^{N_T} - F(\theta^{j-1})^{N_T}) \theta^j - \theta^{k+1} \\ &\quad - F(\theta^k)^{N_T} \theta^k - \sum_{j=k+1}^K (F(\theta^j)^{N_T} - F(\theta^{j-1})^{N_T}) \theta^j + \theta^k \\ &= F(\theta^{k+1})^{N_T} \theta^{k+1} - \theta^{k+1} - F(\theta^k)^{N_T} \theta^k - (F(\theta^{k+1})^{N_T} - F(\theta^k)^{N_T}) \theta^{k+1} + \theta^k \\ &= F(\theta^k)^{N_T} (\theta^{k+1} - \theta^k) - (\theta^{k+1} - \theta^k) \\ &= (\theta^{k+1} - \theta^k) (F(\theta^k)^{N_T} - 1) \\ &< 0. \end{aligned}$$

The LHS is strictly decreasing in θ^k and since $R(\theta^g, N_T) - \theta^g = p_T - p_{T-1}$ by construction, the buyer has no incentive to stop the tournament if $\theta^k < \theta^g$. If $\theta^k > \theta^g$ and all firms do research in the period T it still holds $R(\theta^k, N_T) - \theta^k < p_T - p_{T-1}$, hence the buyer does not want to continue the tournament and, therefore, the buyer stops the tournament in the period $T - 1$ if and only if the highest quality is $\theta^k \geq \theta^g$.

If a round of elimination is ahead and there have been more than N_{t+1} submissions, the buyer obviously chooses the best of them to continue. If not, the buyer chooses randomly. If the buyer has to increase the number of sellers, she chooses $N_{t+1} - N_t$ sellers from the set of previously inactive sellers who paid E_{t+1} entry fee, all which only have worthless innovations to begin with.

Period $t \leq T - 2$

Step 1. Denote with $V_t(\mu, \pi | \theta^k)$ the value to the buyer of having the highest quality

θ^k in period t given that she follows the equilibrium candidate. Then, it follows that $V_t(\mu, \pi|\theta^g) = \theta^g - p_t$. In this step, we will show that also $V_t(\mu, \pi|\theta^g) = F(\theta^g)^{N_{t+1}}V_{t+1}(\mu, \pi|\theta^g) + \sum_{j=g+1}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}})V_{t+1}(\mu, \pi|\theta^j)$ for any t . It suffices to show that

$$\theta^g - p_t = F(\theta^g)^{N_{t+1}}V_{t+1}(\mu, \pi|\theta^g) + \sum_{j=g+1}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}})V_{t+1}(\mu, \pi|\theta^j).$$

Obviously, for any $\theta^j \geq \theta^g$ it holds $V_{t+1}(\mu, \pi|\theta^j) = \theta^j - p_{t+1}$. Substituting

$$\theta^g - p_t = F(\theta^g)^{N_{t+1}}(\theta^g - p_{t+1}) + \sum_{j=g+1}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}})(\theta^j - p_{t+1})$$

$$p_{t+1} - p_t = F(\theta^g)^{N_{t+1}}\theta^g + \sum_{j=g+1}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}})\theta^j - \theta^g$$

which holds by definition for any t .

Step 2. In this step we show that for any pair θ^k, θ^{k+1} such that $\theta^{k+1} \leq \theta^g$ it holds that

$$V_t(\mu, \pi|\theta^{k+1}) - V_t(\mu, \pi|\theta^k) = F(\theta^k)^{N_{t+1}}(V_{t+1}(\mu, \pi|\theta^{k+1}) - V_{t+1}(\mu, \pi|\theta^k)).$$

We have

$$V_t(\mu, \pi|\theta^k) = F(\theta^k)^{N_{t+1}}V_{t+1}(\mu, \pi|\theta^k) + \sum_{j=k+1}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}})V_{t+1}(\mu, \pi|\theta^j)$$

By Step 1, an equivalent expression holds for $V_t(\mu, \pi|\theta^{k+1})$. Expanding $V_t(\mu, \pi|\theta^{k+1}) - V_t(\mu, \pi|\theta^k)$ we get:

$$\begin{aligned} & V_t(\mu, \pi|\theta^{k+1}) - V_t(\mu, \pi|\theta^k) \\ &= F(\theta^{k+1})^{N_{t+1}}V_{t+1}(\mu, \pi|\theta^{k+1}) + \sum_{j=k+2}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}})V_{t+1}(\mu, \pi|\theta^j) \\ &\quad - F(\theta^k)^{N_{t+1}}V_{t+1}(\mu, \pi|\theta^k) - \sum_{j=k+1}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}})V_{t+1}(\mu, \pi|\theta^j) \\ &= F(\theta^k)^{N_{t+1}}(V_{t+1}(\mu, \pi|\theta^{k+1}) - V_{t+1}(\mu, \pi|\theta^k)) \end{aligned}$$

Step 3. In this step we show that for any $t \leq T-2$ and any $\theta^k < \theta^g$, the buyer does not stop the tournament, i.e., there is no profitable one-shot deviation. Stopping the tournament in period t yields the payoff of $\theta^k - p_t$. Thus, it suffices to show that $V_t(\mu, \pi|\theta^k) > \theta^k - p_t$

for any $\theta^k < \theta^g$. Observe that $V_t(\mu, \pi|\theta^g) - \theta^g + p_t = 0$. We will show that $V_t(\mu, \pi|\theta^k) - \theta^k$ is strictly decreasing in θ^k for any $\theta^k < \theta^g$. The result then follows. This is equivalent to

$$V_t(\mu, \pi|\theta^{k+1}) - \theta^{k+1} - (V_t(\mu, \pi|\theta^k) - \theta^k) < 0,$$

where $\theta^{k+1} \leq \theta^g$. By Step 2 we can write

$$F(\theta^k)^{N_{t+1}}(V_{t+1}(\mu, \pi|\theta^{k+1}) - V_{t+1}(\mu, \pi|\theta^k)) - (\theta^{k+1} - \theta^k) < 0$$

Observe that $V_T(\theta^{k+1}) - V_T(\theta^k) = \theta^{k+1} - \theta^k$. Iterating Step 2 we get

$$F(\theta^k)^{N_{t+1}}(V_{t+1}(\theta^{k+1}) - V_{t+1}(\theta^k)) < (\theta^{k+1} - \theta^k)$$

and thus a one-shot deviation is not profitable.

Step 4. In this step we show that the buyer stops the tournament whenever $\theta^k > \theta^g$, i.e. we show that there is no profitable one-shot deviation in this case either. Consider a quality of $\theta^k \geq \theta^g$ in period t and suppose the buyer does not stop the tournament. She will for sure stop the tournament in period $t + 1$, however, as we consider only a one-shot deviation. The buyer stops the tournament if

$$\theta^k - p_t > F(\theta^k)^{N_{t+1}}\theta^k + \sum_{j=k+1}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}}) \theta^j - p_{t+1}$$

or, equivalently, if

$$p_{t+1} - p_t > F(\theta^k)^{N_{t+1}}\theta^k + \sum_{j=k+1}^K (F(\theta^j)^{N_{t+1}} - F(\theta^{j-1})^{N_{t+1}}) \theta^j - \theta^k.$$

From our period T analysis we know that the RHS is strictly decreasing in θ^k , and by construction it is equal to the LHS for $\theta^k = \theta^g$. Hence whenever $\theta^k > \theta^g$, the inequality holds and the buyer stops the tournament.

Step 5. If a round of elimination is ahead and there have been more than N_{t+1} submissions, the buyer obviously chooses the best of them to continue. If not, the buyer chooses randomly. If the buyer has to increase the number of sellers, she chooses $N_{t+1} - N_t$ sellers from the set of previously inactive sellers who paid E_{t+1} entry fee, all which only have worthless innovations to begin with.

There is no profitable one-shot deviation at the research stage for the seller

Let π' be a strategy profile that coincides with the equilibrium candidate π with the exception of the seller i 's action in the investment stage in period t . Thus, it is a one-shot deviation. Recall that

$$p_t = p_1 + \sum_{i=2}^t \Delta(\theta^g, N_i),$$

where

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

In what follows, we show two things. First, for a sufficiently high p_1 there exists a $\theta \in \Theta$ with $\theta^g \leq \theta < \theta^K$ such that investing is optimal for all $\theta^k \leq \theta$. Second, for each p_1 there exists a $\theta \in \Theta$ with $\theta^g \leq \theta < \theta^K$ such that not investing is optimal for all $\theta^k > \theta$. Notice that in any period t the seller could win an amount $z_t \in Z_t = \{0, p_t, p_t/2, \dots, p_t/N_t\}$. Here, winning zero amounts to the tournament ending in that period. Thus, the probability that the tournament continues is given by

$$\left(1 - \sum_{z \in Z_t} P_t^{\mu, \pi}(z|\theta^k)\right) =: W(\mu, \pi|t, \theta^k).$$

Further, let $P_t^{\mu, \pi}(z|\theta^k)$ denote the probability of winning some $z \in Z_t$. We can then define the seller's expected winnings in period t as

$$Z(\mu, \pi|\theta^k, t) = p_t \mathcal{P}^{\mu, \pi}(p_t|\theta^k),$$

where

$$\mathcal{P}^{\mu, \pi}(p_t|\theta^k) = \sum_{n=0}^{N_t-1} \frac{P_t^{\mu, \pi}(p_t/(n+1)|\theta^k)}{n+1}$$

denotes the probability of winning in period t conditional on reaching period t with quality θ^k . Notice that we have

$$\mathcal{P}^{\mu, \pi}(p_t|\theta^i) = \mathcal{P}^{\mu, \pi}(p_t|\theta^j), \quad \forall \theta^i, \theta^j < \theta^g; t < T.$$

Further, for all $\theta^k < \theta^g$ and $t < s < T$ (i.e., any subsequent period after the deviation except the last one) we have

$$\mathcal{P}^{\mu, \pi}(p_s | \theta^k) = \mathcal{P}^{\mu, \pi'}(p_s | \theta^k).$$

In period t , when the deviation takes place, we have

$$\mathcal{P}^{\mu, \pi}(p_t | \theta^k) > \mathcal{P}^{\mu, \pi'}(p_t | \theta^k) = 0$$

for any $\theta^k < \theta^g$ and in the final period we have

$$\mathcal{P}^{\mu, \pi}(p_T | \theta^k) \geq \mathcal{P}^{\mu, \pi'}(p_T | \theta^k)$$

for any $\theta^k < \theta^g$. Further, notice that the investment costs incurred up to any period $s \geq t$ is greater under π than under π' by exactly C , the investment cost from period t . Note that the probability of winning in period $t < T$ does not depend on the quality at the beginning of the period (since we consider one-shot deviations all qualities will be below the threshold), because all that matters is whether a quality above the threshold is drawn in period t . Thus, we can drop the quality and write $\mathcal{P}^{\mu, \pi}(p_t | \theta^k) = \mathcal{P}^{\mu, \pi}(p_t)$, but we need to keep in mind that this is only without loss for $t < T$, as the current quality matters for winning in the last period, too.

Moreover, $I_i(\mu, \pi | \theta^k, t) \in \{0, 1\}$ denotes the seller's investment decision. With this in hand we can define the expected utility of the assessment (μ, π) in period t when having a highest quality θ^k by

$$U_i(\mu, \pi | \theta^k, t) = Z(\mu, \pi | \theta^k, t) - I_i(\mu, \pi | \theta^k, t)C + W(\mu, \pi | t, \theta^k) \tilde{U}_i(\mu, \pi | t + 1),$$

where

$$\begin{aligned} & \tilde{U}_i(\mu, \pi | t + 1) \\ &= I_i(\mu, \pi | \theta^k, t) \left(F(\theta^k) U_i(\mu, \pi | \theta^k, t + 1) + \sum_{x=k+1}^{g-1} (F(\theta^x) - F(\theta^{x-1})) U_i(\mu, \pi | \theta^x, t + 1) \right) \\ &+ (1 - I_i(\mu, \pi | \theta^k, t)) U_i(\mu, \pi | \theta^k, t + 1). \end{aligned}$$

Note that the history is not an argument of the expected utility $U_i(\mu, \pi | \theta^k, t)$. The reason is that any investment cost incurred up to period t are sunk and do not matter for the decision in period t and only the highest quality is relevant, but not how (i.e., at what point in the past) the seller got it. Moreover, the continuation utility $\tilde{U}_i(\mu, \pi | t + 1)$ is different depending on whether or not the seller invests in period t . In case of investment

the seller could start period t with a highest quality ranging from θ^k to θ^{g-1} , but not above, as for any quality above θ^g she would have submitted it in ended the tournament.

Part 1: Investment is optimal up to some θ

We begin by showing that for a sufficiently high p_1 there exists a $\theta \in \Theta$ with $\theta^g \leq \theta < \theta^K$ such that investing is optimal for all $\theta^k \leq \theta$. Thus, we need to show that

$$U_i(\mu, \pi|\theta^k, t) - U_i(\mu, \pi'|\theta^k, t) \geq 0 \quad (6)$$

for all $\theta^k \leq \theta$ for sufficiently high p_1 . We will do so by showing that inequality (6) is satisfied for any $\theta^k < \theta^K$ for sufficiently high p_1 . We can rewrite the LHS of inequality (6) to

$$\begin{aligned} & U_i(\mu, \pi|\theta^k, t) - U_i(\mu, \pi'|\theta^k, t) = \\ & -C + p_t(\mathcal{P}^{\mu, \pi}(p_t) - \mathcal{P}^{\mu, \pi'}(p_t)) \\ & + W(\mu, \pi|t, \theta^k) \left(F(\theta^k)U_i(\mu, \pi|\theta^k, t+1) + \sum_{x=k+1}^{g-1} (F(\theta^x) - F(\theta^{x-1}))U_i(\mu, \pi|\theta^x, t+1) \right) \\ & - W(\mu, \pi'|t, \theta^k)U_i(\mu, \pi|\theta^k, t+1). \end{aligned}$$

Note the following about the term $U_i(\mu, \pi|\theta^k, t+1)$ at the very end. This is the continuation value of the one-shot deviation π' . Yet, the argument in the function is the equilibrium candidate π . The reason for this is that in period $t+1$ the only thing that matters is the strategy (and belief) for periods after and including $t+1$ (and there π and π' coincide) and all that matters from past periods is the highest quality. We need to consider two cases. First the case when the currently highest quality is below θ^g , and second the case when the currently highest quality is at least θ^g .

Case 1 Suppose $\theta^k < \theta^g$. Begin by considering period T . The LHS of inequality (6) then simplifies to

$$\begin{aligned} & U_i(\mu, \pi|\theta^k, T) - U_i(\mu, \pi'|\theta^k, T) = \\ & -C + p_T(\mathcal{P}^{\mu, \pi}(p_T|\theta^k) - \mathcal{P}^{\mu, \pi'}(p_T|\theta^k)) \end{aligned}$$

because the tournament ends for sure. Notice that $\mathcal{P}^{\mu, \pi}(p_T|\theta^k) - \mathcal{P}^{\mu, \pi'}(p_T|\theta^k) > 0$. To see this, note that we are comparing the probabilities of winning the tournament (in a tie or outright) for the case of no research in the last period with the case of another round of research in the last period. Clearly, if the seller invests there is a strictly higher chance of winning. Thus, inequality (6) is satisfied for sufficiently large p_1 .

Now consider any period $t < T$. We can write the expected utility of the equilibrium strategy as

$$p_t \mathcal{P}^{\mu, \pi}(p_t) - C + \left(\sum_{s=t+1}^T (-C + p_s \mathcal{P}^{\mu, \pi}(p_s)) \prod_{i=t}^{s-1} F^{N_i}(\theta^{g-1}) Q^{\mu, \pi}(t) \right)$$

and the utility of the one-shot deviation as

$$\sum_{s=t+1}^T \left(-C + p_s \mathcal{P}^{\mu, \pi'}(p_s) \right) F^{-1}(\theta^{g-1}) \prod_{i=t}^{s-1} F^{N_i}(\theta^{g-1}) Q^{\mu, \pi'}(t),$$

where $Q^{\mu, \pi}(t)$ is the probability that the seller will still be in the tournament given that it continues. This captures the risk of being eliminated. We have $Q^{\mu, \pi'}(t) \leq Q^{\mu, \pi}(t)$ because not investing yields a lower expected innovation and therefore a higher chance of being eliminated.

Recall that we can write prizes as $p_t = p_1 + \sum_{i=2}^T \Delta(\theta^g, N_i)$. Making use of this we can write the expected utility of the equilibrium strategy as

$$\begin{aligned} & p_t \mathcal{P}^{\mu, \pi}(p_t) - C + \left(\sum_{s=t+1}^T (-C + p_s \mathcal{P}^{\mu, \pi}(p_s)) \prod_{i=t}^{s-1} F^{N_i}(\theta^{g-1}) Q^{\mu, \pi}(t) \right) \\ &= p_1 \mathcal{P}^{\mu, \pi}(p_t) + p_1 \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi}(p_s) \prod_{i=t}^{s-1} F^{N_i}(\theta^{g-1}) Q^{\mu, \pi}(t) \right) + K \end{aligned}$$

where K contains all the terms that are not a function of p_1 .

We can proceed accordingly for the expected utility of the one-shot deviation. Then the difference between the equilibrium candidate and the one-shot deviation corresponds exactly to the LHS of inequality (6). Dropping all the terms that do not depend on p_1 this difference then reads

$$\begin{aligned} & p_1 \left(\mathcal{P}^{\mu, \pi}(p_t) + \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi}(p_s) F^{N_t}(\theta^{g-1}) Q^{\mu, \pi}(t) \prod_{i=t+1}^{s-1} F^{N_i}(\theta^{g-1}) Q^{\mu, \pi}(t) \right) \right) \\ & - p_1 \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi'}(p_s) F^{N_t-1}(\theta^{g-1}) Q^{\mu, \pi'}(t) \prod_{i=t+1}^{s-1} F^{N_i}(\theta^{g-1}) Q^{\mu, \pi'}(t) \right) \end{aligned}$$

Recall that $\mathcal{P}^{\mu, \pi}(p_s) = \mathcal{P}^{\mu, \pi'}(p_s)$ for all $t < s < T$ and $\mathcal{P}^{\mu, \pi}(p_T) \geq \mathcal{P}^{\mu, \pi'}(p_T)$ and that $Q^{\mu, \pi'}(t) \leq Q^{\mu, \pi}(t)$. Thus, for the inequality to be satisfied we need

$$\mathcal{P}^{\mu, \pi}(p_t) + \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi}(p_s) F^{N_t}(\theta^{g-1}) Q^{\mu, \pi}(t) \prod_{i=t+1}^{s-1} F^{N_i}(\theta^{g-1}) Q^{\mu, \pi}(t) \right)$$

$$\begin{aligned}
& - \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi'}(p_s) F^{N_t-1}(\theta^{g-1}) Q^{\mu, \pi'}(t) \prod_{i=t+1}^{s-1} F^{N_t}(\theta^{g-1}) Q^{\mu, \pi'}(t) \right) \\
& \geq \mathcal{P}^{\mu, \pi}(p_t) + \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi'}(p_s) F^{N_t}(\theta^{g-1}) Q^{\mu, \pi'}(t) \prod_{i=t+1}^{s-1} F^{N_t}(\theta^{g-1}) Q^{\mu, \pi'}(t) \right) \\
& - \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi'}(p_s) F^{N_t-1}(\theta^{g-1}) Q^{\mu, \pi'}(t) \prod_{i=t+1}^{s-1} F^{N_t}(\theta^{g-1}) Q^{\mu, \pi'}(t) \right) \\
& = \mathcal{P}^{\mu, \pi}(p_t) + F^{N_t}(\theta^{g-1}) \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi'}(p_s) Q^{\mu, \pi'}(t) \prod_{i=t+1}^{s-1} F^{N_t}(\theta^{g-1}) Q^{\mu, \pi'}(t) \right) \\
& - F^{N_t-1}(\theta^{g-1}) \left(\sum_{s=t+1}^T \mathcal{P}^{\mu, \pi'}(p_s) Q^{\mu, \pi'}(t) \prod_{i=t+1}^{s-1} F^{N_t}(\theta^{g-1}) Q^{\mu, \pi'}(t) \right) \\
& = \mathcal{P}^{\mu, \pi}(p_t) - F^{N_t-1}(\theta^{g-1})(1 - F(\theta^{g-1}))R \\
& \geq 0
\end{aligned}$$

to be positive. Note that that $R \leq 1$ as R is the total probability of winning the tournament (either tied or outright). Thus, what remains to be shown in order for the inequality (6) to be satisfied is that

$$\mathcal{P}^{\mu, \pi}(p_t) + F^{N_t-1}(\theta^{g-1})(F(\theta^{g-1}) - 1) \geq 0.$$

We can write

$$\mathcal{P}^{\mu, \pi}(p_t) = \sum_{k=g}^K (F(\theta^k) - F(\theta^{k-1})) \sum_{n=0}^{N_t-1} \binom{N_t-1}{n} \frac{F^{N_t-1-n}(\theta^{k-1}) (F(\theta^k) - F(\theta^{k-1}))^n}{n+1}.$$

The lowest quality with which the seller can win in period $t < T$ is θ^g , which she gets with a probability $F(\theta^g) - F(\theta^{g-1})$. In addition, to win outright, all others must draw at most θ^{g-1} . The probability of winning when tying with one is

$$\binom{N_t-1}{1} \frac{F^{N_t-2}(\theta^{k-1}) (F(\theta^k) - F(\theta^{k-1}))}{2}$$

as $N_t - 2$ sellers must draw at most θ^{g-1} and exactly one must draw θ^g , too, and there are $N_t - 1$ different ways of getting there and then there is a $1/2$ chance of winning the tie.

Proving that inequality (6) is satisfied then boils down to proving

$$F^{N_t-1}(\theta^{g-1})(F(\theta^{g-1}) - 1)$$

$$+ \sum_{k=g}^K (F(\theta^k) - F(\theta^{k-1})) \sum_{n=0}^{N-1} \binom{N-1}{n} \frac{F^{N-1-n}(\theta^{k-1}) (F(\theta^k) - F(\theta^{k-1}))^n}{n+1} \geq 0.$$

To see that this is indeed the case some steps are needed.

$$\begin{aligned} & F^{N_t-1}(\theta^{g-1})(F(\theta^{g-1}) - 1) \\ & + \sum_{k=g}^K (F(\theta^k) - F(\theta^{k-1})) \sum_{n=0}^{N_t-1} \binom{N_t-1}{n} \frac{F^{N_t-1-n}(\theta^{k-1}) (F(\theta^k) - F(\theta^{k-1}))^n}{n+1} \\ & \geq F^{N_t-1}(\theta^{g-1})(F(\theta^{g-1}) - 1) + \sum_{k=g}^K (F(\theta^k) - F(\theta^{k-1})) F^{N_t-1}(\theta^{k-1}) \\ & \geq F^{N_t-1}(\theta^{g-1})(F(\theta^{g-1}) - 1) + F^{N_t-1}(\theta^{g-1}) \sum_{k=g}^K (F(\theta^k) - F(\theta^{k-1})) \\ & = F^{N_t-1}(\theta^{g-1})(F(\theta^{g-1}) - 1) + F^{N_t-1}(\theta^{g-1})(1 - F(\theta^{g-1})) \\ & = 0. \end{aligned}$$

The first inequality follows by considering only the first element of the binomial sum (that is, only the outright wins).

Case 2 Suppose $\theta^g \leq \theta^k < \theta^K$. In this case, the game will end with certainty in this period and the LHS of inequality (6) reads

$$\begin{aligned} & U_i(\mu, \pi|\theta^k, t) - U_i(\mu, \pi'|\theta^k, t) = \\ & - C + (p_1 + \sum_{i=2}^T \Delta(\theta^g, N_i)) (\mathcal{P}^{\mu, \pi}(p_t) - \mathcal{P}^{\mu, \pi'}(p_t)). \end{aligned}$$

Thus, for a sufficiently high p_1 , not investing is not a profitable deviation.

Part 2: Investment is not optimal above some θ

We will now show that for each p_1 there exists a $\theta \in \Theta$ with $\theta^g \leq \theta < \theta^K$ such that not investing is optimal for all $\theta^k > \theta$. It is obvious that for a quality of θ^K it is never optimal to invest, as research is costly. So there will always exist a quality level above which not investing is optimal. What we remains to be shown is that if there is a $\theta < \theta^K$ such that not investing is optimal at θ , then for all $\theta^k \geq \theta$ not investing is optimal, too. We do this by showing that whenever it is optimal not to invest for θ^k , then it is also optimal not to invest for θ^{k+1} . The proof is then completed by induction.

Consider some $\theta^k \geq \theta \geq \theta^g$. Abusing notation, let $Z(\mu, \pi|\theta^k, t)$ denote the expected winnings *after* the investment stage. Thus, we can write the expected utility of having

highest quality θ^k and investing in that period as

$$F(\theta^k)Z(\mu, \pi|\theta^k, t) + \sum_{m=k+1}^K (F(\theta^m) - F(\theta^{m-1}))Z(\mu, \pi|\theta^m, t) - C$$

while the expected utility of having highest quality θ^k and not investing in that period is

$$Z(\mu, \pi|\theta^k, t).$$

By assumption we have

$$Z(\mu, \pi|\theta^k, t) - \left(F(\theta^k)Z(\mu, \pi|\theta^k, t) + \sum_{m=k+1}^K (F(\theta^m) - F(\theta^{m-1}))Z(\mu, \pi|\theta^m, t) - C \right) \geq 0.$$

We now show that this implies

$$Z(\mu, \pi|\theta^{k+1}, t) - \left(F(\theta^{k+1})Z(\mu, \pi|\theta^{k+1}, t) + \sum_{m=k+2}^K (F(\theta^m) - F(\theta^{m-1}))Z(\mu, \pi|\theta^m, t) - C \right) \geq 0.$$

We have

$$\begin{aligned} & Z(\mu, \pi|\theta^{k+1}, t) - \left(F(\theta^{k+1})Z(\mu, \pi|\theta^{k+1}, t) + \sum_{m=k+2}^K (F(\theta^m) - F(\theta^{m-1}))Z(\mu, \pi|\theta^m, t) - C \right) \\ & - Z(\mu, \pi|\theta^k, t) + \left(F(\theta^k)Z(\mu, \pi|\theta^k, t) + \sum_{m=k+1}^K (F(\theta^m) - F(\theta^{m-1}))Z(\mu, \pi|\theta^m, t) - C \right) \\ & = Z(\mu, \pi|\theta^{k+1}, t) - Z(\mu, \pi|\theta^k, t) - F(\theta^{k+1})Z(\mu, \pi|\theta^{k+1}, t) \\ & + F(\theta^k)Z(\mu, \pi|\theta^k, t) + (F(\theta^{k+1}) - F(\theta^k))Z(\mu, \pi|\theta^{k+1}, t) \\ & = (1 - F(\theta^k))(Z(\mu, \pi|\theta^{k+1}, t) - Z(\mu, \pi|\theta^k, t)) \\ & \geq 0 \end{aligned}$$

and therefore not investing is optimal for θ^{k+1} given that not investing was optimal at θ^k .

There is no profitable one-shot deviation at the submission stage for the seller

Let π' denote the one-shot deviation of seller i in period t at the submission stage. Observe that submitting an innovation that has quality below θ^g is never profitable, unless a round of elimination is ahead. Then, any seller should submit their highest innovation as this increases the chance of being allowed to continue. If no round of elimination is ahead, submitting below θ^g is not profitable. Thus we only need to consider the decision of a seller who has an innovation of quality $\theta^k \geq \theta^g$. Let $S_i(\mu, \pi|\theta^k, t) \in \{0, 1\}$ denote seller

i 's submission decision in period t with highest quality θ^k . Given that we consider the submission stage with $\theta^k \geq \theta^g$ we can write utility as follows

$$U_i(\mu, \pi|\theta^k, t) = S_i(\mu, \pi|\theta^k, t)Z(\mu, \pi|\theta^k, t) + (1 - S_i(\mu, \pi|\theta^k, t))Q^{\mu, \pi}(t)\tilde{U}_i(\mu, \pi|t + 1),$$

where

$$Z(\mu, \pi|\theta^k, t) = p_t \sum_{n=0}^{N-1} \frac{P_t^{\mu, \pi}(p_t/(n+1)|\theta^k)}{n+1},$$

and

$$\begin{aligned} \tilde{U}_i(\mu, \pi|t + 1) &= I_i(\mu, \pi|\theta^k, t + 1) \left(F(\theta^k)U_i(\mu, \pi|\theta^k, t + 1) + \sum_{x=k+1}^K (F(\theta^x) - F(\theta^{x-1}))U_i(\mu, \pi|\theta^x, t + 1) - C \right) \\ &\quad + (1 - I_i(\mu, \pi|\theta^k, t + 1))U_i(\mu, \pi|\theta^k, t + 1). \end{aligned}$$

The term $P_t^{\mu, \pi}(z|\theta^k)$ captures the probability of winning prize t in period t when the seller's highest quality is θ^k after the investment stage in t . Thus, $Z(\mu, \pi|\theta^k, t)$ captures the expected winning after the investment stage. Moreover, the cost of investment in period t is already sunk at the submission stage and does not show up. The term $\tilde{U}_i(\mu, \pi|t + 1)$ corresponds to the continuation value the seller receives when she does not submit. Recall that submitting will end the game for sure, as the buyer is playing according to the equilibriums strategy and the seller has a quality $\theta^k \geq \theta^g$. Even if the seller does not submit, the game may still end if another seller submitted a sufficiently high quality or because the seller is eliminated. Thus we weight the continuation value by that probability that the contest indeed continues to the next period and denote this $Q^{\mu, \pi}(t)$. The continuation value itself differs depending on whether the seller will invest in period t or not.

Submitting is profitable if

$$U_i(\mu, \pi|\theta^k, t) - U_i(\mu, \pi'|\theta^k, t) \geq 0. \quad (7)$$

Step 1. Suppose θ^k is sufficiently high so that the seller would not invest in the following period. We can then rewrite the left-hand side of equation (7) to

$$\begin{aligned} &U_i(\mu, \pi|\theta^k, t) - U_i(\mu, \pi'|\theta^k, t) \\ &= Z(\mu, \pi|\theta^k, t) - Q^{\mu, \pi'}(t)Z(\mu, \pi|\theta^k, t + 1) \\ &= p_1 \sum_{n=0}^{N-1} \frac{\left[P_t^{\mu, \pi}(p_t/(n+1)|\theta^k) - Q^{\mu, \pi'}(t)P_{t+1}^{\mu, \pi'}(p_{t+1}/(n+1)|\theta^k) \right]}{n+1} + R_1 - R_2, \end{aligned}$$

where the terms R_1 and R_2 do not depend on p_1 . Further, the term in the brackets in the sum is strictly positive. Hence, increasing p_1 increases the difference $U_i(\mu, \pi|\theta^k, t) - U_i(\mu, \pi'|\theta^k, t)$ and thus, for sufficiently large p_1 , the inequality is satisfied.

Step 2. Suppose θ^k is sufficiently low so that the seller would still invest in the following period. If the seller submits, she will get any prize $z_t \in Z_t = \{0, p_t, p_t/2, \dots, p_t/N_t\}$ and the game will end for sure, as she is submitting a quality $\theta^k \geq \theta^g$. Suppose the state of the world is such that she would receive a price $z_t \in Z_t \setminus \{p_t\}$, i.e., she will not win outright. Not submitting would mean that the contest ends and she receives no prize, as there is (at least) one other seller with a quality $\theta \geq \theta^k \geq \theta^g$. Since we have $z_t \geq 0$ for all $z_t \in Z_t \setminus \{p_t\}$, submitting yields a weakly higher payoff in those states of the world than not submitting. Suppose next that the state of the world is such that the highest valuation among the other sellers is θ and we have $\theta^g \leq \theta < \theta^k$. Thus, the seller would win outright if she submits and if she does not submit she receives nothing and the tournament ends. Thus, submitting yields a strictly higher payoff than not submitting. Finally, suppose the state of the world is such that the highest quality among the other seller is θ and we have $\theta < \theta^g$. Thus, submitting would yield a prize of p_t while not submitting would yield an expected winning of

$$Q^{\mu, \pi'}(t)Z(\mu, \pi|\theta^k, t+1)$$

if she does not invest in the subsequent period or

$$Q^{\mu, \pi'}(t) \left(F(\theta^k)Z(\mu, \pi|\theta^k, t+1) + \sum_{x=k+1}^K (F(\theta^x) - F(\theta^{x-1}))Z(\mu, \pi|\theta^x, t+1) \right)$$

if she submits. Given that she chooses optimally between submitting and not submitting, submitting would yield Z^* , which denotes the maximum of the two above expressions. Essentially, in this state of the world the seller is trading-off winning the higher prize p_{t+1} in the next period with a probability below 1 against winning p_t for sure. Thus, for a sufficiently high prize p_1 she will always choose to submit. ■

Proof of Proposition 6. The proof proceeds in two steps. In the Step 1 we show that extending T is always beneficial for the buyer in case of a dynamic prize tournament. In Step 2 we show by example that extending T may be harmful for the buyer in case of a fixed prize dynamic tournament.

Step 1: The expected costs in a dynamic prize tournament, implementing a global stop-

ping rule θ^g in a T -period tournament are given by

$$EK^g(\theta^g, N, T) = \left(\sum_{t=1}^{T-1} tF^{(t-1)N}(\theta^{g-1})(1 - F^N(\theta^{g-1})) + TF^{(T-1)N}(\theta^{g-1}) \right) NC.$$

Thus, the marginal cost of extending the tournament to $T + 1$ periods is

$$\begin{aligned} EK^g(\theta^g, N, T + 1) - EK^g(\theta^g, N, T) &= \\ &= \left(\sum_{t=1}^T tF^{(t-1)N}(\theta^{g-1})(1 - F^N(\theta^{g-1})) + (T + 1)F^{TN}(\theta^{g-1}) \right) NC \\ &\quad - \left(\sum_{t=1}^{T-1} tF^{(t-1)N}(\theta^{g-1})(1 - F^N(\theta^{g-1})) + TF^{(T-1)N}(\theta^{g-1}) \right) NC \\ &= (-TF^{(T-1)N}(\theta^{g-1})(F^N(\theta^{g-1})) + (T + 1)F^{TN}(\theta^{g-1})) NC \\ &= F^{TN}(\theta^{g-1})NC \end{aligned}$$

and expected quality by

$$EQ^g(\theta^g, N, T) = \sum_{k=1}^K \theta^k h^g(\theta^k | \theta^g, N, T).$$

where

$$h^g(\theta^k | \theta^g, N, T) = \begin{cases} F^{NT}(\theta^k) - F^{NT}(\theta^{k-1}) & k < g, \\ \sum_{t=1}^T F^{N(t-1)}(\theta^{g-1}) (F^N(\theta^k) - F^N(\theta^{k-1})) & k \geq g. \end{cases}$$

Then, the marginal benefit of extending the tournament to $T + 1$ periods is

$$\begin{aligned} EQ^g(\theta^g, N, T + 1) - EQ^g(\theta^g, N, T) &= \\ &= \sum_{k=1}^K \theta^k h^g(\theta^k | \theta^g, N, T + 1) - \sum_{k=1}^K \theta^k h^g(\theta^k | \theta^g, N, T) \\ &= \sum_{k=1}^K \theta^k (h^g(\theta^k | \theta^g, N, T + 1) - h^g(\theta^k | \theta^g, N, T)) \\ &= \sum_{k=1}^{g-1} \theta^k (F^{NT}(\theta^k)(F^N(\theta^k) - 1) - F^{NT}(\theta^{k-1})(F^N(\theta^{k-1}) - 1)) \\ &\quad + \sum_{k=g}^K \theta^k F^{NT}(\theta^{g-1}) (F^N(\theta^k) - F^N(\theta^{k-1})) \end{aligned} \tag{8}$$

The seller benefits from extending the tournament if

$$F^{TN}(\theta^{g-1})NC \leq (EQ^g(\theta^g, N, T+1) - EQ^g(\theta^g, N, T)).$$

From the optimality of θ^g we know that

$$NC \leq \sum_{j=g+1}^K \theta^j (F^N(\theta^j) - F^N(\theta^{j-1})) - \theta^g (1 - F^N(\theta^g))$$

Thus, to show that the seller benefits from extending the tournament, it is sufficient to show that

$$\begin{aligned} F^{TN}(\theta^{g-1}) \left(\sum_{j=g+1}^K \theta^j (F^N(\theta^j) - F^N(\theta^{j-1})) - \theta^g (1 - F^N(\theta^g)) \right) \\ \leq (EQ^g(\theta^g, N, T+1) - EQ^g(\theta^g, N, T)). \end{aligned} \quad (9)$$

Combining (8) and (9) and simplifying, we get that the sufficient condition is

$$F^{TN}(\theta^{g-1})\theta^g(1 - F^N(\theta^{g-1})) \geq \sum_{k=1}^{g-1} \theta^k (F^{NT}(\theta^k)(1 - F^N(\theta^k)) - F^{NT}(\theta^{k-1})(1 - F^N(\theta^{k-1})))$$

Recall that $F(\theta^0) = 0$ and note that we can write this sum as

$$\begin{aligned} & \sum_{k=1}^{g-1} \theta^k (F^{NT}(\theta^k)(1 - F^N(\theta^k)) - F^{NT}(\theta^{k-1})(1 - F^N(\theta^{k-1}))) \\ &= \theta^1 F^{NT}(\theta^1)(1 - F^N(\theta^1)) - \theta^2 F^{NT}(\theta^1)(1 - F^N(\theta^1)) \\ & \quad + \theta^2 F^{NT}(\theta^2)(1 - F^N(\theta^2)) - \theta^3 F^{NT}(\theta^2)(1 - F^N(\theta^2)) \\ & \quad \vdots \\ & \quad + \theta^{g-2} F^{NT}(\theta^{g-2})(1 - F^N(\theta^{g-2})) - \theta^{g-1} F^{NT}(\theta^{g-2})(1 - F^N(\theta^{g-2})) \\ & \quad + \theta^{g-1} F^{NT}(\theta^{g-1})(1 - F^N(\theta^{g-1})) \\ &= - \sum_{k=1}^{g-2} (\theta^{k+1} - \theta^k) F^{NT}(\theta^k)(1 - F^N(\theta^k)) + \theta^{g-1} F^{NT}(\theta^{g-1})(1 - F^N(\theta^{g-1})) \end{aligned}$$

This allows us to rewrite the sufficient condition to

$$(\theta^g - \theta^{g-1})F^{TN}(\theta^{g-1})(1 - F^N(\theta^{g-1})) \geq - \sum_{k=1}^{g-2} (\theta^{k+1} - \theta^k) F^{NT}(\theta^k)(1 - F^N(\theta^k))$$

which always holds because $\theta^{k+1} > \theta^k$.

Step 2: To construct an example of a harmful extension of T in case of a fixed prize tournament we consider a setting with $N = 2$, $\Theta = \{0, 1\}$ and extend $T = 2$ to $T + 1 = 3$. Let the probability of drawing $\theta^1 = 1$ be given by π . We will choose parameters such that the optimal individual threshold is $\theta^i = 1$. The expected costs in the case of $T = 2$ are given by

$$EK(2, 2) = 2(\pi C + 2(1 - \pi)C)$$

and the expected quality

$$EQ(2, 2) = 1 - (1 - \pi)^4.$$

In the case of $T = 3$ we have

$$EK(2, 3) = 2(\pi C + 2(1 - \pi)\pi C + 3(1 - \pi)^2 C)$$

and the expected quality

$$EQ(2, 3) = 1 - (1 - \pi)^6.$$

Hence, the change in expected surplus for the buyer is given by

$$EQ(2, 2) - EK(2, 2) - EQ(2, 3) + EK(2, 3) = (1 - \pi)^2 (2C - (2 - \pi)\pi(1 - \pi)^2)$$

which is negative for $C = 1/10$ and $\pi = 2/5$. Since

$$EQ(2, 2) - EK(2, 2) = \frac{16}{25} - \frac{8}{25} > 0$$

the individual threshold is $\theta^i = 1$ is indeed optimal. ■