

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/81600>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Contests for Experimentation*

Marina Halac[†]

Navin Kartik[‡]

Qingmin Liu[§]

February 19, 2016

Abstract

We study contests for innovation with learning about the innovation’s feasibility and opponents’ outcomes. We characterize contests that maximize innovation when the designer chooses a prize-sharing scheme and a disclosure policy. A “public winner-takes-all contest” dominates public contests—where any success is immediately disclosed—with any other prize-sharing scheme as well as winner-takes-all contests with any other disclosure policy. Yet, jointly modifying prize sharing and disclosure can increase innovation. In a broad class of mechanisms, it is optimal to share the prize with disclosure following a certain number of successes; under simple conditions, a “hidden equal-sharing” contest is optimal.

*We thank Martin Cripps, Florian Ederer, Alex Frankel, Qiang Fu, Kobi Glazer, Ed Green, Yingni Guo, Hugo Hopenhayn, Johannes Hörner, Ilan Kremer, Jingfeng Lu, Sofia Moroni, Pauli Murto, Motty Perry, Sven Rady, Ilya Segal, Francesco Squintani, Juuso Välimäki, Julia Wirtz, Tom Wiseman, and various seminar and conference audiences for helpful comments. We appreciate Curtis Taylor’s valuable discussion of the paper. We have also profited from the thoughtful feedback of a Co-Editor and anonymous referees. Andrew Kosenko and Teck Yong Tan provided excellent research assistance. Kartik gratefully acknowledges financial support from the NSF.

[†]Columbia University and University of Warwick. Email: mhalac@columbia.edu.

[‡]Columbia University. Email: nkartik@columbia.edu.

[§]Columbia University. Email: qingmin.liu@columbia.edu.

1. Introduction

Contests or prize awards are practical and proven mechanisms to procure innovations. Used since at least the 18th century, their popularity has surged in recent decades (McKinsey & Company, 2009). The internet television company Netflix generated significant buzz in 2006 by announcing a \$1 million prize to induce a 10% improvement in the accuracy of its movie recommendation algorithm. There is an ongoing \$30 million Google Lunar X Prize for landing a private spacecraft on the surface of the Moon and sending “Mooncasts” back to Earth. In the public sector, President Barack Obama signed the America COMPETES Reauthorization Act in 2011 to grant U.S. government agencies the authority to conduct contests to spur innovation. There have also been renewed calls to reform the patent system by using prizes to avoid the deadweight losses of monopoly power (e.g. Stiglitz, 2006).¹

This paper studies the design of contests for specific innovations. Previous work on contest design has focused on settings in which there is no uncertainty about the environment. By contrast, we emphasize endogenous learning about the desired innovation. We are motivated by applications in which the viability or feasibility of the innovation is uncertain at the outset. Agents update their beliefs over time through their own *experimentation*—exerting costly effort and observing their outcomes—and also based on what they learn about other agents’ outcomes.

Our model builds on the workhorse exponential-bandit framework (Keller, Rady, and Cripps, 2005). A principal and a set of ex-ante homogenous agents (or contestants) are initially uncertain whether an innovation is feasible or not. If the innovation is feasible, an agent’s instantaneous probability of obtaining the innovation—hereafter, synonymous with “a success”—depends on the agent’s effort. If the innovation is not feasible, success cannot obtain. At each instant of time, each agent chooses how much effort to covertly exert. Whether an agent succeeds or not is only directly observed by that agent and the principal, not by any other agent. All parties are risk neutral.

The principal chooses a prize and, given that prize, a contest design that maximizes innovation, viz., the probability of obtaining one success. Contest design consists of two

¹Debates about the merits of patents versus prizes (versus grants) to encourage innovation date back to at least the 19th century (Machlup and Penrose, 1950, particularly pp. 19–20). In 2011, the U.S. Senator Bernie Sanders proposed two bills that together would create innovation prize funds of 0.57% of U.S. GDP—over \$80 billion at the time—for targeted medical research; the bills have not yet been put to vote in Congress. In some domains—e.g., the development of new antibiotics—prizes are advocated not to mitigate the deadweight loss from monopoly power, but rather because the lure of market exclusivity has been insufficient to promote research (Emanuel, 2015).

instruments. First, the principal chooses a *prize-sharing scheme*, which specifies how the prize will be divided among successful agents as a function of when each agent succeeds. For example, a “winner-takes-all” contest awards the entire prize to the first agent who succeeds, whereas an “equal-sharing” contest gives every agent who succeeds by a deadline an equal share of the prize. Second, the principal chooses a *disclosure policy*, which specifies what information she discloses over time about agents’ outcomes. For example, a “public” contest reveals to all agents, at every point of time, whether each agent has succeeded or not, whereas a “hidden” contest does not reveal any information until the end of the contest.

In light of the agents’ risk neutrality and the principal valuing only one success, an intuitive solution to the design problem is to use a *public winner-takes-all* contest. After all, sharing the prize in any other fashion lowers a contestant’s expected reward from success, which should depress effort incentives. Not disclosing success immediately would lead contestants to fear that another contestant may have already succeeded, which should also lower incentives to exert effort. Consistent with this intuition, we show that a public winner-takes-all contest dominates a public contest with any other sharing scheme as well as a winner-takes-all contest with any other disclosure policy.

However, we find that it is possible to increase innovation by *jointly* modifying both the prize-sharing scheme and the disclosure policy: a *hidden equal-sharing* contest can dominate a public winner-takes-all contest. The intuition turns on a tradeoff that arises from the uncertainty in the environment. On the one hand, the principal wants to increase each agent’s expected reward from success; this familiar force pushes in favor of using a winner-takes-all contest. On the other hand, the principal also wants to buttress agents’ beliefs about the innovation’s feasibility; this force, which owes entirely to learning, pushes in favor of hiding information—specifically, not disclosing the lack of success by other agents. Crucially, though, the gains from hiding information can only be harnessed by also sharing the prize.

We explain the above intuition using a two-period example in [Section 2](#), before presenting the paper’s main model in [Section 3](#). Our central results on optimal contest design are in [Section 4](#), with our analysis restricted to symmetric equilibria.

Among contests with public and hidden information disclosure, the optimal contest is either public winner-takes-all or hidden equal-sharing. We provide intuitive conditions for when one of these contests dominates the other; for example, hidden equal-sharing is preferred if the value of the innovation is large enough. More broadly, we prove that among contests with rank-monotonic prize schemes—those in which an agent

is rewarded no less than another agent who succeeds later—and arbitrary deterministic and symmetric information disclosure policies, it is optimal to use a *cutoff-disclosure equal-sharing* contest. A cutoff-disclosure policy consists of the principal making an announcement as soon as a critical number of agents have succeeded while remaining silent otherwise. The intuitive conditions that rank public winner-takes-all and hidden equal-sharing contests are also sufficient for optimality of these contests among all contests in the class we study. While we are unable to solve for the optimal contest in full generality, and the principal may be able to do better with contests we do not study, we believe our results contain useful economic lessons and qualify the conventional presumption in favor of winner-takes-all contests.

A notable implication of our results is that the principal may obtain the innovation with higher probability when there are more agents in the contest, despite our model abstracting away from any exogenous forces (e.g., heterogeneity) that favor having multiple agents. Having more agents can be (second-best) optimal because it allows the principal to better harness the benefits from hiding information and sharing the prize.

[Section 5](#) addresses the principal’s optimal choice of the prize and discusses some extensions of the model. Among other things, we show that our main results remain valid under alternative observability structures, such as when only the principal or only an agent observes success directly, and when agents are allowed to communicate their successes to each other. The insight that a hidden equal-sharing contest can dominate a public winner-takes-all contest also applies to a designer who internalizes effort costs, so long as the value of the innovation is larger than the maximum prize available.

The economic forces and insights uncovered in our analysis of contest design are relevant to other contexts in which agents work on projects of uncertain quality or feasibility. [Section 6](#) draws implications for first-to-file versus first-to-invent rules in patent law, optimal task allocation in organizations, and the design of contract awards in government procurement. Although none of these settings fits our setup perfectly, together they illustrate the breadth of applications of our insights on the benefits of limiting disclosure and sharing prizes.

Related literature

A subset of the prior work on contest design concerns research contests rather than innovation contests. The distinction is articulated well by [Taylor \(1995, p. 874\)](#): “in a research tournament, the terminal date is fixed, and the quality of innovations varies, while in an innovation race, the quality standard is fixed, and the date of discovery is variable.” The research-contest literature includes both static ([Fullerton and McAfee, 1999](#);

Moldovanu and Sela, 2001; Che and Gale, 2003) and dynamic models (Taylor, 1995). Krishna and Morgan (1998) study a setting where the principal has a fixed budget.

There is a sizable literature on different aspects of innovation or patent races, pioneered by Loury (1979) and Dasgupta and Stiglitz (1980). The focus in this literature is typically only on a winner-takes-all structure and most of it is without learning. Design questions are addressed, for example, by Bhattacharya, Glazer, and Sappington (1990), Moscarini and Smith (2011) and Judd, Schmedders, and Yeltekin (2012); see also the references therein.

Our paper is more closely related to work on innovation contests with learning. Choi (1991), Malueg and Tsutsui (1997), Mason and Välimäki (2010), and Moscarini and Squintani (2010) focus on winner-takes-all contests rather than contest design. Choi (1991) considers a multi-stage innovation process and notes that learning about a competitor’s success has a “positive effect” of making an agent more optimistic about the return to his own effort. In our setting, agents do not learn about opponents’ outcomes in a hidden contest; rather, the benefit obtains from an agent’s *conjecture* about opponents’ success when the prize is shared. In concurrent work, Bimpikis, Ehsani, and Mostagir (2014) compare certain information disclosure policies and reward schemes for the first stage of a two-stage contest. Without characterizing optimal contests, they show that the principal can benefit from hiding information in the first stage.² Due to discounting and the presence of the second stage, there is social value to having multiple agents succeed in the first stage, whereas we identify (second-best) reasons to use a contest that induces multiple successes even when only the first success is socially valuable. Moroni (2015) studies optimal contracts with multiple agents and stages; see Section 5.1 for a connection with our results.³

More broadly, the exponential-bandit learning framework we follow is now widely used to study multi-agent strategic experimentation (e.g., Keller, Rady, and Cripps, 2005; Keller and Rady, 2010; Bonatti and Hörner, 2011; Murto and Välimäki, 2011; Cripps and Thomas, 2014), as an alternative to the brownian-motion formulation of Bolton and Harris (1999). Some authors have analyzed how partial information disclosure or strategic communication about experimentation outcomes can improve aggregate learning; see Bimpikis and Drakopoulos (2014), Che and Hörner (2015), Heidhues, Rady, and Strack

²Their second stage has no learning and has a public winner-takes-all structure. The authors consider three regimes for the first stage: public winner-takes-all, hidden information with no (direct) rewards, and a structure with disclosure only after a given time.

³Contracting with a single agent in the exponential-bandit framework has been studied by Bergemann and Hege (1998, 2005), Gomes, Gottlieb, and Maestri (2015), Hörner and Samuelson (2014), and Halac, Kartik, and Liu (2016).

(2015), and also [Kremer, Mansour, and Perry \(2014\)](#) in a non-exponential-bandit framework. All these papers consider a fixed payoff-interdependence structure. Our work stresses the importance of jointly designing both information disclosure and payoff interdependence, the latter determined by the prize-sharing scheme.

Finally, how much information a principal should disclose about agents' outcomes has been tackled in other contexts. Feedback in multi-stage contests without learning is studied by [Aoyagi \(2010\)](#), [Ederer \(2010\)](#), [Goltsman and Mukherjee \(2011\)](#), and [Wirtz \(2013\)](#). [Yildirim \(2005\)](#), [Gill \(2008\)](#), [Rieck \(2010\)](#), and [Akcigit and Liu \(2015\)](#) address the incentives for contestants to themselves disclose their outcomes to opponents, an issue we take up in [Section 5](#); [Campbell, Ederer, and Spinnewijn \(2014\)](#) consider related matters in a moral-hazard-in-teams setting. [Manso \(2011\)](#) discusses how much feedback a single agent should be provided about his own experimentation outcomes, and [Ely \(2015\)](#) studies dynamic information disclosure about an exogenous state variable.

2. The Main Idea

This section explains the intuition for our results in a simplified example. A principal wants to obtain a specific innovation. The innovation's feasibility depends on a binary state—either good or bad—that is persistent and unobservable. The prior probability of the good state is $p_0 \in (0, 1)$. There are two periods, $t = 0, 1$, no discounting, and two risk-neutral agents. In each period each agent covertly chooses whether to work or shirk. If an agent works in a period and the state is good, the agent succeeds in that period with probability $\lambda \in (0, 1)$; if either the agent shirks or the state is bad, the agent does not succeed.⁴ Working in a period costs an agent $c > 0$. Successes are conditionally independent across agents given the state and are observed only by the principal and the agent who succeeds. The principal wants to induce both agents to work until at least one succeeds; an additional success provides no extra benefit. The principal has a prize \bar{w} to pay the agents.

In this illustrative setting, we consider four contests. They vary by whether the entire prize is allocated to the first successful agent or divided equally among all agents who succeed by the end of the second period, and by whether an agent's success in the first period is publicly disclosed or kept hidden.

⁴Without loss, an agent who succeeds in the first period cannot—more precisely, will not—exert effort in the second period. Whenever we refer to an agent working in the second period, it is implicitly in the event that he has not succeeded in the first period.

Public winner-takes-all. Suppose the principal awards the entire prize \bar{w} to the first agent who obtains a success, and she publicly discloses all successes at the end of each period. If both agents succeed simultaneously, the prize is equally divided (or allocated to either agent with equal probability). In this mechanism, neither agent will work in the second period if either succeeded in the first period. Thus, in any period, if there has been no earlier success and the opponent is exerting effort, an agent's expected reward for success is $\hat{w} := \lambda\bar{w}/2 + (1 - \lambda)\bar{w}$. If $p_0\lambda\hat{w} > c$, it is a dominant strategy for an agent to work in the first period; assume for the rest of this section that this condition holds. If neither agent succeeds in the first period, both agents work in the second period if and only if

$$p_1\lambda\hat{w} \geq c, \tag{1}$$

where $p_1 := \frac{p_0(1-\lambda)^2}{p_0(1-\lambda)^2 + 1 - p_0}$ is the agents' belief in the second period that the state is good given that neither succeeded in the first period having exerted effort.

Public equal-sharing. Suppose the principal discloses all successes at the end of each period, but she now divides the prize \bar{w} equally between the two agents if both succeed by the end of the second period no matter their order of success (while still allocating the entire prize to an agent if he is the only one to succeed). If an agent succeeds in the first period, the opponent is certain in the second period that the state is good and, due to the shared-prize scheme, the opponent's reward for success is $\bar{w}/2$. Thus, when $\lambda\bar{w}/2 < c$, an agent does not work in the second period if his opponent succeeds in the first period; in this case, the contest is equivalent to public winner-takes-all (WTA). On the other hand, when $\lambda\bar{w}/2 > c$, an agent will work in the second period if the opponent succeeds in the first period. This "duplication effort" does not benefit the principal because she values only one success; moreover, as compared to public WTA, agents' incentives to work in the first period can now be lower due to two reasons: free-riding—an agent may want to wait for the other agent to experiment and reveal information about the state—and a lower expected reward for first-period success due to the opponent's duplication effort.

In any case, observe that if both agents work in the first period and neither succeeds, the incentive constraint in the second period is still given by (1). Therefore, a public equal-sharing (ES) contest cannot improve on public WTA in the sense that the former cannot induce effort by both agents in both periods when the latter cannot.⁵

⁵ An artifact of this two-period, two-agent example is that whenever a public WTA contest induces effort by both agents in both periods, so will a public ES contest. This equivalence is not a general property.

Hidden winner-takes-all. Suppose the principal awards the entire prize \bar{w} to the first successful agent (and splits the prize in case of simultaneous success) but now she does not disclose any information about first-period successes until the end of the game. Plainly, an agent works in the first period if he is willing to work in the second period. However, because first-period successes are hidden, an agent's second-period decision must now take into account the possibility that the opponent may have already succeeded and secured the entire prize. When both agents work in the first period, an agent i 's incentive constraint for effort in the second period if he did not succeed in the first is

$$\Pr[j \text{ failed} \mid i \text{ failed}]p_1\lambda\hat{w} \geq c, \quad (2)$$

where j denotes i 's opponent. Clearly, constraint (2) is more demanding than (1). In other words, a hidden WTA contest cannot improve on public WTA; moreover, for some parameters, a public WTA contest improves on hidden WTA (see [Appendix A.7](#)).

Hidden equal-sharing. Now suppose the principal combines the equal-sharing prize scheme described earlier with not disclosing any information about first-period successes until the end of the game. Although the prize is shared, there is now no free-riding concern: an agent cannot learn from his opponent's success when that is not disclosed. In fact, because there is nothing to be learned about the opponent after the first period, it is without loss that an agent works in the first period if he works at all.⁶

Suppose there is an equilibrium in which both agents work in the first period and consider an agent's incentive to work in the second period if he did not succeed in the first. The agent does not know whether the opponent succeeded or failed in the first period. In the event that the opponent succeeded, the agent's posterior belief that the state is good is one while his reward for success becomes half the prize. On the other hand, in the event the other agent failed in the first period, the agent's posterior belief is $p_1 < 1$ but the expected reward for success is $\hat{w} > \bar{w}/2$. Hence, an agent i 's incentive constraint in the second period is

$$\Pr[j \text{ succeeded} \mid i \text{ failed}]\lambda\frac{\bar{w}}{2} + \Pr[j \text{ failed} \mid i \text{ failed}]p_1\lambda\hat{w} \geq c. \quad (3)$$

It can be checked (see [Appendix A.7](#)) that there are parameters under which inequality (3) holds while (1) does not. In other words, there are parameters under which both agents

⁶“Without loss” in the sense that if there is an equilibrium in which an agent works in the second period (and possibly in the first period too), then there is an outcome-equivalent equilibrium in which this agent works in the first period (and possibly in the second period too).

work in both periods in a hidden ES contest but not in a public WTA contest. For these parameters, a hidden ES contest improves on public WTA.

What is the intuition behind why hidden ES can improve on public WTA although neither public ES nor hidden WTA can? On the one hand, holding fixed an agent’s belief about the state, it is clear that a WTA prize scheme maximizes effort incentives. On the other hand, the nature of learning—specifically, failure is bad news—implies that the principal would like to hide information about the opponent’s first-period outcome to bolster an agent’s second-period belief in the only event that matters to the principal, viz. when the opponent fails in the first period. Hiding information but still using WTA is counter-productive, however, because when an agent conditions on obtaining a reward in the second period, he deduces that the opponent must have failed. Consequently, harnessing the benefits of hiding information requires some sharing of the prize. On the flip side, sharing the prize while maintaining public disclosure is not beneficial either because this change from public WTA only alters second-period incentives when the principal does not value additional effort, viz. when the innovation has obtained in the first period.

It bears emphasis that public WTA would be an optimal contest were it certain that the state is good: if there is no learning, there is no benefit to hiding information. Public WTA would also be optimal if agents did learn but only from their own outcomes and not from others’, i.e. if their experimentation “arms” were uncorrelated because their approaches to innovation had no connection. If the principal were to value obtaining a success by each agent, then a public ES contest would sometimes be optimal—sharing the prize has an obvious benefit when success by multiple agents is desired. What is surprising is that ES can be optimal, when paired with hidden information, despite the principal valuing only one success.

These intuitions in hand, we turn to a more general analysis.

3. The Model

3.1. Environment

A principal wants to obtain a specific innovation. Whether the innovation is feasible depends on the state of nature, $\theta \in \{G, B\}$, where G represents “good” and B represents “bad”. This state is persistent and unobservable to all parties. There are $N \geq 1$ agents who can work on the principal’s project. Time is continuous and runs from 0 up to some end date $T \in \mathbb{R}_+$, which is chosen by the principal. At every moment $t \in [0, T]$, each agent $i \in \mathcal{N} := \{1, 2, \dots, N\}$ covertly chooses effort $a_{i,t} \in [0, 1]$ at instantaneous cost $ca_{i,t}$, where

$c > 0$. Denote $A_t := \sum_i a_{i,t}$. If $\theta = G$ and agent i exerts effort $a_{i,t}$ at time t , he succeeds with instantaneous probability $\lambda a_{i,t}$ at t , where $\lambda > 0$ is a commonly known parameter. No success can be obtained if $\theta = B$. Successes are conditionally independent given the state. We assume that successes are observable only to the agent who succeeds and to the principal; [Section 5](#) discusses alternative scenarios.

When a success is obtained, the principal receives a lump-sum payoff $v > 0$; the agents do not intrinsically care about success. The principal values only one success: additional successes have no social value. All parties are risk neutral, have quasi-linear preferences, and are expected-utility maximizers. To make our analysis and insights more transparent, we assume no discounting.⁷

Let $p_0 \in (0, 1)$ be the commonly-known prior probability that the state is good. Assume the ex-ante expected marginal benefit of effort is larger than the marginal cost: $p_0 \lambda v > c$. This means that some experimentation is efficient, even though conditional on the bad state the marginal benefit of effort is zero. Denote by p_t the posterior probability that the state is good when no agent has succeeded by time t given a (measurable) effort profile $\{a_{i,t}\}_{i,t}$. We refer to p_t as the *public belief*. By Bayes' rule,

$$p_t = \frac{p_0 e^{-\int_0^t \lambda A_z dz}}{p_0 e^{-\int_0^t \lambda A_z dz} + 1 - p_0}, \quad (4)$$

which is decreasing over time.

Let $a_i^t := \int_0^t a_{i,z} dz$ be the cumulative effort, or experimentation, by agent i up to time t conditional on him not having succeeded by t , and $A^t := \sum_i a_i^t = \int_0^t A_z dz$ the *aggregate cumulative effort up to t* given no success by t . The (aggregate) *probability of success* is

$$p_0 \left(1 - e^{-\lambda A^t}\right). \quad (5)$$

It follows from (4) that (5) is equivalent to $p_0 - (1 - p_0) \frac{p_T}{1 - p_T}$. Hence, as is intuitive:

Remark 1. For any set of parameters, the probability of success is increasing in aggregate cumulative effort, A^T . Moreover, for any prior belief p_0 , a lower public belief at the deadline, p_T , corresponds to a higher probability of success.

⁷ As elaborated in [Section 5](#), our finding that an optimal public WTA contest can be dominated by a hidden ES contest is robust to discounting.

3.2. First best

Since it is socially optimal to never exert effort after a success has been obtained, the social optimum is derived by maximizing

$$\int_0^\infty (p_t \lambda v - c) A_t e^{-\int_0^t p_z \lambda A_z dz} dt.$$

To interpret this expression, note that $e^{-\int_0^t p_z \lambda A_z dz}$ is the probability that no success is obtained by time t , and p_t is the probability that the state is good given no success by t . Conditional on the good state, a success then occurs at t with instantaneous probability λA_t , yielding a value v . Since the public belief p_t is decreasing over time, a social-welfare maximizing effort profile is $a_{i,t} = 1$ for all $i \in \mathcal{N}$ if $p_t \lambda v \geq c$, and $a_{i,t} = 0$ for all $i \in \mathcal{N}$ otherwise. The first-best *stopping (posterior) belief* is given by

$$p^{FB} := \frac{c}{\lambda v}. \quad (6)$$

3.3. Contests and strategies

The principal designs a mechanism to incentivize the agents. As is common, we endow the principal with commitment power and impose limited liability for the agents (i.e., each agent must receive a non-negative payment). In general, a mechanism specifies a deadline $T \geq 0$ and a vector of payments $(w_1, \dots, w_N) \in \mathbb{R}_+^N$ that, without loss, are made at T as a function of the principal's information at T (when each agent succeeded, if ever). In addition, a mechanism specifies an information disclosure policy: a distribution over messages for each agent at each time, which may depend on the principal's information at that time.

We are interested in a sub-class of mechanisms, which we call contests, that restrict both payments and disclosure policies. With regards to payments, a contest specifies a prize amount, simply a prize hereafter, $\bar{w} > 0$ and how that prize is allocated to the agents. Let s_i denote the time at which agent i succeeds; by convention, $s_i = \emptyset$ if i does not succeed, and we take $\emptyset > T$. The *prize-sharing scheme* is specified by a tuple of functions $(w_i(\mathbf{s}))_{i \in \mathcal{N}}$, where $w_i(\mathbf{s})$ is the payment to agent i when the vector of success times is \mathbf{s} .⁸ We require the scheme to satisfy three properties: (i) for all i , $w_i(\mathbf{s}) = w(s_i, \mathbf{s}_{-i})$, with $w(s_i, \mathbf{s}_{-i}) = w(s_i, \sigma(\mathbf{s}_{-i}))$ for any permutation σ ; (ii) $w(\emptyset, \cdot) = 0$; and (iii) $\mathbf{s} \neq (\emptyset, \dots, \emptyset) \implies \sum_{i=1}^N w_i(\mathbf{s}) = \bar{w}$.

⁸We use bold symbols to denote vectors. Note that an agent is paid the same regardless of whether he succeeds once or more than once.

Requirement (i) says that the prize-sharing scheme must be anonymous.⁹ An agent’s payment can be interpreted as an expected payment, in which case anonymity implies ex-ante symmetry while permitting ex-post asymmetries via randomization, as elaborated below. Requirement (ii) says that an agent who does not succeed is paid zero; as usual under limited liability, this is without loss of generality. Lastly, requirement (iii) says that the principal must pay out the entire prize if at least one success is obtained.¹⁰ This requirement is a natural property of contests—they are quintessentially about prize allocation—and is consistent with prize awards observed in the real world. Notwithstanding, [Section 5.1](#) describes how our main insights can be extended without this restriction.

As foreshadowed in [Section 2](#), two simple prize-sharing schemes will be of particular interest. A *winner-takes-all* (WTA) contest is one in which the first agent to succeed receives the entire prize: for all i , $w_i(\mathbf{s}) = \bar{w}$ if $i = \arg \min_{j \in \mathcal{N}} s_j$ and $w_i(\mathbf{s}) = 0$ otherwise.¹¹ An *equal-sharing* (ES) contest shares the prize equally among all successful agents: for all i , $w_i(\mathbf{s}) = \frac{\bar{w}}{|\{j \in \mathcal{N}: s_j \neq \emptyset\}|}$ if $s_i \neq \emptyset$ and $w_i(\mathbf{s}) = 0$ otherwise. An ES contest can equivalently be implemented by a lottery that gives each successful agent an equal probability of receiving the entire prize.

Some notation is needed to define information disclosure policies. Let $o_{i,t} = 1$ if agent i succeeds at time t and $o_{i,t} = 0$ otherwise, $\mathbf{o}_t = (o_{1,t}, \dots, o_{N,t})$ be a profile of outcomes at time t , and $\mathbf{o}^t = (\mathbf{o}_z)_{z < t}$ be the history of outcome profiles up to time t . Let O^t be the set of all possible histories of outcome profiles up to time t . A *disclosure policy* is a sequence $(M_t, \mu_t)_{t \in [0, T]}$, where each M_t is a measurable message space for each time t , and each $\mu_t : O^t \rightarrow M_t$ is a measurable function capturing the time t disclosure rule.¹² At each time t , all agents observe a common message $m_t = \mu_t(\mathbf{o}^t) \in M_t$. Strictly speaking, this setup precludes disclosure policies that treat agents differently depending on their history of

⁹ Relaxing this requirement would only broaden the set of parameters for which an optimal public WTA contest is dominated by a contest that shares the prize and hides information; see [fn. 21](#).

¹⁰ This requirement is without loss when a contest designer takes the prize as given and simply maximizes the probability of obtaining the innovation. This is arguably the appropriate objective for contest designers when contests are funded by third parties, as is increasingly frequent ([McKinsey & Company, 2009](#)). An interesting parallel is [Maskin’s \(2002\)](#) formulation of the UK government choosing an optimal mechanism to minimize total pollution, subject to the government having a fixed budget to spend on the task. Another interpretation of requirement (iii) is that the prize for success is not monetary. For example, contestants’ incentives often stem from factors such as the publicity received from recognition by the contest (cf. [MacCormack, Murray, and Wagner, 2013](#), p. 27). Contest design can control how the publicity is allocated among successful contestants. In a different setting but with related motivation, [Easley and Ghosh \(2013\)](#) study a problem of “badge design”.

¹¹ Notice that the prize scheme is undefined when $|\arg \min_{j \in \mathcal{N}} s_j| > 1$. As the cumulative probability of simultaneous success is zero, there is no loss in ignoring simultaneous success here and elsewhere.

¹² Each O^t can be endowed with any sigma-algebra that contains each \mathbf{o}^t as an element.

success; this is, however, without loss of generality because agents never exert effort after succeeding, so messages sent to successful agents do not matter. Thus, our definition of a disclosure policy rules out disclosure that is stochastic or asymmetric across agents, but is otherwise general.

Two simple disclosure policies will be of particular interest. A *public* contest is one in which all information about the outcome history is disclosed at all times: for all t , $M_t = O^t$ and $\mu_t(o^t) = o^t$. A *hidden* contest is one in which no information is disclosed until the end of the contest: for all t , $|M_t| = 1$. (In either case, the contest runs until its deadline.)

The principal designs a contest to maximize her expected payoff gain, which, given requirement (iii) of the prize-sharing scheme, is the expectation of

$$(v - \bar{w})p_0 \left(1 - e^{-\lambda A^T}\right), \quad (7)$$

where A^T is the aggregate cumulative effort induced by the contest. The principal's problem can be decomposed into two steps: first, for any given prize \bar{w} , solve for the optimal prize-sharing scheme and information disclosure policy; second, use the first-step solution to solve for the optimal prize. Note that given any $\bar{w} < v$, the principal's objective in the first step is to simply maximize the probability of obtaining a success, i.e. to maximize the expectation of (5).

Let $h_i^t = (m_z, o_{i,z}, a_{i,z})_{z < t}$ be the private history of agent i at time t . An agent i 's (pure) strategy is a measurable function that specifies, for each history h_i^t , a choice of effort at time t , $a_{i,t}$. Without loss, we interpret $a_{i,t}$ as i 's effort at t conditional on him not having succeeded by t , as an agent will never exert effort after succeeding. Our solution concept is Nash equilibrium; we restrict attention to symmetric equilibria (viz., equilibria in which all agents use the same strategy).¹³ We say that an agent i uses a *stopping strategy* with *stopping time* t if the agent exerts full effort (i.e., effort of one) until t so long as he has not learned that any agent (including himself) has succeeded, followed by no effort after t .

¹³ Agents will be playing sequentially rationally in the equilibria we characterize of our focal contests, and hence our analysis would not be affected by standard refinements such as (appropriately defined versions of) subgame perfection or perfect Bayesian equilibrium. Regarding symmetry: for some parameters, some contests will have asymmetric equilibria that induce more experimentation than symmetric equilibria. It is common to focus on symmetric equilibria in symmetric contest and experimentation environments (e.g., Moldovanu and Sela, 2001; Bonatti and Hörner, 2011).

4. Optimal Contests

In this section, we take as given an arbitrary prize, \bar{w} , and solve for optimal contests given the prize, i.e. those that maximize the probability of success given \bar{w} . Our main insights concern this step of the principal's problem; [Section 5.1](#) endogenizes the principal's choice of prize \bar{w} . Without loss, we take $\bar{w} < v$ (as the principal will never choose $\bar{w} \geq v$) and also assume $p_0 \lambda \bar{w} > c$, as otherwise no contest can induce experimentation. We begin our analysis by studying public and hidden contests in [Section 4.1](#) and [Section 4.2](#) respectively, and then compare them in [Section 4.3](#). [Section 4.4](#) tackles contests within a more general class and contains the paper's main theoretical result.

4.1. Public contests

In a public contest, an agent's success is immediately disclosed to all other agents. Agents therefore update their beliefs based on their outcomes as well as their opponents' outcomes, given the equilibrium strategies.

Consider a public contest with an arbitrary prize scheme $w(s_i, \mathbf{s}_{-i})$. Let $A_{-i,z}$ denote (i 's conjecture of) the aggregate effort exerted by i 's opponents at time z so long as no agent has obtained a success by z . We denote by $w_{i,t}$ the expected reward agent i receives if he is the first one to succeed at t , which depends on $w(s_i, \mathbf{s}_{-i})$ and the continuation strategies of the opponents who may continue to exert effort and share the prize. If some agent besides i is the first agent to succeed at t , we denote agent i 's expected continuation payoff by $u_{i,t}$. We suppress the dependence of the relevant variables on the strategy profile to save on notation. Agent i 's problem can then be written as

$$\max_{(a_{i,t})_{t \in [0,T]}} \int_0^T [(p_{i,t} \lambda w_{i,t} - c) a_{i,t} + p_{i,t} \lambda A_{-i,t} u_{i,t}] e^{-\int_0^t p_{i,z} \lambda (a_{i,z} + A_{-i,z}) dz} dt, \quad (8)$$

where $p_{i,t}$ is i 's belief that the state is good at time t (so long as success has not been obtained), given by the following analog of (4):

$$p_{i,t} = \frac{p_0 e^{-\int_0^t \lambda (a_{i,z} + A_{-i,z}) dz}}{p_0 e^{-\int_0^t \lambda (a_{i,z} + A_{-i,z}) dz} + 1 - p_0}.$$

To interpret the objective (8), note that $e^{-\int_0^t p_{i,z} \lambda (a_{i,z} + A_{-i,z}) dz}$ is the agent's belief that no success will obtain by time t . Conditional on the good state and no success by t , the instantaneous probability that the agent is the first to succeed at t is $\lambda a_{i,t}$, and the instantaneous probability that an agent besides i is the first to succeed at t is $\lambda A_{-i,t}$.

Solving for equilibria in an arbitrary public contest is not straightforward; instead, we take an indirect but also more insightful approach. First, observe that agent i can ensure $u_{i,t} \geq 0$ (by exerting no effort after t). It follows that the agent chooses $a_{i,t} = 0$ if $p_{i,t}\lambda w_{i,t} < c$, and thus $a_{i,t} > 0$ requires

$$p_{i,t} \geq \frac{c}{\lambda w_{i,t}} \geq \frac{c}{\lambda \bar{w}}, \quad (9)$$

where the second inequality is because $w_{i,t}$ must be less than the prize \bar{w} . Hence, the lowest belief at which agent i is willing to exert effort in a public contest is $p_{i,t} = \frac{c}{\lambda \bar{w}}$.

Consider now a public WTA contest, where the full prize is awarded to the first agent who succeeds: $w_{i,t} = \bar{w}$ and $u_{i,t} = 0$ for all $t \in [0, T]$. Since the agent's belief $p_{i,t}$ is decreasing over time, the unique solution to (8) in this case is $a_{i,t} = 1$ if $p_{i,t} \geq p^{PW}$ and $a_{i,t} = 0$ otherwise,¹⁴ where

$$p^{PW} := \frac{c}{\lambda \bar{w}}. \quad (10)$$

It follows that in a public WTA contest with deadline T , there is a unique equilibrium: all agents exert full effort until either a success is obtained, or the public belief, viz. expression (4) with $A_t = N$, reaches p^{PW} (or the deadline T binds), and they exert zero effort thereafter. To maximize experimentation, the deadline T is optimal if and only if $T \geq T^{PW}$, where T^{PW} is the time at which the public belief reaches p^{PW} given that all agents exert full effort, i.e.,

$$\frac{p_0 e^{-N\lambda T^{PW}}}{p_0 e^{-N\lambda T^{PW}} + 1 - p_0} = \frac{c}{\lambda \bar{w}}. \quad (11)$$

Comparing with condition (9) above, we see that an optimal public WTA contest induces effort by all agents until their belief reaches the lowest belief at which any agent is willing to exert positive effort in a public contest. It follows that this contest maximizes the probability of success (see [Remark 1](#)).

Proposition 1. *A winner-takes-all contest is optimal among public contests. In an optimal public winner-takes-all contest, each agent uses a stopping strategy with stopping time T^{PW} defined by (11). T^{PW} is increasing in p_0 and \bar{w} , decreasing in c and N , and non-monotonic in λ . The probability of success is increasing in p_0 , \bar{w} and λ , decreasing in c , and independent of N .*

(All proofs are in the [Appendix](#).)

¹⁴ Throughout, we break indifference in favor of exerting effort whenever this is innocuous.

The non-monotonicity of T^{PW} with respect to λ is due to two countervailing effects: on the one hand, for any given belief $p_{i,t}$, the marginal benefit of effort is larger if λ is higher; on the other hand, the higher is λ , the faster each agent updates his belief down following a history of effort and no success (cf. [Bobtcheff and Levy, 2014](#); [Halac, Kartik, and Liu, 2016](#)). Nevertheless, the stopping belief, p^{PW} , is decreasing in λ , as seen immediately from (10), and as a result the probability of obtaining a success is increasing in λ . It is also intuitive why p^{PW} is independent of the number of agents: the likelihood that multiple agents succeed at the same instant is second order, hence the only effect of higher N on an agent's incentives at time t (so long as no one has succeeded yet) is to lower the public belief at t .¹⁵

Remark 2. A public WTA contest implements the first best if and only if $\bar{w} = v$, as seen by comparing (6) and (10).

4.2. Hidden contests

In a hidden contest, an agent's success is not disclosed until the end of the contest. Agents therefore update their beliefs based only on their own outcomes.

Given a hidden contest with prize scheme $w(s_i, \mathbf{s}_{-i})$, denote by $w_{i,t}$ the expected reward agent i receives if he succeeds at time t , which depends on $w(s_i, \mathbf{s}_{-i})$ and the strategies of the opponents. Then agent i 's problem reduces to

$$\max_{(a_{i,t})_{t \in [0, T]}} \int_0^T \left(p_{i,t}^{(1)} \lambda w_{i,t} - c \right) a_{i,t} e^{-\int_0^t p_{i,z}^{(1)} \lambda a_{i,z} dz} dt, \quad (12)$$

where $p_{i,t}^{(1)}$ is i 's belief that the state is good at time t given that he has not succeeded by t , which is given by the following analog of (4):

$$p_{i,t}^{(1)} = \frac{p_0 e^{-\int_0^t \lambda a_{i,z} dz}}{p_0 e^{-\int_0^t \lambda a_{i,z} dz} + 1 - p_0}.$$

To interpret the objective (12), note that $e^{-\int_0^t p_{i,z}^{(1)} \lambda a_{i,z} dz}$ is the agent's belief that he will not succeed by time t . Conditional on the good state and the agent not succeeding by t , the agent succeeds at t and receives $w_{i,t}$ with instantaneous probability $\lambda a_{i,t}$.

Consider now a hidden ES contest, where the prize is shared equally among all successful agents. Since i 's expected reward for success, which we denote w_i^{HS} , is independent

¹⁵ [Keller, Rady, and Cripps \(2005\)](#) also have a result that, among "simple" equilibria, the amount of experimentation is invariant to the number of agents. The structure of payoff interdependence in their setting is different from a public WTA contest, however; their key force is free-riding. Owing to the differences, they find that more experimentation can be induced with more agents in "complex" equilibria of their model.

of when he succeeds, it is immediate from (12) that an optimal strategy for i is a stopping strategy where $a_{i,t} = 1$ if $p_{i,t}^{(1)} \lambda w_i^{HS} \geq c$ and $a_{i,t} = 0$ otherwise. Under a stopping strategy, $p_{i,t}^{(1)} = \frac{p_0 e^{-\lambda t}}{p_0 e^{-\lambda t} + 1 - p_0}$.

Let p_i^{HS} and T_i^{HS} be the stopping belief and time respectively. Deadline T is optimal if and only if $T \geq \max_{i \in \mathcal{N}} T_i^{HS}$, in which case an agent i 's stopping belief satisfies

$$p_i^{HS} \lambda w_i^{HS} = c, \quad (13)$$

where

$$p_i^{HS} = \frac{p_0 e^{-\lambda T_i^{HS}}}{p_0 e^{-\lambda T_i^{HS}} + 1 - p_0}. \quad (14)$$

In a symmetric equilibrium, T_i^{HS} and w_i^{HS} are independent of i and can be denoted T^{HS} and w^{HS} . Should agent i succeed at any time, he learns that the state is good, in which event (i) the probability he ascribes to any opponent succeeding (resp., not succeeding) by T^{HS} is $1 - e^{-\lambda T^{HS}}$ (resp., $e^{-\lambda T^{HS}}$); and (ii) he views opponents' successes as independent. Thus, when all opponents use a stopping time T^{HS} , an agent's expected reward for success in a hidden ES contest is

$$w^{HS} = \bar{w} \sum_{n=0}^{N-1} \frac{1}{n+1} \binom{N-1}{n} (1 - e^{-\lambda T^{HS}})^n e^{-(N-1-n)\lambda T^{HS}}. \quad (15)$$

The summation in (15) is the expected share of the prize that an agent receives for success. It can be shown (see the [proof of Proposition 2](#)) that this expected-share expression simplifies to $\frac{1 - e^{-N\lambda T^{HS}}}{(1 - e^{-\lambda T^{HS}})^N}$. Substituting this simplification and (14) into (13) implicitly defines the equilibrium stopping time, T^{HS} , by

$$\frac{1 - e^{-N\lambda T^{HS}}}{(1 - e^{-\lambda T^{HS}})^N} \frac{p_0 e^{-\lambda T^{HS}}}{p_0 e^{-\lambda T^{HS}} + 1 - p_0} = \frac{c}{\lambda \bar{w}}. \quad (16)$$

Each of the two fractions on the left-hand side of (16) is strictly decreasing in T^{HS} ; hence a hidden ES contest has a unique equilibrium among symmetric equilibria in stopping strategies. It can be shown that any symmetric equilibrium must be outcome equivalent to this one in the sense that the probability of success by each agent—equivalently, the private belief reached by the deadline T by each agent, and hence the public belief at T (see [Remark 1](#))—is the same.¹⁶

¹⁶ For $T \leq T^{HS}$, the unique equilibrium is the symmetric equilibrium in stopping strategies with stopping time T . When $T > T^{HS}$, even though best responses always exist in stopping strategies, there will be

Instead of using equal sharing, the principal may also consider dividing the prize in other ways to increase experimentation; for example, rewarding agents who succeed earlier with larger shares of the prize. However:

Proposition 2. *An equal-sharing contest is optimal among hidden contests. In an optimal hidden equal-sharing contest, there is an equilibrium in which each agent uses a stopping strategy with stopping time T^{HS} defined by (16); furthermore, all symmetric equilibria are outcome equivalent. T^{HS} is increasing in p_0 and \bar{w} , decreasing in c and N , and non-monotonic with respect to λ . The probability of success is increasing in p_0 , \bar{w} and λ and decreasing in c . An increase in N can increase or decrease the probability of success.*

Let us sketch the argument why it is optimal to share the prize equally among hidden contests. Using the fact that an agent's expected reward for success is independent of when he succeeds in a hidden ES contest, the [proof of Proposition 2](#) shows that each agent's (dynamic) incentive constraint binds at each $t \in [0, T^{HS}]$; that is, at each moment in an optimal hidden ES contest, an agent is indifferent over how much effort to exert given that he has exerted full effort in the past and will behave optimally in the future. Intuitively, exerting no effort at time t precludes success at t but increases the continuation payoff after t (as the private belief does not decrease); these effects cancel when the reward for success is constant over time. That each agent's incentive constraint is binding at each point is shown to imply that no hidden contest can satisfy all agents' incentives up to T^{HS} if any agent works beyond T^{HS} ; thus, no hidden contest can induce more cumulative effort from each agent than T^{HS} .

The comparative statics in [Proposition 2](#) are largely intuitive, so we only note two points. First, the non-monotonicity of T^{HS} in λ owes to the same countervailing forces that were noted after [Proposition 1](#). Second, the probability of obtaining a success may increase or decrease when N increases, as shown in [Figure 1](#). Intuitively, holding fixed T^{HS} , conditional on at least one opponent having succeeded (which reveals the state to be good), a larger number of opponents having succeeded only lowers the expected benefit of effort for an agent at T^{HS} . However, an increase in N also lowers T^{HS} , which by itself decreases the probability that any opponent has succeeded by T^{HS} conditional on the good state.

symmetric equilibria in which agents do not play stopping strategies. The reason is that when $T > T^{HS}$ there are multiple strategies by which an agent can arrive at T with a private belief corresponding to the stopping strategy with stopping time defined by (16). However, in any symmetric equilibrium, the cumulative effort by each agent must be T^{HS} ; this is because an agent's expected share of the prize from success is strictly decreasing in each opponent's cumulative effort, and his own private belief at T is strictly decreasing in his own cumulative effort.




Figure 1 – Probability of success in an optimal hidden ES contest. Parameters are $p_0 = 0.95$, $c = 0.4$, $\bar{w} = 1$, and $\lambda = 1.5$; the maximum is obtained at $N = 3$.

4.3. Public versus hidden contests

[Proposition 1](#) shows that a WTA contest is optimal among public contests and [Proposition 2](#) shows that an ES contest is optimal among hidden contests. To compare public contests with hidden contests, it therefore suffices to compare the stopping times identified in those propositions: the principal strictly prefers hidden to public if and only if $T^{HS} > T^{PW}$, where T^{PW} is given by [\(11\)](#) and T^{HS} is given by [\(16\)](#). Since the left-hand side of [\(11\)](#) and that of [\(16\)](#) are each decreasing functions of T^{PW} and T^{HS} respectively, $T^{HS} > T^{PW}$ if and only if the left-hand side of [\(16\)](#) would be strictly larger than its right-hand side were $T^{HS} = T^{PW}$.

Proposition 3. *The principal strictly prefers an optimal hidden contest to an optimal public contest if and only if*

$$\frac{1 - e^{-\lambda N T^{PW}}}{(1 - e^{-\lambda T^{PW}})N} \frac{p_0 e^{-\lambda T^{PW}}}{p_0 e^{-\lambda T^{PW}} + 1 - p_0} \lambda \bar{w} > c, \quad (17)$$

where T^{PW} is defined by [\(11\)](#).

To see the intuition behind [\(17\)](#), rewrite it as¹⁷

¹⁷ Inequality [\(17\)](#) is equivalent to, for any $i \in \mathcal{N}$,

$$c < \Pr[\text{some } j \neq i \text{ succeeded by } T^{PW} \mid i \text{ did not}] \lambda \bar{w} L \\ + \Pr[\text{no } j \neq i \text{ succeeded by } T^{PW} \mid i \text{ did not}] \Pr[G \mid \text{no success by } T^{PW}] \lambda \bar{w},$$

$$\lambda \bar{w} \sum_{n=1}^{N-1} \frac{\Pr[n \text{ opponents succeeded by } T^{PW} \mid G]}{\Pr[\text{at least one opponent succeeded by } T^{PW} \mid G]} \left(\frac{1}{n+1} \right) > c. \quad (18)$$

Assume all opponents are using a stopping time T^{PW} in a hidden ES contest. At T^{PW} , conditional on all opponents having failed, agent i is indifferent over his effort (by definition of T^{PW}). So, he strictly prefers to continue if and only if he strictly prefers to continue conditional on at least one opponent having succeeded. The left-hand side of (18) is i 's expected benefit from effort at T^{PW} conditional on some opponent having succeeded (which implies the state is good); the right-hand side of (18) is the cost.

When $N = 2$, inequality (18), and hence inequality (17), simplifies to

$$\lambda \frac{\bar{w}}{2} > c. \quad (19)$$

This condition is transparent: it says that an agent would continue experimenting if he knew his only opponent had already succeeded, in which case he infers the state is good but success will only earn him half the prize.

For $N > 2$, (19) is a necessary condition for the optimal hidden ES contest to dominate the optimal public WTA contest, but it is not sufficient. Inspecting (18), a simple sufficient condition is

$$\lambda \frac{\bar{w}}{N} \geq c. \quad (20)$$

This condition says that an agent would continue experimenting if he knew that all his opponents have succeeded. The example in Figure 1 shows that condition (20) is not necessary, as there, hidden ES dominates public WTA for all $N \in \{2, \dots, 13\}$ even though (20) fails when $N > 3$.¹⁸

An increase in p_0 decreases the left-hand side of (17), making the dominance of hidden ES less likely (in the sense of weakly shrinking the set of other parameters for which the inequality holds).¹⁹ This reinforces the intuition that the gains from using hidden ES stem from bolstering agents' beliefs that the innovation may be feasible despite their own failures, which is more important to the principal when the prior is lower; were $p_0 = 1$, in

where L denotes an agent's expected share of the prize for success at T^{PW} in a hidden ES contest given that all opponents use a stopping time T^{PW} and at least one opponent has succeeded. The definition of T^{PW} in (11) implies $\Pr[G \mid \text{no success by } T^{PW}] \lambda \bar{w} = c$. Thus, (17) is equivalent to just $\lambda \bar{w} L > c$, which is (18).

¹⁸ Recall that, as shown in Proposition 1, the probability of success in an optimal public WTA contest is independent of N ; hence, it is equal to that under an optimal hidden ES contest when $N = 1$.

¹⁹ To verify this claim, use (18): an increase in p_0 increases T^{PW} , which decreases the left-hand side of (18) because of a first-order stochastically dominant shift of the relevant probability distribution.

which case there would be no learning, public WTA would be an optimal contest.²⁰ The necessary and sufficient conditions (19) and (20) also reveal that hidden ES dominates (resp., is dominated by) public WTA when $\frac{c}{\lambda w}$ is sufficiently small (resp., large).

The discussion above assumes a fixed number of agents. If the principal can instead choose the number of agents, then an optimal hidden ES contest always does at least as well as any public WTA contest. This is because the principal can replicate the public WTA outcome by setting $N = 1$ and using hidden ES.²¹ An implication of Proposition 3 is that combining hidden ES with an optimally chosen $N > 1$ can be strictly better than using public WTA with any N ; this is the case if and only if condition (17) holds. Furthermore, as seen in Figure 1, it can be optimal to set $N > 2$; this observation contrasts, for example, with a result of Che and Gale (2003) in a different contest environment. The rationale for why multiple agents can be beneficial in our setting appears novel. Our model shuts down standard channels like heterogeneity among agents, convex effort costs, and discounting, so that the number of agents is irrelevant in the first best. Nevertheless, having multiple agents allows the principal to harness the benefits from hiding information and sharing the prize.

4.4. General information disclosure policies

We now turn to disclosure policies beyond the extremes of public and hidden disclosure. A contest is *rank monotonic* if its prize-sharing scheme is such that if agent i succeeds earlier than agent j , then i receives a weakly larger share of the prize than j . Formally, in a rank-monotonic contest, the prize-scheme $w(\cdot)$ is such that for any profile of success times \mathbf{s} , $s_i < s_j \implies w(s_i, \mathbf{s}_{-i}) \geq w(s_j, \mathbf{s}_{-j})$. Recall that our convention is $\emptyset > T$. Proposition 1 and Proposition 2 show that among contests with public and hidden disclosure, the optimal contest is rank monotonic.

For any history of outcomes \mathbf{o}^t , denote by $\eta(\mathbf{o}^t)$ the number of successes up to time t . A disclosure policy is *cutoff disclosure* with cutoff n if at any time $t \in [0, T]$, $M_t = \{0, 1\}$ and $\mu_t(\mathbf{o}^t) = 1$ if $\eta(\mathbf{o}^t) \geq n$ while $\mu_t(\mathbf{o}^t) = 0$ if $\eta(\mathbf{o}^t) < n$. That is, the principal simply makes an announcement ($m_t = 1$) as soon as n or more agents have succeeded and is otherwise

²⁰ If condition (20) holds, then hidden ES would also be optimal were $p_0 = 1$.

²¹ The principal can effectively reduce the number of agents by using a non-anonymous prize scheme that simply offers no reward for success to some agents. An optimal hidden contest without our restriction to anonymous prize schemes is thus always weakly better than public WTA. Moreover, as seen in Figure 1, a hidden contest with a non-anonymous prize scheme can strictly improve on both hidden ES and public WTA (by effectively reducing $N > 3$ to $N = 3$ in the figure). Note that the proof of Proposition 1 shows that among public contests, non-anonymous prize schemes cannot improve on WTA.

silent ($m_t = 0$). Among all possible cutoffs, the following one turns out to be key:

$$n^* := \max \left\{ n \in \{1, \dots, N\} : \lambda \frac{\bar{w}}{n} \geq c \right\}.$$

Note that $\lambda \bar{w} \geq c$ by assumption, and hence $n^* \geq 1$ is well defined. The value n^* is such that at an ES contest's deadline, an agent would be willing to exert effort if he knew that some opponent but less than n^* opponents have succeeded; however, if he knew that n^* or more opponents have succeeded, he would not be willing to exert effort.

We can now state the paper's main result.

Proposition 4. *A cutoff-disclosure equal-sharing contest with cutoff n^* is optimal among rank-monotonic contests.*

By definition of n^* , a cutoff-disclosure ES contest with cutoff n^* has the property that all agents stop exerting effort when the principal makes the announcement that n^* agents have succeeded. Moreover, so long as the principal is silent, an agent would be willing to exert effort if he conditions on at least one opponent having succeeded; in fact, the proof of [Proposition 4](#) establishes that an agent always has a best response in stopping strategies, i.e., in which the agent exerts full effort until either the principal makes the announcement or he reaches a stopping time (or he succeeds).

To see the logic behind [Proposition 4](#), consider (symmetric) equilibria in the class of aforementioned best responses. Denote by T^* the agents' stopping time in an optimal cutoff-disclosure ES contest.²² At T^* , each agent's cumulative effort (given no success) is equal to T^* and his static incentive constraint for effort binds. If another rank-monotonic contest induced a strictly higher probability of success in some (symmetric) equilibrium, there would be a time t in that equilibrium at which each agent's cumulative effort is also equal to T^* but his static incentive constraint for effort is slack. However, we show that holding fixed any profile of cumulative efforts, the static incentive constraint cannot be relaxed by moving away from cutoff disclosure with cutoff n^* . The reason is that, by rank monotonicity, an agent who succeeds at t receives a reward no larger than an equal share of the prize when shared among those who succeeded by t , and, by definition of n^* , an agent expects to benefit from exerting effort conditional on the good state and equal sharing if and only if the sharing is with $n^* - 1$ or fewer opponents. As a result, given

²² Similar to our discussion of hidden ES contests, particularly in [fn. 16](#), one can show that all symmetric equilibria are outcome equivalent in an optimal cutoff-disclosure ES contest. If the deadline is chosen to be T^* , the unique equilibrium is the symmetric stopping-strategy equilibrium with stopping time T^* .

a history of no success, an agent's incentive to exert effort is increased if the principal reveals n^* or more successes and is reduced if the principal reveals $n^* - 1$ or less successes.

Observe that $n^* = 1$ when $\lambda\bar{w}/2 < c$, whereas $n^* = N$ when $\lambda\bar{w}/N > c$. A cutoff-disclosure ES contest with cutoff $n^* = 1$ is equivalent to a public WTA contest, while one with cutoff $n^* = N$ is equivalent to a hidden ES contest. Thus, we obtain the following generalization of the comparison in [Section 4.3](#):

Corollary 1. *Among rank-monotonic contests, a public winner-takes-all contest is optimal if $\lambda\bar{w}/2 < c$ and a hidden equal-sharing contest is optimal if $\lambda\bar{w}/N > c$.*

Analogous to [Proposition 2](#), the probability of success in an optimal cutoff-disclosure ES contest need not be monotonic in the number of agents. If the principal can choose the number of agents, she may optimally choose N such that $\lambda\bar{w}/N < c$ (in which case $n^* < N$), but would never choose N such that $\lambda\bar{w}/(N + 1) > c$ because increasing N would be beneficial. A hidden ES contest with $N = n^*$ agents may induce a higher or lower probability of success than a cutoff-disclosure ES contest with cutoff n^* and $N > n^*$ agents.

[Proposition 4](#) allows for any (deterministic and symmetric) disclosure policy, but restricts attention to rank-monotonic prize schemes. Within certain classes of information disclosure, rank monotonicity of the prize scheme is not restrictive from the viewpoint of optimality. For example, say that an information disclosure policy is *intermittently public* if there is a set $\mathcal{T} \subseteq [0, T]$ such that: (i) at any time $t \in \mathcal{T}$, the principal sends a message $m_t = (o_{1,z}, \dots, o_{N,z})_{z < t}$; and (ii) at any time $t \notin \mathcal{T}$, the principal sends a message $m_t = \emptyset$. In other words, the principal reveals the entire history of outcomes at each $t \in \mathcal{T}$, while she is silent at each $t \notin \mathcal{T}$. Intermittently-public contests have a particularly simple disclosure structure and subsume both public and hidden contests. We show in the online appendix (available on the journal's website) that rank monotonicity is without loss of optimality among intermittently-public contests; it follows that no intermittently-public contest can improve on the optimal cutoff-disclosure ES contest.

Finally, given the salience and widespread use of WTA contests, we record:

Proposition 5. *A public contest is optimal among winner-takes-all contests.*

The logic is as discussed earlier: the principal only cares about agents' effort incentives following a history of no successes; hiding any information about this history in a WTA contest reduces effort incentives because when an agent conditions on some opponent having succeeded, his expected reward for success is zero. Indeed, as is clear from the proposition's proof, the result also applies to stochastic and asymmetric disclosure.

5. Discussion

Having established our main points, we now study the principal’s problem of choosing the prize \bar{w} and discuss some extensions and variations of our model.

5.1. Optimal prize

Section 4 solved for an optimal contest—one that maximizes the probability of success—given any prize $\bar{w} < v$. The second stage of the principal’s problem is to choose an optimal prize. The principal chooses \bar{w} to maximize (7), where A^T is the aggregate cumulative effort induced by an optimal rank-monotonic contest associated with prize \bar{w} . Building on the preceding analysis, the solution to this problem is a cutoff-disclosure ES contest with prize $\bar{w} \in (0, v)$. Furthermore:

Proposition 6. *Fix any set of parameters (p_0, λ, c, N) and consider rank-monotonic contests. When the value of the innovation, v , is large (resp., small) enough, the principal maximizes (7) by choosing a prize $\bar{w} \in (0, v)$ and a hidden equal-sharing (resp., public winner-takes-all) contest.*

The idea is that the larger is v , the larger is the principal’s gain from inducing more experimentation, and hence the larger is the optimal prize. For v large enough, the principal optimally chooses a prize \bar{w} large enough that $\lambda\bar{w}/N > c$; for v small enough, the optimal prize is small enough that $\lambda\bar{w}/2 < c$. Proposition 6 then follows from Corollary 1.

Our definition of contests in Section 3.3 required the total payment by the principal to be constant so long as there is at least one success (i.e., we have required the principal’s prize scheme to satisfy: $\mathbf{s} \neq (\emptyset, \dots, \emptyset) \implies \sum_{i=1}^N w_i(\mathbf{s}) = \bar{w}$). It is this requirement that implied that given any prize $\bar{w} < v$, the principal simply seeks to maximize the probability of success. We could instead assume the principal need not pay out the entire prize after a success (in order to save on payments), but require the total payment to satisfy a budget constraint. That is, given a budget $\bar{W} \geq 0$, the principal chooses a contest to maximize the expectation of

$$\left(v - \sum_{i=1}^N \mathbb{E}[w(s_i, \mathbf{s}_{-i}) | \mathbf{s} \neq (\emptyset, \dots, \emptyset)] \right) p_0 \left(1 - e^{-\lambda A^T} \right),$$

subject to a payment scheme satisfying, for all \mathbf{s} , $\sum_{i=1}^N w(s_i, \mathbf{s}_{-i}) \leq \bar{W}$ (and the other requirements described in Section 3.3). For arbitrary parameters of this problem, the optimal payment scheme need not resemble a “contest” in the sense of simply dividing a fixed total prize among successful contestants. Nevertheless, we can show that when

$\lambda\bar{W}/N > c$ in this budget-constrained problem, a hidden ES contest with prize \bar{W} is asymptotically optimal as the value of the innovation v becomes large.²³ As the total payment the principal makes conditional on at least one success is constant in a hidden ES contest, this implies that restricting attention to payment schemes with this property is almost without loss of optimality when the principal’s budget constraint is tight relative to her value of innovation. Furthermore, any public contest is strictly sub-optimal. Our finding contrasts with [Moroni \(2015\)](#), who shows that a public (i.e., full-disclosure) mechanism is optimal in the absence of a budget constraint. The takeaway is that even when the principal values saving money, hiding information is beneficial when the innovation’s value is (sufficiently) larger than the principal’s budget.

5.2. Observability of success

Our model has assumed that a success is observed by both the agent who succeeds and the principal. We now consider what happens if only the principal or only the agent directly observes a success and can choose whether and when to verifiably reveal it to the other party. We also discuss agents’ ability to verifiably reveal a success to their opponents. In all these cases, we will see that the main results of [Section 4](#) continue to hold.

Only principal observes success. Suppose the principal observes an agent’s success but the agent does not. The principal can choose when to reveal a success to the agent. This scenario is relevant, for example, to the Netflix contest: there, contestants had to submit an algorithm whose performance was evaluated by Netflix on a proprietary “qualifying dataset” to determine whether the 10% improvement target had been achieved. We suppose that if a success is obtained, the principal must reveal it by the end of the contest; she cannot harness the innovation’s benefits without paying out the prize. We now interpret the arguments of the prize-sharing scheme $w(s_i, s_{-i})$ as the times at which the principal reveals agents’ successes.

It is readily seen that our results extend to this setting. As the principal only values one success and optimally chooses a prize $\bar{w} < v$, she has no incentive to not reveal a success immediately to an agent who succeeds. At each time t , an agent conditions on not having obtained a success by t unless the principal has revealed otherwise. Consequently, the analysis of [Section 4](#) applies without change. Indeed, verifiable revelation by the principal is not essential: the same outcomes can be supported even if the principal is only able to make cheap-talk or unverifiable statements about an agent’s success.

²³ More precisely, for any $\varepsilon > 0$, a hidden ES contest with prize \bar{W} is ε -optimal when v is large enough.

Only agent observes success. Suppose next that the principal does not observe success directly; rather, any agent who succeeds can choose voluntarily when to verifiably reveal his success to the principal. This assumption is obviously relevant for many applications. The payments $w(s_i, s_{-i})$ are now interpreted as a function of the times at which agents reveal their success.

It is weakly dominant for an agent to immediately reveal his success in a WTA contest. Immediate revelation to the principal is also optimal in a hidden ES contest, and more generally in a cutoff-disclosure ES contest with cutoff n^* . Given the analysis in [Section 4](#), it follows that in any of these contests, there exist symmetric equilibria where agents follow stopping strategies and reveal their successes immediately, inducing the same outcome as when both the principal and the agent directly observe success. Naturally, verifiability is important here; the same outcome cannot be obtained with cheap talk by the agents.

Agents can reveal success to opponents. Lastly, suppose agents can verifiably reveal their success to other agents. Would they have the incentive to do so? While the issue is moot in public contests, it is paramount in hidden contests, because such an incentive would unravel the principal's desire to keep successes hidden.

In a hidden ES contest, a successful agent wants to deter opponents from continuing experimenting so that he can secure a larger share of the prize. Revealing a success has two opposing effects: it makes opponents more optimistic about the innovation's feasibility but decreases their expected prize shares from their own success.²⁴ An agent's incentive to reveal that he has succeeded (so long as no other agent has already done so) will thus generally involve a tradeoff, and the tradeoff's resolution could potentially harm the principal. However, if condition (20) holds, the resolution is unambiguous: revealing a success always increases experimentation by other agents. Therefore, that sufficient condition for hidden ES to be optimal also ensures that agents have no incentives to reveal their success to opponents, and hence the principal can indeed implement hidden ES when (20) holds. More generally, note that in a cutoff-disclosure ES contest, agents have no incentives to reveal their successes to opponents if the cutoff is n^* : revealing success before the principal has announced that n^* agents have succeeded increases other agents' incentives to experiment.

The foregoing discussion presumes that an agent can verifiably reveal a success to his opponents without actually making the innovation available to them. This would be dif-

²⁴ Once an agent reveals success, all other successful agents will have strict incentives to reveal too: with the uncertainty about the innovation's feasibility resolved, the only effect of revelation is to lower opponents' expected prize shares.

difficult in many contexts; for example, a contestant in the Netflix contest could probably not prove that he has succeeded without sharing his algorithm. Clearly, if verifiable revelation implies sharing the innovation, then an agent would never reveal a success to his opponents in a hidden ES contest. Moreover, revelation cannot be credibly obtained if messages are cheap talk.

5.3. Socially efficient experimentation

We have focused on contest design for a principal who does not internalize agents' effort costs. But our analysis also implies that public WTA contests can be dominated by a hidden ES contest even for a social planner who does internalize these costs. As noted in [Remark 2](#), a public WTA contest implements the first-best solution if and only if the prize \bar{w} is set to be equal to the social value of a success v . Thus, if the social planner has a binding budget constraint (cf. [Section 5.1](#)), this contest may not be efficient. In particular, if condition (17) holds—so that an optimal hidden ES contest induces a later stopping time than any public WTA contest—and if v is significantly larger than the available budget, then hidden ES will be preferred to public WTA net of effort costs: even though hidden ES induces wasteful effort after the innovation is first obtained, it increases agents' incentives to experiment. It is likely that in various circumstances, the social value of innovation is substantially larger than the prize available to a contest designer, e.g. for medical innovations or scientific discoveries.

5.4. Other issues

Discounting. Our analysis has abstracted away from discounting. Incorporating a discount rate would introduce additional forces, making the tradeoffs that we highlight less transparent. For example, in the presence of discounting, having multiple agents simultaneously experiment would be desirable in the first best (to speed up experimentation), whereas we show that this feature can arise in the optimal contest even when it provides no social value. From a robustness perspective, however, discounting would not qualitatively alter our main points. Analogous to our analysis of [Section 4.3](#), we derive in the online appendix a necessary and sufficient condition for the probability of success in an optimal hidden ES contest to be higher than that in an optimal public WTA contest, given an arbitrary discount rate $r \geq 0$. This condition is continuous in r and reduces to condition (17) in [Proposition 3](#) when $r = 0$. Discounting also affects the computation of the principal's ex-ante payoff: one advantage of public WTA contests when there is discounting is that the principal can profit from an agent's innovation immediately following success, whereas hidden ES contests must be run until their deadline. Yet, because

a hidden ES contest can induce a higher ex-ante probability of the innovation, a tradeoff still arises, and our insights from comparing these contests remain relevant even with discounting.

Convex effort costs. Our analysis has assumed a linear cost of effort. Incorporating a convex instantaneous effort cost would introduce new forces: for example, in a public WTA contest, agents' equilibrium effort would be front-loaded. Nevertheless, we show in the online appendix that the stopping beliefs induced by a public WTA contest and a hidden ES contest (given by (10) and (13) respectively), and hence the induced probabilities of success, are invariant to convex costs, because stopping beliefs are determined by marginal costs. Therefore, our main points are robust: if (and only if) condition (17) holds, the principal strictly prefers an optimal hidden ES contest to an optimal public WTA contest, regardless of whether the instantaneous effort cost is linear or convex.

Multistage contests. Innovations sometimes require a sequence of successes. Suppose, for example, the principal only gains the profit v when an agent successfully completes two tasks or stages. In the online appendix, we illustrate how our analysis may be useful even in such settings. Intuitively, the second-stage "continuation value" augments any first-stage prize. We show that, under certain conditions, it is preferable to implement some sharing of the overall prize with hiding information about outcomes.

6. Applications

Our focus has been on designing a contest to obtain an innovation. More broadly, our framework and conclusions may be usefully applied to various settings in which agents work on projects of uncertain quality or feasibility. We discuss three applications in this section. In each case, we focus only on an aspect of the application that this paper's analysis sheds some light on, setting aside other aspects that are pertinent in practice.

6.1. Patent law

As it is common for multiple inventors to work on the same invention, the priority rule that is used to allocate patent rights is of substantial economic significance. Until 2013, U.S. patent law followed the principle of granting patents to those who were the "first to invent" (FTI), provided the date of first invention could be documented; some other countries also used to follow the FTI principle. However, all countries now use the "first to file" (FTF) system, under which a patent is issued to the inventor who first files an application with the patent office, regardless of priority in discovery.

FTF is generally advocated on the grounds that it is less costly to administer—it is difficult to prove priority in discovery—and it encourages earlier patent applications. Given the disclosure requirements for patentability, earlier applications mean earlier disclosure of inventions to the public. The disclosure benefits of FTF are stressed, for example, by [Scotchmer and Green \(1990\)](#). Sharing knowledge increases innovation by allowing innovators to build upon others’ discoveries and reducing duplication, but, for these same reasons, inventors have incentives not to disclose information to their competitors so as to retain a competitive advantage. Inventors will thus tend to delay filing under FTI—as they can still claim priority if a competitor later files—and FTF can improve upon FTI by inducing earlier filing and disclosure.

Our model illuminates a novel point: more disclosure can depress innovation because agents become pessimistic about their own likelihood of success when they don’t observe others filing. More specifically, since innovators have incentives to file immediately under an FTF rule, and their innovations become public upon filing, the FTF system can be interpreted as akin to a public WTA contest in our framework, where the fixed prize corresponds to the competitive advantage that the patent confers. On the other hand, an FTI system can be viewed as a hidden shared-prize contest: it is hidden because innovators delay filing and thus disclosure, and it entails prize sharing because documentation on invention date is typically inconclusive, so either first or later inventors may be granted the patent in an interference proceeding. In fact, in a large share of cases, parties negotiate cross-licensing agreements to settle the dispute, so they indeed end up sharing the prize ([Lerner, 1995](#)). Hence, even if neither FTF nor FTI are optimal mechanisms, their comparison can be related to our comparison of public WTA and hidden shared-prized contests. Our results imply that one advantage of FTI over FTF is that it can produce more innovation when learning is important—not despite the fact that FTI limits disclosure, but rather precisely because of it, combined with the sharing of rewards FTI induces.²⁵

6.2. Organizational economics

An important problem in organizations is task allocation. The seminal work of [Holmström and Milgrom \(1991\)](#) highlights the issue of performance measurability and argues that it is never optimal to assign two agents to the same task. [Itoh \(1991\)](#), on the other hand, emphasizes specialization and agents’ incentives to help each other, and obtains a result that optimal job design consists of either unambiguous division of labor or sub-

²⁵ There have been other arguments in favor of FTI; for example, it protects small inventors who are more resource constrained and hence slower in transforming an invention into an application ([Abrams and Wagner, 2013](#)).

stantial teamwork. Another strand of the literature studies task allocation as a response to workers' cognitive limits; see [Garicano and Prat \(2013\)](#) for a survey.

Our framework provides insight into a different consideration in job design, viz., learning about task difficulty. Consider a principal who must assign two tasks of uncertain and independent difficulty to two agents. One option is to assign one task to each agent and reward each agent if he succeeds in completing his task. This scheme is equivalent to implementing two public WTA contests, one for each task.²⁶ Alternatively, the principal can assign both tasks to both agents, in which case each agent must decide how to allocate his time between the two tasks. If agents are rewarded based on an equal-sharing scheme—with an agent receiving the full reward for a task if he is the only one to succeed and half the reward if both succeed—and no information about their progress is disclosed until a specified deadline, this scheme effectively implements two hidden ES contests.²⁷ Our results on contest design imply that making the two agents jointly responsible for the two tasks can be beneficial, as it allows the principal to keep the agents optimistic that they will be able to successfully complete their jobs and receive the corresponding reward.

6.3. Government procurement

The AMERICA Competes Reauthorization Act mentioned in the [Introduction](#) is one of several recent efforts of the U.S. government to encourage greater innovation in federal contracting. One of the goals is to help agencies reach out beyond traditional contractors and “pay contractors for results, not just best efforts.”²⁸ The “incentive prize” model of procurement promotes innovation by offering a monetary reward upon completion of a specific task. Examples include the Astronaut Glove Challenge run by NASA in 2009 and the Progressive Insurance Automotive X Prize sponsored by the Department of Energy in 2011. Another model is “challenge-based acquisition,” under which firms must demonstrate product performance for a contract award or task orders for additional refinement or production of their proposed solution. An example is the MRAP II pre-solicitation launched by the Marine Corps Systems Command in 2007 for the development of a new model of mine resistant ambush protected (MRAP) vehicles.

While agencies' needs and their specific context introduce additional features, our study of contests can be used to understand some key elements of these procurement

²⁶ Recall that the outcome of a public WTA contest is independent of the number of agents in the contest.

²⁷ It can be verified that given this hidden ES structure, there is a symmetric equilibrium in which each agent works for the same amount of time on each of the two tasks in the absence of success.

²⁸ See *Innovative Contracting Case Studies*, Office of Science and Technology Policy (OSTP) and Office of Management and Budget's Office of Federal Procurement Policy (OFPP), 2014 for details on the programs and examples described in this section. The quote is from p. 2.

processes. The prize is the monetary reward offered by the agency in the incentive prize model or the value of the contract in the challenge-based acquisition model. Under both models, federal agencies specify a date by or at which participants must submit their proposals/demonstrate their solutions. Since the evaluation is made at the deadline, in general there is no information disclosed about participants' progress before this date; in fact, participants themselves have incentives to keep their achievements hidden to prevent others from replicating their solutions.

Regarding the allocation of the prize, sharing is particularly common in challenge-based acquisitions.²⁹ In the case of MRAP II, initial testing disqualified vehicles by five companies that did not meet the requirements; the other two participating firms, BAE Systems and Ideal Innovations, were both awarded orders. More generally, in recent years the Federal Acquisition Regulation has placed greater emphasis on multiple award contracts, a type of indefinite-quantity contracts which are awarded to multiple contractors from a single solicitation (Fusco, 2012). There are various reasons for the use of multiple suppliers, including ensuring a stable supply and maintaining continuous competition. Our results imply that in settings with uncertainty, contract sharing can be beneficial beyond these considerations (given the aforementioned hidden disclosure): it better incentivizes firms to invest in the desired innovation in the first place.

7. Conclusion

This paper has studied contest design for innovations in environments with learning, using the exponential-bandit framework of experimentation. We have shown that a winner-takes-all contest in which any successful innovation is disclosed immediately ("public WTA") can be dominated by a contest in which no information is disclosed until the end, at which point all successful agents equally share the prize ("hidden ES"). More generally, among rank-monotonic contests with any deterministic and symmetric disclosure policy, a contest in which the principal announces when a critical number of agents have succeeded and successful agents equally share the prize is optimal; simple sufficient

²⁹ Procurement often combines elements of both innovation races and research tournaments: agencies specify minimum requirements that must be met to qualify for the prize, but among those who succeed in meeting these objectives by the deadline, winners may be selected according to their relative performance. This was the case in the Astronaut Glove Challenge and the Progressive Insurance Automotive X Prize (see the contests' rules, available at <http://astronaut-glove.tripod.com> and <http://auto.xprize.org> respectively). Relative performance evaluations incentivize participants to continue exerting effort towards improving their solutions after achieving the specified objectives. Furthermore, to the extent that a participant's performance beyond the (typically high) minimum requirements, conditional on meeting these requirements, is difficult to predict, this allocation can also be taken to implement prize-sharing in expectation.

conditions guarantee optimality of either public WTA or hidden ES.

Our formal analysis is within the confines of a specific model. Naturally, we do not suggest that all the particulars of our results will continue to apply verbatim in richer environments; for example, it will not be *precisely* equal sharing that is optimal among hidden (or cutoff-disclosure) contests when there is discounting or realistic asymmetries among agents. However, we believe the fundamental intuition we have explored—a tradeoff between the reward an agent expects to receive should he succeed and his belief about the likelihood of success—is valid quite generally. Our work suggests that the common default assumption (for theory) or prescription (for policy) of using WTA schemes deserves further scrutiny when the feasibility of the innovation is uncertain and successful innovations are not automatically public information. Our broad message is to consider hiding information, at least partially, while also sharing the prize in some form. In particular, alternative patent schemes that implement some version of “sharing the prize” may be warranted for certain kinds of R&D. On the other hand, in contexts where innovations are publicly observable, our analysis implies that a WTA contest is optimal; note though that the principal would sometimes be willing to pay a cost to alter the observability structure.

We hope that future research will make progress on dropping some restrictions we have had to make for tractability; specifically, it is desirable for the scope of contest design to encompass stochastic and asymmetric disclosure and non-rank-monotonic prize schemes.

We have taken the number of contestants, N , to be fixed (or to be specified by the principal). An alternative that may be more appropriate for some contexts would be to consider a fixed entry cost—any contestant must incur some cost to either register for the contest or to get started with experimentation—and endogenously determine N through free entry. We believe our main themes would extend to such a specification, but endogenizing the number of contestants in this manner may yield additional insights. Naturally, the entry cost itself can also be set by the principal.

It would also be interesting to incorporate heterogeneity among agents into the current framework. For example, agents may be privately informed about their “ability” (λ) or their cost of effort (c). While introducing the latter is likely to have intuitive effects—agents with higher c stop experimenting sooner—the former would be more subtle because of the countervailing effects on stopping times discussed following [Proposition 1](#).

A. Appendix

A.1. Proof of Proposition 1

The first part of the result was proven in the text. For the last part: the probability of obtaining a success in an optimal public WTA contest is given by expression (5) with $A^{T^{PW}} = NT^{PW}$. By Remark 1, the comparative statics for the probability of success follow immediately from the left-hand side (LHS, hereafter) of (10) being increasing in c , decreasing in λ and \bar{w} , and independent of p_0 and N .

The comparative statics of T^{PW} are obtained through straightforward manipulations of equation (11). Specifically, the comparative statics with respect to c and \bar{w} follow from the fact that the right-hand side (RHS, hereafter) of (11) is increasing in c and decreasing in \bar{w} , while the LHS is independent of these parameters and decreasing in T^{PW} . Similarly, the comparative statics with respect to p_0 and N follow from the fact that the LHS of (11) is increasing in p_0 , decreasing in N and decreasing in T^{PW} , while the RHS is independent of these parameters. Finally, to compute the comparative static with respect to λ , note that (11) gives

$$T^{PW} = \frac{1}{\lambda N} \log \left(\frac{p_0}{1-p_0} \left(\frac{\lambda \bar{w}}{c} - 1 \right) \right),$$

and thus,

$$\frac{\partial T^{PW}}{\partial \lambda} = \frac{\lambda \bar{w} - (\lambda \bar{w} - c) \log \left(\frac{p_0(\lambda \bar{w} - c)}{c(1-p_0)} \right)}{\lambda^2 N (\lambda \bar{w} - c)},$$

where the logarithm in the numerator is non-negative because $p_0 \lambda \bar{w} \geq c$. Hence, $\frac{\partial T^{PW}}{\partial \lambda}$ is positive for λ small enough (i.e., $p_0 \lambda \bar{w} \approx c$) and negative if λ is large.

A.2. Proof of Proposition 2

The first part of the proof establishes that an ES contest is optimal among hidden contests. The second part of the proof establishes the comparative statics.

We begin with the following lemma:

Lemma 1. *Take any symmetric equilibrium of a hidden contest with prize scheme $w(s_i, \mathbf{s}_{-i})$ and optimal deadline T . If this equilibrium does not have full effort by all agents from 0 to T , there is another scheme $w'(\cdot)$ with an optimal deadline T' that has a symmetric equilibrium in which each agent exerts full effort (so long as he has not succeeded) from 0 until T' , and where the aggregate cumulative effort is the same as under scheme $w(\cdot)$ and deadline T .*

Proof. First, without loss, take $T = \sup \{t : a_{i,t} > 0\}$. Next, suppose that each agent's effort $a_{i,t}$ is not constantly 1 over $[0, T]$. Let each agent's private belief at T be $p_{i,T}$. We choose a sub-interval $[0, T']$ such that each agent's private belief at T' conditional on no success before T' and all agents exerting full effort for the whole sub-interval is $p'_{i,T'} = p_{i,T}$. We find a prize scheme $w'(\cdot)$ such that exerting full effort for the whole sub-interval $[0, T']$ is a Nash equilibrium.

To this end, define a function $\tau : [0, T] \rightarrow \mathbb{R}_+$ by

$$\tau(z) = \int_0^z a_{i,t} dt. \quad (21)$$

Note that τ is weakly increasing. Take $T' = \tau(T)$. By convention, for any $t \in [0, T']$, we let $\tau^{-1}(t) = \inf\{z : \tau(z) = t\}$. Denote by $p'_{i,t}$ the private belief at time $t \in [0, T']$ under full effort. It is straightforward that for $t \in [0, T']$, $p'_{i,t} = p_{i,\tau^{-1}(t)}$.

We find $w'(\cdot)$ such that for any t' , agent i 's payoff by following the new equilibrium, $a'_{i,t} = 1$, over $[t', T']$ is the same as his payoff from following the old equilibrium over $[\tau^{-1}(t'), T]$ under $w(\cdot)$. The latter payoff is:

$$p_{i,\tau^{-1}(t')} \int_{\tau^{-1}(t')}^T (w_{i,t} \lambda - c) a_{i,t} e^{-\int_{\tau^{-1}(t')}^t \lambda a_{i,z} dz} dt - (1 - p_{i,\tau^{-1}(t')}) c \int_{\tau^{-1}(t')}^T a_{i,t} dt,$$

where $w_{i,t}$ is the expected reward if agent i succeeds at t given scheme $w(\cdot)$, i.e. suppressing the dependence on equilibrium strategies,

$$\begin{aligned} w_{i,t} &= \mathbb{E}_{s_{-i}} w(t, \mathbf{s}_{-i}) \\ &= \int_{[0,T]^{N-1}} w(t, s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_N) \left(\prod_{j \neq i} a_{j,s_j} \right) \lambda^{N-1} e^{-\lambda \sum_{j \neq i} \int_0^{s_j} a_{j,z} dz} d\mathbf{s}_{-i}. \end{aligned}$$

The former payoff is:

$$p_{i,t'} \int_{t'}^{T'} (w'_{i,t} \lambda - c) a'_{i,t} e^{-\int_{t'}^t \lambda a'_{i,z} dz} dt - (1 - p_{i,t'}) c \int_{t'}^{T'} a'_{i,t} dt,$$

where we find scheme $w'(\cdot)$ such that $w'_{i,t} = w_{i,\tau^{-1}(t)}$:

$$\begin{aligned} w'_{i,t'} &= \mathbb{E}_{s'_{-i}} w'(t', \mathbf{s}'_{-i}) \\ &= \int_{[0,T']^{N-1}} w'(t', \mathbf{s}'_{-i}) \lambda^{N-1} e^{-\lambda \sum_{j \neq i} \int_0^{s'_j} a_{j,z} dz} d\mathbf{s}'_{-i} \end{aligned}$$

$$\begin{aligned}
&= \int_{[0,T]^{N-1}} w'(\tau(t), \tau(\mathbf{s}'_{-i})) \left(\prod_{j \neq i} \tau'(s_j) \right) \lambda^{N-1} e^{-\lambda \sum_{j \neq i} \tau(s_j)} d\mathbf{s}_{-i} \\
&= \int_{[0,T]^{N-1}} w'(\tau(t), \tau(\mathbf{s}'_{-i})) \left(\prod_{j \neq i} a_{j,s_j} \right) \lambda^{N-1} e^{-\lambda \sum_{j \neq i} \tau(s_j)} d\mathbf{s}_{-i},
\end{aligned}$$

where $\tau(\mathbf{s}'_{-i}) := (\tau(s_1), \dots, \tau(s_{i-1}), \tau(s_{i+1}), \dots, \tau(s_N))$. We want to find $w'(\cdot)$ such that for any i and t ,

$$\begin{aligned}
&\int_{[0,T]^{N-1}} w'(\tau(t), \tau(\mathbf{s}_{-i})) \left(\prod_{j \neq i} a_{j,s_j} \right) \lambda^{N-1} e^{-\lambda \sum_{j \neq i} \tau(s_j)} d\mathbf{s}_{-i} \\
&= \int_{[0,T]^{N-1}} w(t, \mathbf{s}_{-i}) \left(\prod_{j \neq i} a_{j,s_j} \right) \lambda^{N-1} e^{-\lambda \sum_{j \neq i} \int_0^{s_j} a_{j,z} dz} d\mathbf{s}_{-i}. \tag{22}
\end{aligned}$$

Note that $\tau(s_j) = \int_0^{s_j} a_{j,z} dz$. Hence, (22) is equivalent to

$$\int_{[0,T]^{N-1}} [w'(\tau(t), \tau(\mathbf{s}_{-i})) - w(t, \mathbf{s}_{-i})] \left(\prod_{j \neq i} a_{j,s_j} \right) e^{-\lambda \sum_{j \neq i} \int_0^{s_j} a_{j,z} dz} d\mathbf{s}_{-i} = 0.$$

It thus suffices that for any t and \mathbf{s}_{-i} ,

$$w'(\tau(t), \tau(s_1), \dots, \tau(s_{i-1}), \tau(s_{i+1}), \dots, \tau(s_N)) = w(t, s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_N).$$

This defines the contest $w'(\cdot)$; note that because τ is a continuous function with image $[0, T']$ (see (22)), $w'(\cdot)$ has been defined over $[0, T']^N$.

We now argue that the full-effort profile, $a'_{i,z}$, is optimal (i.e. a best response). Suppose, per contra, that there is a strictly better strategy $\hat{a}_{i,z}$. Let us show that there is a profitable deviation from the original equilibrium. For any $z \in [0, T]$, define

$$\tilde{a}_i^z := \int_0^{\tau(z)} \hat{a}_{i,t} dt \tag{23}$$

as the total amount of effort that agent i should exert by time z .

Note that

$$\tilde{a}_i^z = \int_0^z \tilde{a}_{i,t} dt, \tag{24}$$

where (21) and (23) imply that, almost everywhere,

$$\tilde{a}_{i,z} := \frac{d\tilde{a}_i^z}{dz} = \frac{d\tau(z)}{dz} \hat{a}_{i,\tau(z)} = a_{i,z} \hat{a}_{i,\tau(z)} \in [0, 1].$$

It follows from (23) and (24) that

$$\text{for any } t', t \in [0, T']: \int_{t'}^t \hat{a}_{i,z} dz = \int_{\tau^{-1}(t')}^{\tau^{-1}(t)} \tilde{a}_{i,z} dz. \quad (25)$$

We claim that (25) implies that the payoff at any $t \in [0, T']$ under $\hat{a}_{i,z}$ and at $\tau^{-1}(t) \in [0, T]$ under $\tilde{a}_{i,z}$ are the same. This follows from the same argument as before, because we have used the same $\tau(\cdot)$ as before and hence, at any z , the total amount of effort by agents $-i$ at z is also the same as that at $\tau^{-1}(z)$, and, by construction, agent i 's private belief and total effort at z is the same as that at $\tau^{-1}(z)$.

Since $\hat{a}_{i,z}$ is a profitable deviation from $a'_{i,z}$, we conclude that $\tilde{a}_{i,z}$ is a profitable deviation from $a_{i,z}$, a contradiction. *Q.E.D.*

The proof that an ES contest is optimal among hidden contests proceeds in two steps. Consider a hidden contest with prize scheme $w(s_i, s_{-i})$ and associated optimal deadline T as defined in Lemma 1. Given full effort from 0 to T (as is without loss by Lemma 1), the contest induces a sequence of expected rewards for success at each time $t \in [0, T]$ as shown in the proof of Lemma 1; denote this sequence for agent i by $\{w_{i,t}\}$.

Step 1. We show that in an optimal hidden ES contest, each agent's incentive constraint for effort is binding at each time $t \in [0, T^{HS}]$.

An agent i 's continuation payoff at any time $t' \in [0, T]$ is

$$U_{t'} := p_{i,t'} \int_{t'}^T (w_{i,t} \lambda - c) a_{i,t} e^{-\int_{t'}^t \lambda a_{i,z} dz} dt - (1 - p_{i,t'}) c \int_{t'}^T a_{i,t} dt.$$

Consider the continuation payoff at t' from a strategy $\tilde{a}_{i,t} = 1 \forall t \in [0, T] \setminus (t', t' + \varepsilon)$ and 0 otherwise:

$$U_{t'}(\varepsilon) := p_{i,t'} \int_{t'+\varepsilon}^T (w_{i,t} \lambda - c) e^{-(t-t'-\varepsilon)\lambda} dt - (1 - p_{i,t'}) c (T - t' - \varepsilon).$$

We compute

$$\begin{aligned} U_{t'}'(\varepsilon) &= p_{i,t'} \left[-(w_{i,t'+\varepsilon} \lambda - c) + \int_{t'+\varepsilon}^T (w_{i,t} \lambda - c) e^{-(t-t'-\varepsilon)\lambda} \lambda dt \right] + (1 - p_{i,t'}) c, \\ U_{t'}'(0) &= -(p_{i,t'} w_{i,t'} \lambda - c) + p_{i,t'} \int_{t'}^T (w_{i,t} \lambda - c) e^{-(t-t')\lambda} \lambda dt. \end{aligned}$$

Note that because of Nash equilibrium, $U_{t'}(0) \geq U_{t'}(\varepsilon)$ for any $\varepsilon \geq 0$ and thus $U_{t'}'(0) \leq 0$.

If $w_{i,t} = w_i$ for all t , the expression above simplifies to

$$\begin{aligned}
U'_{t'}(0) &= -(p_{i,t'} w_i \lambda - c) + p_{i,t'} \int_{t'}^T (w_i \lambda - c) e^{-(t-t')\lambda} \lambda dt \\
&= -(p_{i,t'} w_i \lambda - c) + p_{i,t'} (w_i \lambda - c) \left(1 - e^{-(T-t')\lambda}\right) \\
&= c(1 - p_{i,t'}) - p_{i,t'} (w_i \lambda - c) e^{-(T-t')\lambda}.
\end{aligned}$$

In an optimal hidden ES contest, $T = T^{HS}$, $w_i = w^{HS}$, and $p_{i,T^{HS}} w^{HS} \lambda = c$; hence

$$p_{i,t'} e^{-(T^{HS}-t')\lambda} (w^{HS} \lambda - c) = c(1 - p_{i,t'}).$$

Therefore, we obtain $U'_{t'}(0) = 0$, showing that each agent's incentive constraint is binding at each time $t' \leq T^{HS}$ in an optimal hidden ES contest.

Step 2. We show that any hidden contest (with full effort up to a deadline T) is weakly dominated by hidden ES. The incentive constraint for agent i at any $z \in [0, T]$ implies

$$U'_z(0) = -(p_{i,z} w_{i,z} \lambda - c) + p_{i,z} \lambda \int_z^T (w_{i,t} \lambda - c) e^{-(t-z)\lambda} dt \leq 0. \quad (26)$$

Suppose, to contradiction, $T > T^{HS}$. For any $z \leq T^{HS} < T$, rewrite (26) as

$$U'_z(0) = -(p_{i,z} w_{i,z} \lambda - c) + p_{i,z} \lambda \left[\int_z^{T^{HS}} (w_{i,t} \lambda - c) e^{-(t-z)\lambda} dt + \int_{T^{HS}}^T (w_{i,t} \lambda - c) e^{-(t-z)\lambda} dt \right] \leq 0. \quad (27)$$

By Step 1, it holds that in an optimal hidden ES contest, $U'_z(0) = 0$ at any time $z \leq T^{HS}$. Subtracting this binding constraint from (27), we obtain for any $z \leq T^{HS}$:

$$-p_{i,z} (w_{i,z} - w^{HS}) \lambda + p_{i,z} \lambda \left[\int_z^{T^{HS}} (w_{i,t} - w^{HS}) \lambda e^{-(t-z)\lambda} dt + \int_{T^{HS}}^T (w_{i,t} \lambda - c) e^{-(t-z)\lambda} dt \right] \leq 0.$$

Define

$$t' = \sup \left\{ t : \left\{ z \in \left(t - \frac{1}{m}, t \right) : w_{i,z} \leq w^{HS} \right\} \text{ has positive Lebesgue measure for any } m > 0 \right\}.$$

Such a t' is well-defined because otherwise $w_{i,t} > w^{HS}$ almost everywhere, which is not possible given a prize of \bar{w} .

Consider first the case where $t' \leq T^{HS}$. Then for any m , consider any $z \in (t' - \frac{1}{m}, t')$

such that $w_{i,z} \leq w^{HS}$. Define

$$\begin{aligned} \Delta_z &:= -p_{i,z} (w_{i,z} - w^{HS}) \lambda + p_{i,z} \lambda \int_z^{t'} (w_{i,t} - w^{HS}) \lambda e^{-(t-z)\lambda} dt \\ &\quad + p_{i,z} \lambda \int_{t'}^{T^{HS}} (w_{i,t} - w^{HS}) \lambda e^{-(t-z)\lambda} dt + p_{i,z} \lambda \int_{T^{HS}}^T (w_{i,t} \lambda - c) e^{-(t-z)\lambda} dt. \end{aligned}$$

Note that $-p_{i,z} (w_{i,z} - w^{HS}) \lambda \geq 0$ and $\int_{t'}^{T^{HS}} (w_{i,t} - w^{HS}) \lambda e^{-(t-z)\lambda} dt \geq 0$ by the definition of t' . Moreover,

$$\left| \int_z^{t'} (w_{i,t} - w^{HS}) \lambda e^{-(t-z)\lambda} dt \right| \leq 2\lambda \bar{w} (t' - z)$$

because the prize being \bar{w} implies $|w_{i,t} - w^{HS}| \leq 2\bar{w}$. Note $\gamma := \int_{T^{HS}}^T (w_{i,t} \lambda - c) e^{-(t-z)\lambda} dt > 0$, because the incentive constraint implies that for any t , $p_{i,t} w_{i,t} \lambda - c \geq 0$, and $p_{i,t} < 1$. Now take $\frac{1}{m} < \frac{\gamma}{2\lambda \bar{w}}$. Then $\Delta_z > 0$, a contradiction.

Consider next the case where $t' > T^{HS}$. Then for m sufficiently large and $z \in (t' - \frac{1}{m}, t')$ such that $w_{i,z} \leq w^{HS}$,

$$U'_z(0) = -(p_{i,z} w_{i,z} \lambda - c) + p_{i,z} \lambda \int_z^T (w_{i,t} \lambda - c) e^{-(t-z)\lambda} dt > 0.$$

The first term is strictly positive because $p_{i,T^{HS}} w^{HS} \lambda - c = 0$, $p_{i,z} < p_{i,T^{HS}}$, and $w_{i,z} \leq w^{HS}$. The second term is strictly positive for the same reasons as before. Thus, again, we reach a contradiction.

Comparative statics. Having established optimality of an ES contest among hidden contests, it remains to prove the proposition's comparative statics.

Step 3. For any $\kappa \in [0, 1]$ and $N \geq 2$, we claim $\sum_{n=0}^{N-1} \left(\frac{1}{n+1}\right) \binom{N-1}{n} (1-\kappa)^n \kappa^{N-1-n} = \frac{1-\kappa^N}{(1-\kappa)N}$, which immediately implies that the summation in (15) is equivalent to $\frac{1-e^{-N\lambda T^{HS}}}{(1-e^{-\lambda T^{HS}})N}$ for any $N \geq 1$. The claim is verified as follows:

$$\begin{aligned} \sum_{n=0}^{N-1} \left(\frac{1}{n+1}\right) \binom{N-1}{n} (1-\kappa)^n \kappa^{N-1-n} &= \sum_{n=0}^{N-1} \frac{(N-1)! (1-\kappa)^n \kappa^{N-1-n}}{(N-1-n)! n! (n+1)} \\ &= \frac{1}{(1-\kappa)N} \sum_{n=0}^{N-1} \binom{N}{n+1} (1-\kappa)^{n+1} \kappa^{N-(n+1)} \\ &= \frac{1}{(1-\kappa)N} \left(1 - \binom{N}{0} (1-\kappa)^0 \kappa^N\right) = \frac{1-\kappa^N}{(1-\kappa)N}. \end{aligned}$$

Step 4. Letting $\kappa := e^{-\lambda T^{HS}} \in (0, 1)$ and $q := \frac{1-p_0}{p_0} > 0$, rewrite (16) as

$$\frac{1 - \kappa^N}{(1 - \kappa)N} \frac{\kappa}{\kappa + q} = \frac{c}{\lambda \bar{w}}. \quad (28)$$

We show that the LHS of (28) is increasing in κ . We compute $\frac{\partial \left(\frac{1 - \kappa^N}{(1 - \kappa)N} \right)}{\partial \kappa} = \frac{1 - \kappa^N - N\kappa^{N-1}(1 - \kappa)}{(1 - \kappa)^2 N}$, which is positive if and only if

$$1 \geq N\kappa^{N-1} - (N - 1)\kappa^N. \quad (29)$$

Differentiation shows that the RHS of (29) is increasing in κ since $\kappa \in (0, 1)$. As (29) holds with $\kappa = 1$, it follows that $\frac{\partial \left(\frac{1 - \kappa^N}{(1 - \kappa)N} \right)}{\partial \kappa} > 0$. Moreover, $\frac{\kappa}{\kappa + q}$ is also increasing in κ because $q > 0$; hence the LHS of (28) is increasing in κ .

A similar argument establishes that the LHS of (28) is decreasing in N .

Step 5. Using Step 4, we now derive the comparative statics of T^{HS} . As the LHS of (28) is increasing in κ and κ is decreasing in T^{HS} , the LHS of (28) is decreasing in T^{HS} . Moreover, (i) the RHS of (28) is increasing in c and decreasing in \bar{w} while the LHS is independent of these parameters, and (ii) q is decreasing in p_0 and thus the LHS of (28) is increasing in p_0 , while the RHS is independent of this parameter. We thus obtain that T^{HS} is increasing in p_0 and \bar{w} and decreasing in c . Similarly, T^{HS} is decreasing in N because the LHS of (28) is decreasing in both N and T^{HS} while the RHS is independent of both parameters.

Lastly, we show that T^{HS} is non-monotonic with respect to λ by providing an example. Let $p_0 = \frac{1}{2}$ and $N = 2$. Then (16) becomes

$$\frac{1 - e^{-2\lambda T^{HS}}}{2(1 - e^{-\lambda T^{HS}})} \frac{e^{-\lambda T^{HS}}}{e^{-\lambda T^{HS}} + 1} = \frac{c}{\lambda \bar{w}},$$

which simplifies to $T^{HS} = \frac{1}{\lambda} \log \left(\frac{\lambda \bar{w}}{2c} \right)$. Differentiating, $\frac{\partial T^{HS}}{\partial \lambda} = \frac{1}{\lambda^2} \left(1 - \log \left(\frac{\lambda \bar{w}}{2c} \right) \right)$, which is positive for small λ (i.e., $p_0 \lambda \bar{w} \approx c$) and negative for large λ .

Step 6. We can now show the comparative statics for the probability of obtaining a success. The probability of success in an optimal hidden ES contest is given by expression (5) with $A^{T^{HS}} = NT^{HS}$. The comparative static with respect to c is immediate: as shown in Step 5, if c increases, T^{HS} decreases, which implies that NT^{HS} and thus the probability of success decreases. An analogous argument shows that the probability of success is increasing in \bar{w} and p_0 .

We next show that the probability of success increases with λ . From Step 5, T^{HS} may increase or decrease when λ increases. However, note that λT^{HS} must increase when λ increases: if λT^{HS} decreases, the LHS of (16) increases, while the RHS decreases when λ increases, leading to a contradiction. Therefore, λT^{HS} increases with λ , implying that the probability of obtaining a success increases with λ .

Finally, the ambiguous effect of an increase in N on the probability of success is seen through the example reported in Figure 1.

A.3. Proof of Proposition 3

Proposition 1 shows that in an optimal public contest, agents follow a stopping strategy with stopping time T^{PW} given by (11). Proposition 2 shows that in an optimal hidden contest, agents follow a stopping strategy with stopping time T^{HS} given by (16). Hence, the principal strictly prefers an optimal hidden contest to an optimal public contest if and only if $T^{HS} > T^{PW}$. Using (11) and (16), this condition is equivalent to (17), as discussed in the text before Proposition 3.

A.4. Proof of Proposition 4

For any $n \in \{0, \dots, N\}$ and time t , we will write $n_t = n$ as shorthand for $\{o^t : \eta(o^t) = n\}$, the event that n agents have succeeded by t . Similarly, we will write $n_t < n$ as shorthand for $\{o^t : \eta(o^t) < n\}$, the event that strictly less than n agents have succeeded by t .

Consider first a cutoff-disclosure ES contest with cutoff n^* and a large deadline T . By agent i following a stopping strategy with stopping time t in such a contest, we mean agent i exerting effort $a_{i,z} = 1$ if $z \leq t$ and the principal has sent message $m_\ell = 0$ at all $\ell \leq z$ (and the agent has not succeeded by z), and exerting effort $a_{i,z} = 0$ otherwise. The stopping time T^* is defined by the static incentive constraint for effort at T^* binding when all agents follow stopping strategies with stopping time T^* :

$$\sum_{n \in \{0, \dots, n^* - 1\}} \Pr [n_{T^*} = n | n_{T^*} < n^*] \Pr [G | n_{T^*} = n] \lambda \frac{\bar{w}}{n+1} = c. \quad (30)$$

Three points bear clarification: (i) since the principal only cares about efforts in a history of no success, we have suppressed above (and will suppress elsewhere) that agent i conditions on him not having succeeded; (ii) we have suppressed (and will suppress elsewhere) that agent i conditions on his conjectures about opponents' effort and on his own effort; (iii) the only information disclosed by the principal along the path in a history of no success is that $n_{T^*} < n^*$.

Lemma 2. *In a cutoff-disclosure ES contest with cutoff n^* and a large enough deadline T , there is a symmetric equilibrium in stopping strategies with stopping time T^* defined by (30). If the deadline is chosen to be T^* , this equilibrium is unique.*

Proof. We first show that there is a symmetric equilibrium in stopping strategies. To prove this, it is sufficient to show that the set of best responses to stopping strategies always includes a stopping strategy (and then standard arguments imply existence).

Fixing any profile of stopping strategies for an agent i 's opponents, consider the following auxiliary problem. Let τ denote the random time at which the number of opponents who have succeeded first equals n^* , assuming each opponent simply works throughout according to what he was supposed to given the 0 (null) message. By convention, $\tau = \infty$ if there is no time by which n^* opponents succeed. If agent i succeeds at $t < T$, the event in which agent i is among the first n^* agents who succeed before T is $\{\tau > t\}$. Therefore, in the original problem, if agent i is constrained to exert cumulative effort α , a strategy $(a_{i,t})_{t \in [0, T]}$ is optimal if and only if it solves

$$\max_{(a_{i,t})_{t \geq 0} \text{ s.t. } \int_0^T a_{i,t} dt = \alpha} p_0 \int_0^T \Pr(\tau > t) (\mathbb{E}[w_i(t, \tau) | \tau > t] - c) a_{i,t} e^{-\lambda a_i^t} dt - (1 - p_0)c\alpha, \quad (31)$$

where $w_i(t, \tau)$ is i 's expected reward for succeeding at t as a function of τ . To interpret this objective, recognize that agent i will optimally stop exerting effort when it is disclosed that n^* opponents have succeeded or when he himself has succeeded, and thus any effort that is planned after τ arrives can be treated as having zero payoff. Note that $\Pr(\tau > t) (\mathbb{E}[w_i(t, \tau) | \tau > t] - c)$ can be rewritten as

$$\Pr(\tau > T) (\mathbb{E}[w_i(t, \tau) | \tau > T] - c) + \Pr(t < \tau \leq T) (\mathbb{E}[w_i(t, \tau) | t < \tau \leq T] - c),$$

where $\mathbb{E}[w_i(t, \tau) | t < \tau \leq T] = \lambda \bar{w} / n^*$ is independent of t and τ , and

$$\mathbb{E}[w_i(t, \tau) | \tau > T] = \mathbb{E} \left[\frac{\bar{w}}{n+1} \mid n < n^* \text{ opponents succeed by } T \right]$$

only depends on the opponents' strategy profile. Since $\Pr(t < \tau \leq T)$ is non-increasing in t , it follows that $\Pr(\tau > t) (\mathbb{E}[w_i(t, \tau) | \tau > t] - c)$ is non-increasing in t . Hence, a stopping strategy is a solution to problem (31).

By the previous claim, the existence of a symmetric equilibrium in stopping strategies follows if no agent wants to deviate to a stopping strategy with different cumulative effort. Equation (30) implies that no agent wants more nor less cumulative effort.

Finally, to show the uniqueness claim, suppose, per contra, that T^* is the deadline and there is an equilibrium in which an agent i 's cumulative effort conditional on no success is $T_i < T^*$. The same argument as in Lemma 1 implies that there is another equilibrium in which agent i uses a stopping strategy with stopping time T_i . By the definition of T^* in (30), agent i strictly prefers to work at T^* , a contradiction. Q.E.D.

Consider now the stopping-strategy equilibrium with stopping time T^* in a cutoff-disclosure ES contest with cutoff n^* . For any t , $\Pr[G|n_t = n] = 1$ for all $n \geq 1$ and, by definition, $p_{i,t} = \Pr[G|n_t = 0]$ is the probability that agent i assigns to the good state conditional on no agent having succeeded by time t . Letting $p^* := p_{i,T^*}$, we can accordingly rewrite (30) as

$$\Pr[n_{T^*} = 0 | n_{T^*} < n^*] (p^* \lambda \bar{w} - c) + \sum_{n \in \{1, \dots, n^* - 1\}} \Pr[n_{T^*} = n | n_{T^*} < n^*] \left(\lambda \frac{\bar{w}}{n+1} - c \right) = 0. \quad (32)$$

Suppose, per contra, that the proposition is not true. Then there is a contest C with a rank-monotonic prize-sharing scheme and deadline \hat{T} (and with a deterministic and symmetric disclosure policy) that induces a symmetric equilibrium with aggregate cumulative effort absent success $A^{\hat{T}} > NT^*$. Since each agent's cumulative effort is continuous and starts at 0, there is a time $\hat{t} \in [T^*, \hat{T}]$ such that $A^{\hat{t}} = NT^*$ and $A^{\hat{t} + \varepsilon} > NT^*$ for all $\varepsilon > 0$ small enough. Given a history of no success (which is an event with strictly positive probability), each agent i exerts strictly positive effort in the neighborhood $(\hat{t}, \hat{t} + \varepsilon)$; by continuity, agent i 's dynamic incentive constraint for effort must be satisfied at \hat{t} . Moreover, that the agent's dynamic incentive constraint is satisfied and the agent exerts strictly positive effort after \hat{t} implies the static incentive constraint at \hat{t} holds strictly:

$$\sum_{n \in \{0, \dots, N-1\}} \Pr[n_{\hat{t}} = n | m_{\emptyset}^{\hat{t}}] \Pr[G | n_{\hat{t}} = n] \left(\lambda w_{i,\hat{t}}^n - c \right) > 0, \quad (33)$$

where $m_{\emptyset}^{\hat{t}}$ denotes the history of messages that agent i receives up to \hat{t} when there is no success and $w_{i,\hat{t}}^n$ is the agent's equilibrium expected reward for succeeding at time \hat{t} (computed at such time) given $n_{\hat{t}} = n$ and the message history $m_{\emptyset}^{\hat{t}}$. (To reduce notation, we do not explicitly include the dependence of $w_{i,\hat{t}}^n$ on $m_{\emptyset}^{\hat{t}}$.) Note that we need not condition on $m_{\emptyset}^{\hat{t}}$ in the $\Pr[G | n_{\hat{t}} = n]$ term because it is redundant there given the conditioning on $n_{\hat{t}} = n$; also, we have again suppressed the dependence of probabilities on equilibrium efforts.

By rank monotonicity of the prize scheme in contest C , $w_{i,\hat{t}}^n \leq \frac{\bar{w}}{n+1}$ (as agent i 's reward

for success at \hat{t} cannot be larger than that of any of his opponents who succeeded prior to \hat{t} for any profile of successes); thus, (33) implies

$$\sum_{n \in \{0, \dots, N-1\}} \Pr \left[n_{\hat{t}} = n | m_{\varnothing}^{\hat{t}} \right] \Pr \left[G | n_{\hat{t}} = n \right] \left(\lambda \frac{\bar{w}}{n+1} - c \right) > 0. \quad (34)$$

Since \hat{t} is defined as a time at which $A^{\hat{t}} = NT^*$, it holds that $\Pr[G | n_{\hat{t}} = 0] = p^*$ in contest C , and hence, analogous to (32), we can rewrite constraint (34) as

$$\Pr \left[n_{\hat{t}} = 0 | m_{\varnothing}^{\hat{t}} \right] (p^* \lambda \bar{w} - c) + \sum_{n \in \{1, \dots, N-1\}} \Pr \left[n_{\hat{t}} = n | m_{\varnothing}^{\hat{t}} \right] \left(\lambda \frac{\bar{w}}{n+1} - c \right) > 0. \quad (35)$$

We seek a contradiction between (35) and (32). To this end, note that since $\lambda \frac{\bar{w}}{n+1} - c < 0$ for $n \geq n^*$, (35) implies

$$\Pr \left[n_{\hat{t}} = 0 | m_{\varnothing}^{\hat{t}}, n_{\hat{t}} < n^* \right] (p^* \lambda \bar{w} - c) + \sum_{n \in \{1, \dots, n^*-1\}} \Pr \left[n_{\hat{t}} = n | m_{\varnothing}^{\hat{t}}, n_{\hat{t}} < n^* \right] \left(\lambda \frac{\bar{w}}{n+1} - c \right) > 0. \quad (36)$$

Consider the static incentive constraint in the cutoff-disclosure ES contest with cutoff n^* . Let $E \subset \{0, \dots, N\}$ be any subset that contains 0, and write $n_{T^*} \in E$ to denote the event in which the number of agents who have succeeded by T^* is in E . Since $\lambda \frac{\bar{w}}{n+1} - c \geq 0$ for $n < n^*$, (32) implies

$$\begin{aligned} & \Pr \left[n_{T^*} = 0 | n_{T^*} \in E, n_{T^*} < n^* \right] (p^* \lambda \bar{w} - c) \\ & + \sum_{n \in \{1, \dots, n^*-1\}} \Pr \left[n_{T^*} = n | n_{T^*} \in E, n_{T^*} < n^* \right] \left(\lambda \frac{\bar{w}}{n+1} - c \right) \leq 0. \end{aligned} \quad (37)$$

Since the disclosure policy in C is deterministic and symmetric, the event $m_{\varnothing}^{\hat{t}}$ in (36) corresponds to a set of success counts; note that this set includes 0. As E in (37) is arbitrary, we can take the event $n_{T^*} \in E$ in (37) to be that corresponding to $m_{\varnothing}^{\hat{t}}$. Then, by definition of \hat{t} (i.e., $A^{\hat{t}} = NT^*$) and our restriction to symmetric equilibria, it follows that for any $n \in \{0, \dots, n^*-1\}$, $\Pr \left[n_{T^*} = n | n_{T^*} \in E, n_{T^*} < n^* \right]$ in the cutoff-disclosure ES contest with cutoff n^* is equal to $\Pr \left[n_{\hat{t}} = n | m_{\varnothing}^{\hat{t}}, n_{\hat{t}} < n^* \right]$ in contest C . Thus, the inequalities (36) and (37) are in contradiction.

A.5. Proof of Proposition 5

Consider a WTA contest $C = (w(\cdot), D, T)$, where $w(\cdot)$ is the WTA prize scheme, D is an arbitrary information disclosure policy, and T is the deadline. Denote by $I_{i,t}$ the information that the principal has disclosed to agent i by time t ; if D is stochastic, let $I_{i,t}$ correspond to any given realization of D . Denote agent i 's belief at time t by $p_{i,t}$ and his expected reward for success at time t by $w_{i,t}$. We let (agent i 's conjecture of) the aggregate cumulative effort up to t given no success by t be A^t .

At any time t , agent i can ensure a positive continuation payoff by exerting no effort. The agent thus chooses $a_{i,t} > 0$ only if

$$c \leq p_{i,t} \lambda w_{i,t}. \quad (38)$$

Let

$$\alpha_{i,t} := \Pr[\text{some } j \neq i \text{ succeeded by } t \mid i \text{ did not, } I_{i,t}] \in [0, 1].$$

In a WTA contest, (38) is equivalent to

$$c \leq (1 - \alpha_{i,t}) \Pr[G \mid \text{no success, } A^t] \lambda \bar{w}.$$

Consider now an optimal public WTA contest. As shown in Section 4, agent i chooses $a_{i,t} = 1$ if

$$c \leq \Pr[G \mid \text{no success, } A^t] \lambda \bar{w}.$$

It follows that, given aggregate cumulative effort A^t , if agent i exerts positive effort in contest C , he exerts full effort in an optimal public WTA contest. Therefore, if the aggregate cumulative effort induced by contest C is A^T and that induced by an optimal public WTA contest is A^{PW} , then $A^{PW} \geq A^T$.

A.6. Proof of Proposition 6

The principal chooses \bar{w} and a rank-monotonic contest to maximize (7). By Proposition 4 and its proof, we can take the optimal contest to be a cutoff-disclosure ES contest with cutoff n^* , in which agents follow a stopping strategy with stopping time defined by (30). Let $A^T(\bar{w})$ be the aggregate cumulative effort in this contest under prize \bar{w} . $A^T(\cdot)$ is continuous and unbounded. When v is large enough, an optimal prize must induce arbitrarily large aggregate cumulative effort, and hence $\lambda \bar{w}/N > c$ at an optimum. On the other hand, $\lambda \bar{w}/2 < c$ when v is small enough because an optimal prize must satisfy

$\bar{w} < v$. The result now follows from [Corollary 1](#).³⁰

A.7. Details for [Section 2](#)

We provide here details for the claims made in the illustrative setting of [Section 2](#). First, we show that there are parameters $(p_0, \lambda, c, \bar{w})$ under which inequality (1) holds while (2) does not hold. This establishes our claim that there are parameters under which a public WTA contest improves on hidden WTA.

In a hidden contest, an agent's belief that the state is good in the second period given that he has failed in the first period (having exerted effort) is $p_1^{(1)} := \frac{p_0(1-\lambda)}{p_0(1-\lambda)+1-p_0}$. Since $\Pr[j \text{ failed} \mid i \text{ failed}] = 1 - p_1^{(1)}\lambda$, we can rewrite inequality (2) as

$$(1 - p_1^{(1)}\lambda)p_1\lambda\hat{w} \geq c.$$

For (1) to hold while (2) does not, parameters must be such that

$$p_1\lambda\hat{w} \geq c > (1 - p_1^{(1)}\lambda)p_1\lambda\hat{w}. \quad (39)$$

Here is an example: $p_0 = \lambda = \frac{1}{2}$, $c = \frac{1}{15}$, and $\bar{w} = 1$. Then $\hat{w} = \frac{3}{4}$, $p_1 = \frac{1}{5}$, and $p_1^{(1)} = \frac{1}{3}$, so (39) is satisfied: $\frac{3}{40} > \frac{1}{15} > \frac{1}{16}$.

Next, we show that there are parameters $(p_0, \lambda, c, \bar{w})$ under which $p_0\lambda\hat{w} > c$ and inequality (3) holds while (1) does not hold. This establishes our claim that there are parameters under which a hidden ES contest improves on public WTA.

We can rewrite inequality (3) as

$$p_1^{(1)}\lambda^2\frac{\bar{w}}{2} + (1 - p_1^{(1)}\lambda)p_1\lambda\hat{w} \geq c.$$

For (3) to hold while (1) does not, parameters must be such that

$$p_1^{(1)}\lambda^2\frac{\bar{w}}{2} + (1 - p_1^{(1)}\lambda)p_1\lambda\hat{w} \geq c > p_1\lambda\hat{w}. \quad (40)$$

There are parameters for which $p_0\lambda\hat{w} > c$ and the inequalities in (40) hold. For example, take the same parameter values as above but with $c = \frac{1}{10}$. Then $p_0\lambda\hat{w} = \frac{3}{16} > \frac{1}{10}$ and (40) is satisfied: $\frac{5}{48} > \frac{1}{10} > \frac{3}{40}$.

³⁰ Recall that no experimentation can be induced if $p_0\lambda\bar{w} \leq c$. Thus, for sufficiently small v , no contest can incentivize experimentation; however, so long as $p_0 > 1/2$, there will be parameters under which public WTA is optimal while still inducing experimentation.

References

- Abrams, David and R. Polk Wagner. 2013. "Poisoning the Next Apple? The America Invents Act and Individual Inventors." *Stanford Law Review* 65:517–564.
- Akcigit, Ufuk and Qingmin Liu. 2015. "The Role of Information in Innovation and Competition." *Journal of the European Economic Association* Forthcoming.
- Aoyagi, Masaki. 2010. "Information Feedback in a Dynamic Tournament." *Games and Economic Behavior* 70:242–260.
- Bergemann, Dirk and Ulrich Hege. 1998. "Venture Capital Financing, Moral Hazard, and Learning." *Journal of Banking & Finance* 22 (6-8):703–735.
- . 2005. "The Financing of Innovation: Learning and Stopping." *RAND Journal of Economics* 36 (4):719–752.
- Bhattacharya, Sudipto, Jacob Glazer, and David Sappington. 1990. "Sharing Productive Knowledge in Internally Financed R&D Contests." *Journal of Industrial Economics* 39 (2):187–208.
- Bimpikis, Kostas and Kimon Drakopoulos. 2014. "Disclosing Information in Strategic Experimentation." Manuscript, Graduate School of Business, Stanford University.
- Bimpikis, Kostas, Shayan Ehsani, and Mohamed Mostagir. 2014. "Information Provision in Dynamic Innovation Tournaments." Manuscript, Graduate School of Business, Stanford University.
- Bobtcheff, Catherine and Raphaël Levy. 2014. "More Haste, Less Speed: Signalling through Investment Timing." Manuscript, Toulouse School of Economics.
- Bolton, Patrick and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica* 67 (2):349–374.
- Bonatti, Alessandro and Johannes Hörner. 2011. "Collaborating." *American Economic Review* 101 (2):632–663.
- Campbell, Arthur, Florian Ederer, and Johannes Spinnewijn. 2014. "Delay and Deadlines: Freeriding and Information Revelation in Partnerships." *American Economic Journal: Microeconomics* 6 (2):163–204.
- Che, Yeon-Koo and Ian Gale. 2003. "Optimal Design of Research Contests." *American Economic Review* 93 (3):646–671.
- Che, Yeon-Koo and Johannes Hörner. 2015. "Optimal Design for Social Learning." Manuscript, Department of Economics, Columbia University.
- Choi, Jay P. 1991. "Dynamic R&D Competition under 'Hazard Rate' Uncertainty." *RAND Journal of Economics* 22 (4):596–610.

- Cripps, Martin and Caroline Thomas. 2014. "Strategic Experimentation in Queues." Manuscript, Department of Economics, University College London.
- Dasgupta, Partha and Joseph Stiglitz. 1980. "Uncertainty, Industrial Structure, and the Speed of R&D." *The Bell Journal of Economics* :1–28.
- Easley, David and Arpita Ghosh. 2013. "Incentives, Gamification, and Game Theory: An Economic Approach to Badge Design." In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 359–376.
- Ederer, Florian. 2010. "Feedback and Motivation in Dynamic Tournaments." *Journal of Economics & Management Strategy* 19 (3):733–769.
- Ely, Jeffrey. 2015. "Beeps." Manuscript, Department of Economics, Northwestern University.
- Emanuel, Ezekiel J. 2015. "How to Develop New Antibiotics." *The New York Times* February 24.
- Fullerton, Richard and R. Preston McAfee. 1999. "Auctioning Entry into Tournaments." *Journal of Political Economy* 107 (3):573–605.
- Fusco, Michael. 2012. "An Analysis of Multiple Award Indefinite Delivery Indefinite Quantity Contracts at the Army Contracting Command Aberdeen Proving Ground." Naval Postgraduate School, Thesis.
- Garicano, Luis and Andrea Prat. 2013. "Organizational Economics with Cognitive Costs." In *Advances in Economics and Econometrics, Tenth World Congress of the Econometric Society*, edited by D. Acemoglu, M. Arellano, and E. Dekel. Cambridge: Cambridge University Press, 342–388.
- Gill, David. 2008. "Strategic Disclosure of Intermediate Research Results." *Journal of Economics & Management Strategy* 17 (3):733–758.
- Goltsman, Maria and Arijit Mukherjee. 2011. "Interim Performance Feedback in Multi-stage Tournaments: The Optimality of Partial Disclosure." *Journal of Labor Economics* 29:229–265.
- Gomes, Renato, Daniel Gottlieb, and Lucas Maestri. 2015. "Experimentation and Project Selection: Screening and Learning." Manuscript, Toulouse School of Economics.
- Halac, Marina, Navin Kartik, and Qingmin Liu. 2016. "Optimal Contracts for Experimentation." Manuscript, Graduate School of Business, Columbia University.
- Heidhues, Paul, Sven Rady, and Philipp Strack. 2015. "Strategic Experimentation with Private Payoffs." *Journal of Economic Theory* 159 (5):531–551.
- Holmström, Bengt and Paul R. Milgrom. 1991. "Multi-Task Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, and Organization* 7 (Special Issue):24–52.

- Hörner, Johannes and Larry Samuelson. 2014. "Incentives for Experimenting Agents." *RAND Journal of Economics* 44 (4):632–663.
- Itoh, Hideshi. 1991. "Incentives to Help in Multi-Agent Situations." *Econometrica* 59 (3):611–636.
- Judd, Kenneth, Karl Schmedders, and Şevin Yeltekin. 2012. "Optimal Rules for Patent Races." *International Economic Review* 53 (1):23–52.
- Keller, Godfrey and Sven Rady. 2010. "Strategic Experimentation with Poisson Bandits." *Theoretical Economics* 5 (2):275–311.
- Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73 (1):39–68.
- Kremer, Ilan, Yishay Mansour, and Motty Perry. 2014. "Implementing the 'Wisdom of the Crowd'." *Journal of Political Economy* 122 (5):988–1012.
- Krishna, Vijay and John Morgan. 1998. "The Winner-Take-All Principle in Small Tournaments." In *Advances in Applied Microeconomics: Contests*, vol. 7, edited by Michael R. Baye. Jai Pr, 61–74.
- Lerner, Josh. 1995. "Patenting in the Shadow of Competitors." *Journal of Law and Economics* 38 (2):463–495.
- Loury, Glenn. 1979. "Market Structure and Innovation." *Quarterly Journal of Economics* 93 (3):395–410.
- MacCormack, Alan, Fiona Murray, and Erika Wagner. 2013. "Spurring Innovation through Competition." *MIT Sloan Management Review* 55 (1):25–32.
- Machlup, Fritz and Edith Penrose. 1950. "The Patent Controversy in the Nineteenth Century." *Journal of Economic History* 10 (1):1–29.
- Malueg, David and Shunichi Tsutsui. 1997. "Dynamic R&D Competition with Learning." *RAND Journal of Economics* 28 (4):751–772.
- Manso, Gustavo. 2011. "Motivating Innovation." *Journal of Finance* 66 (5):1823–1860.
- Maskin, Eric. 2002. "How to Reduce Greenhouse Gas Emissions: An Application of Auction Theory." Nancy L. Schwartz Memorial Lecture, Northwestern University.
- Mason, Robin and Juuso Välimäki. 2010. "Learning in Contests." Manuscript, Graduate School of Business, University of Exeter.
- McKinsey & Company. 2009. "And the Winner Is...: Capturing the Promise of Philanthropic Prizes." Sydney, Australia: McKinsey & Company, 2009.
- Moldovanu, Benny and Aner Sela. 2001. "The Optimal Allocation of Prizes in Contests." *American Economic Review* :542–558.

- Moroni, Sofia. 2015. "Experimentation in Organizations." Manuscript, Department of Economics, University of Pittsburgh.
- Moscarini, Giuseppe and Lones Smith. 2011. "Optimal Dynamic Contests." Manuscript, Department of Economics, Yale University.
- Moscarini, Giuseppe and Francesco Squintani. 2010. "Competitive Experimentation with Private Information: The Survivor's Curse." *Journal of Economic Theory* 145 (2):639–660.
- Murto, Pauli and Juuso Välimäki. 2011. "Learning and Information Aggregation in an Exit Game." *The Review of Economic Studies* 78 (4):1426–1461.
- Rieck, Thomas. 2010. "Information Disclosure in Innovation Contests." Manuscript, Department of Economics, University of Bonn.
- Scotchmer, Suzanne and Jerry Green. 1990. "Novelty and Disclosure in Patent Law." *RAND Journal of Economics* 21 (1):131–146.
- Stiglitz, Joseph. 2006. "Give Prizes not Patents." *New Scientist* 2569:20.
- Taylor, Curtis. 1995. "Digging for Golden Carrots: An Analysis of Research Tournaments." *American Economic Review* 85 (4):872–890.
- Wirtz, Julia. 2013. "Feedback and Learning in Tournaments." Manuscript, Department of Economics, University of Warwick.
- Yildirim, Huseyin. 2005. "Contests with Multiple Rounds." *Games and Economic Behavior* 51:213–227.