INSEAD
The Business School for the World®

Working Paper Series

2015/64/TOM

Sourcing Innovation: Public and Private Feedback in Contests

Jürgen Mihm
INSEAD, jurgen.mihm@insead.edu

Jochen Schlapp
University of Mannheim, jschlapp@bwl.uni-mannheim.de

Contests, in which contestants compete for a prize offered by a contest holder, have become a popular way to source innovation. Despite great interest from the academic community, many important managerial aspects of contests have received very little formal inquiry. The most important of these is feedback from the contest holder to the contestants while the contest unfolds. This paper sets out to establish a comprehensive understanding of how to give feedback in a contest by answering the questions of when to give feedback and when not to give feedback and which type of feedback to give, public (which all solvers can observe) or private (which only the concerned party can observe). We find that feedback will not affect the behavior of competing problem solvers unless the contest holder credibly pre-commits to a truthful feedback policy. We then set up a framework that reduces the feedback decision to a pair of conceptual questions. First: Is the contest's ultimate objective to increase average quality or to find the best solution? Second: How uncertain are outcomes for the solvers? We show that no feedback or public feedback generally dominate private feedback. However, if the host is interested exclusively in the best performance and if the contest displays large uncertainties, private feedback is optimal.

Keywords: Open Innovation; Innovation Contest; Innovation Incentives; Information Transmission; Feedback; Research and Development

We wish to acknowledge Stylianos Kavadias and Nektarios Oraiopoulos as well as the attendees of the presentation of the Darden & Cambridge Judge Entrepreneurship and Innovation Research Conference for comments.

Electronic copy available at: http://ssrn.com/abstract=2659171

1. Introduction

Firms increasingly source their innovation from beyond their own boundaries (Chesbrough 2003, von Hippel 2005). Although doing so enables them to exploit a much broader knowledge base, it also creates substantial challenges for managing the innovation process. Most notably, the firm yields its immediate control over the process and the inventors. *Innovation contests* provide the firm with a mechanism for retaining some control. At the outset of the contest, the firm specifies its goals; at the end of the contest, it grants an award—typically to the inventor with the best solution. Thus the competitive setting incentivizes the inventors. These elements are foundational aspects of all contests. But the firm also makes other more discretionary choices about exerting control. For instance, the firm may provide feedback during the innovation process: it may dynamically inform the inventors about how much it values their (preliminary) innovations, thereby gradually eliminating any uncertainty about its preferences and thus affecting the inventors' decisions as the contest unfolds. Yet this possibility confronts the firm with challenging decisions: When sourcing innovation through a contest, should the firm give intermediate performance feedback or not? If it does, should the feedback be public, so that everyone can observe it, or private, so that only the concerned party is aware? Answering these questions is the focus of our study.

Consider the example of Kaggle. In recent years, many companies have begun to collect vast amounts of data. However, few have built their own machine learning capabilities sufficient to the task of fully exploiting this data. Building such capabilities is difficult because technology evolves rapidly, and so different groups of data scientists worldwide leapfrog each other with respect to knowledge about certain contexts. In this dynamic environment, Kaggle provides a platform that companies can use to access precious outside knowledge and resources. More precisely, it allows companies to provide data (typically a small sample of the full problem) and organize a contest. For example, in 2012 the Heritage Provider Network offered a reward of $500k to the team that would best "[i]dentify patients who will be admitted to a hospital within the next year using historical claims data." When setting up such a contest, companies must define the rules of engagement. In particular, those rules establish the relevant metric—typically, out-of-sample accuracy of the algorithm's predictions—and specify the reward(s) offered. In response to the contest announcement, data scientists compete against each other by developing at their own expense (of time and money) algorithms that perform the required task. The group of scientists that ultimately provides the best-performing algorithm wins the prize. During the competition, data scientists can enter preliminary versions of their code and receive feedback on how well it performs. Yet Kaggle not only provides absolute performance feedback to the team itself; it also maintains a public “leaderboard” so that each participant or team can study its performance relative to all competing submissions.

Many such contests have been devised that involve data applications. For example, the Netflix Prize, organized by Netfilx itself, was established to award $1 million to the group that built the best algorithm for predicting how users would rate movies. But contests and platforms for contests have permeated many other fields as well. In the software context, for example, platforms such as Devpost (which focuses on API development) and IEEEmadC (for mobile applications) have sprung up. The SAP Google Glass Challenge rewarded development teams that created applications integrating SAP platform technologies with Google Glass, and the Microsoft Imagine Cup encouraged applications that could have a major impact on our future life.

All these examples follow a common underlying script. A firm has a problem that requires a tailored innovation, a problem with some inherent uncertainty. The firm advertises a reward for solving this problem. Many different solvers can devise a solution. In providing a solution, each solver exerts effort for which he incurs some private cost; but given the task’s inherent uncertainty, the solver cannot predict exactly how his performance is influenced by effort. Furthermore, the solver cannot perfectly anticipate which features of a solution the firm values and which it does not. In short, the consequences of effort are stochastic. Conversely, the firm immediately recognizes the value of a solution, once presented with it, but has no way of knowing how much effort the solver needed to expend. So the firm cannot hire the solver and directly compensate him for those efforts. Finally, the solution process is inherently dynamic and therefore extended over time.

Their commonalities notwithstanding, the examples just given differ in two key respects: the type of feedback and the firm’s objective. Kaggle and the Netflix Prize give public feedback—that is, the feedback they provide on each entry can be observed by all solvers. In contrast, the software development platforms Devpost and IEEEmadC provide feedback in decidely different ways: IEEEmadC informs developers privately about the quality of their own initial submission but without publishing information on the performance of competitors; even more restrictively, Devpost provides no interim performance feedback at all. And whereas the Netflix Prize’s objective was to identify the *best* solution, the SAP Google Glass Challenge and the Microsoft Imagine Cup sought to incentivize raising the community’s *average* solution quality.

The common script, known as a “contest”, is a well-studied one (Lazear and Rosen 1981, Moldovanu and Sela 2001, Terwiesch and Xu 2008, Siegel 2009). However, the aspect of feedback has received only sparse attention (for notable exceptions, see Aoyagi 2010, Ederer 2010, Goltsman and Mukherjee 2011). In particular, we do not know what type of feedback, public or private, is optimal for which type of objective, best or average performance. To fill this gap, we aim in this paper to establish when and how feedback can be used to increase the value of innovation contests.

Our study makes several contributions to the existing literature. First, it puts on a solid theoretical foundation the answer to a fundamental question: Can the firm make feedback decisions *dynamically* or must it rather commit, at the outset, to a particular feedback mapping (a translation of first-round performance to feedback signal)? We find that the contest holder must commit to a feedback mapping in advance of the contest. Second, the paper extends existing insights from the literature on public feedback and average solution quality to innovation contests in which the firm is interested only in the best outcomes. Third, our paper's main contribution is offering a solution to the technically challenging problem of private feedback, and comparing the case of private feedback to the cases of public feedback and no feedback. We uncover a non trivial relationship between the contest characteristics and the benefits of private, public, and no feedback. Thus our paper makes a first step toward a comprehensive understanding of feedback in contests—and hence toward a comprehensive theory of sourcing innovation.

2. Related Literature

The question of how to motivate innovation and creativity has become a central topic of academic inquiry (see e.g. Manso 2011, Erat and Krishnan 2012, Ederer and Manso 2013, Erat and Gneezy 2015). In particular, contests as a mechanism for eliciting opportunities have become a focal point of attention. The literature on contests is broad, spanning both the economics and the operations management literatures.[1] Starting with the pioneering work of Lazear and Rosen (1981), Green and Stokey (1983), and Nalebuff and Stiglitz (1983), the contest has become the primitive for studying a wide variety of settings; these include lobbying, promotional competition, litigation, military conflict, sports, education, internal labor markets, and of course R&D management (for an overview of applications, see Konrad 2009). It will be useful for our purposes to classify works in the contest literature along two dimensions.[2] The first dimension classifies this research in terms

[1] The operations management literature uses the term "tournaments". The economics literature prefers the term "contest", of which "all-pay auctions" and "tournaments" are special cases. We shall use "contest" throughout because it is the more established term.

[2] In this section we focus on theoretical work, but there is also a broad empirical and experimental literature (for a comprehensive survey of the experimental work, see Dechenaux et al. 2014). Many researchers have investigated the implications of architectural decisions that are made at the outset of a contest. These decisions include, among others, the award structure (Bull et al. 1987, Liu et al. 2014), the number of solvers (Garcia and Tor 2009, Boudreau et al. 2011, Lim et al. 2014), talent differences (Schotter and Weigelt 1992, Fonseca 2009, Jeppesen and Lakhani 2010, Bockstedt et al. 2015), the distribution of information (Brookins and Ryvkin 2014), and problem uncertainty (Boudreau et al. 2011). More recently, researchers have begun to pay closer attention to decisions that are made in the course of a contest. Most notably, the center of attention has shifted toward the consequences of providing

of its major technical assumptions and thus of the situations to which the contest applies; the second dimension classifies the literature in terms of the particular contest design question being addressed.

With regard to the first classifying dimension, the most important distinction between different innovation-related contest models is based on the link between a solver's effort and his contest performance. This link between a solver's action and the outcome is one that ranges from purely deterministic to purely stochastic, with many hybrid forms that incorporate both deterministic and stochastic components. To a lesser extent we further differentiate between models that focus on the best performance versus average performance.

The economics literature has largely (although not exclusively) focused on deterministic links. Two substreams have emerged. The literature on *complete* information contests assumes that all characteristics of the players are observable. Early work has seeked to characterize equilibria in symmetric contests (e.g., Glazer and Hassin 1988), but that aim has been expanded to settings that feature asymmetric cost functions or valuations (e.g., Che and Gale 2003, Siegel 2009, 2010, 2014b). The literature on *incomplete* information contests assumes that each solver holds some piece of private information which is often framed as some type of individual capability. Research in this substream has yielded results both for symmetric settings (i.e., for ex ante identical solvers; see Krishna and Morgan 1997, Moldovanu and Sela 2001) and for asymmetric settings (i.e., for ex ante different solvers; see Amann and Leininger 1996, Parreiras and Rubinchik 2010, Siegel 2014a). The central concern in much of this literature is exploring how solvers can be induced to provide good solutions on average (Krishna and Morgan 1997, Moldovanu and Sela 2001). That being said, there is some research (e.g., Moldovanu and Sela 2006) that focuses also on the best entry's performance.

A second stream of literature assumes that there is a stochastic link between solvers' actions and contest outcomes. This stream was originated by Taylor (1995) in the context of R&D contests and later generalized by Fullerton and McAfee (1999). An extreme is presented by Gaba et al. (2004) and

in-contest feedback. Several studies highlight the motivation effect of feedback. For instance, Berger and Pope (2011) show—in their empirical investigation of data from basketball games—that the feedback of being slightly behind motivates a team to fight harder. In contrast, Casas-Arce and Asis Martinez-Jerez (2009) present empirical evidence from sales contests that public feedback results in lower effort if the performance difference becomes substantial, a finding that is confirmed by the laboratory experiments of Ludwig and Lünser (2012). Interestingly, experimental evidence by Kuhnen and Tymula (2012) indicates that even the mere announcement of future feedback has an effect on the behavior of contest problem solvers. Ederer and Fehr (2009) and Gürtler and Harbring (2010) investigate the implications of manipulatable feedback; each study finds that solvers react to feedback even when it is not truthful (although that reaction is much weaker than when the feedback is truthful). Wooten and Ulrich (2014) demonstrate that although directed feedback raises the average quality of all submissions, it may reduce the best submission's quality. We find it intriguing that none of the works just cited has studied the implications of private feedback.

Tsetlin et al. (2004) who investigate contests that lack any deterministic relation between actions and outcomes. Closer to our work are those studies that consider a combined deterministic and stochastic relationship between actions and outcomes. This approach has gained traction among those who study contests in the context of innovation (e.g., Terwiesch and Xu 2008, Terwiesch and Ulrich 2009, Ales et al. 2014, Körpeoğlu and Cho 2015). These researchers have also shifted the focus from fostering average performance to maximizing the best entry's performance.

The technical classification in terms of how effort and outcomes are related is important on a conceptual level. First of all, assuming a stochastic link between effort and outcome seems appropriate when modeling innovation and R&D contexts. The search for new products, methods, and technologies is inherently uncertain, so disregarding that aspect will call model predictions into serious question. A second point reinforces this notion. As Siegel (2009, p. 72) emphasizes, analyzing a stochastic link between effort and outcomes requires fundamentally different assumptions than in the deterministic case—especially with regard to interaction effects between effort and uncertainty. These assumptions make the models applicable to different types of settings yet entail very different predictions about a solver's behavior.

The existing contest literature can also be classified according to a second criterion—namely, the (contest) design issue being addressed. Prominent among such issues is whether or not access to the contest ought to be limited: Should (or shouldn't) the contest be open to everybody? The literature that assumes a deterministic effort-to-outcome link tends to favor a degree of restrictiveness (e.g., Taylor 1995, Moldovanu and Sela 2001), and some researchers have even postulated that—under a wide array of assumptions—the optimal number of solvers is two (Fullerton and McAfee 1999, Che and Gale 2003). However, the literature focusing on innovation contests takes a more nuanced view. In particular, these studies emphasize that the contest designer must balance two countervailing forces: a larger number of entrants yields a larger number of trials, but at the cost of each solver expending less effort on his respective trial (Terwiesch and Xu 2008).[3] Bid caps have been studied as a means for limiting access to a contest (Gavious et al. 2002, Che and Gale 2003) and so have more advanced mechanisms such as auctions for the right to participate (Fullerton and McAfee 1999). Another prominent issue is the optimal award structure. Conditions for the optimality of one prize versus many prizes depend on the solvers' cost function (Moldovanu and Sela 2001), on performance uncertainty (Ales et al. 2014), and on whether the firm is seeking the best solution or merely to improve the average solution (Moldovanu and Sela 2006). Another perspective is

[3] As an exception to this general concept Ales et al. (2014) and Körpeoğlu and Cho (2015) give examples of contests for which individual solution efforts are *increasing* in the number of competitors.

Figure 1 Classification of the Contest Literature

		Without feedback	**Public feedback**	**Private feedback**
Stochastic effort consequences	Best solution	Taylor (1995), Fullerton and McAfee (1999), Terwiesch and Xu (2008), Erat and Krishnan (2012), Ales et. al. (2014)		
	Average solution	Lazear and Rosen (1981), Green and Stokey (1983)	Aoyagi (2010), Ederer (2010), Goltsman and Mukherjee (2011)	
Deterministic effort consequences		Amann and Leininger (1996), Krishna and Morgan (1997), Moldovanu and Sela (2001, 2006), Che and Gale (2003), Siegel (2009, 2010, 2014a, 2014b), Parreiras and Rubinchik (2010), Körpeoğlu and Cho (2015)	Yildirim (2005), Gershkov and Perry (2009)	

Focus of this paper

offered by Che and Gale (2003) who discuss *menus* of prizes. The last major issue is the contest's temporal structure. Should the contest designer offer one big collective contest or instead a series of "cascading" contests?—see Moldovanu and Sela (2006) and Konrad and Kovenock (2009).

All of these questions presume that the contest holder is relatively passive during the course of the contest. More recently, attention has been shifting to the actions that a contest holder could take as the contest unfolds—for instance, to preclude collusion among different solvers (Gürtler et al. 2013). However, the most important in-contest decision is whether or not to provide interim performance feedback.

The literature on feedback in contests is sparse. The paper generally acknowledged to be the first in this area (Yildirim 2005) does not address feedback per se, but instead focuses on information disclosure as a strategic choice made by solvers. Gershkov and Perry (2009) are likewise not primarily concerned with feedback as we understand it here; instead, they are concerned with optimally aggregating scores by combining intermediate and final reviews when the review process itself is noisy. However, there are three papers that do address feedback in contests in a more narrow sense. Goltsman and Mukherjee (2011) explore a setting in which solvers compete for a single price by fulfilling two tasks at which they can either fail or succeed. Closer to our work, both Aoyagi (2010) and Ederer (2010) examine settings in which a firm provides feedback to solvers who have to make continuous effort choices.

It is noteworthy that past work on feedback in contests shares three aspects that bear strongly on our work. First, all authors casually assume that—whatever its feedback policy—the firm should pre-commit to a truthful feedback mapping; that is to a truthful translation of first-round performance to feedback signal. In other words, rather than dynamically choosing the feedback signal that optimizes outcomes dynamically once the contest is under way the firm should credibly pre-commit at the start of the contest to giving truthful feedback. Second, the papers cited here all take the firm's main interest to be *average* effort provision and not the value of the *best* performance.Third, and most importantly, none of these works consider private feedback. In short, past work on feedback in contests restricts attention to the question of how public feedback influences a solver's average solution quality.

We extend previous work along those three dimensions. First, we establish the necessity of pre-commitment in feedback by first solving our game without such pre-commitment. Second, we investigate how feedback affects the outcome of innovation contests—that is, contests in which the firm's interest is in strictly the best solution. Our third and main contribution is introducing private feedback into the analysis, which allows us to make full comparisons among different feedback policies. In sum, our analysis establishes which kind of feedback (no, public, or private feedback) should be used for which kind of contest and which kind of objective (average vs. best solution).

Figure 1 summarizes the foregoing discussion, and it makes clear that further analysis is required if we are to develop a more comprehensive understanding of how feedback should be used in innovation contests.

3. Model Setup

Let us recapitulate our model's requirements in terms of both the firm and the solvers. The firm understands its own preference structure sufficiently that, when presented with a solution, it can express how much it values the solution. However, the firm cannot know the effort level expended by solvers in achieving a given performance because the link between performance and effort has a stochastic component. In contrast, each solver knows how much effort he expends and also understands that expected performance increases in effort. Yet, solvers still experience ex ante uncertainty about how, exactly, effort links to performance; ex post, solvers are uncertain about the firm's preference structure and so, even after devising a solution, they cannot truly evaluate its performance. This ex post uncertainty follows because for any true innovation effort the firm cannot *fully* specify, in advance, what criteria it values or how they should be weighted.

Consider the Microsoft Imagine Cup, which seeks software applications that "change the future of software solutions". The stated goal is clearly so broad that Microsoft cannot pre-specify exactly what it wants; once presented with a solution, however, the firm will know how to evaluate the result. In contrast, the solvers are naturally unable to anticipate precisely how the firm will react to their software applications. Although the implemented functionality will, of course, partly depend

Figure 2 **Structure of the Innovation Contest with Feedback**

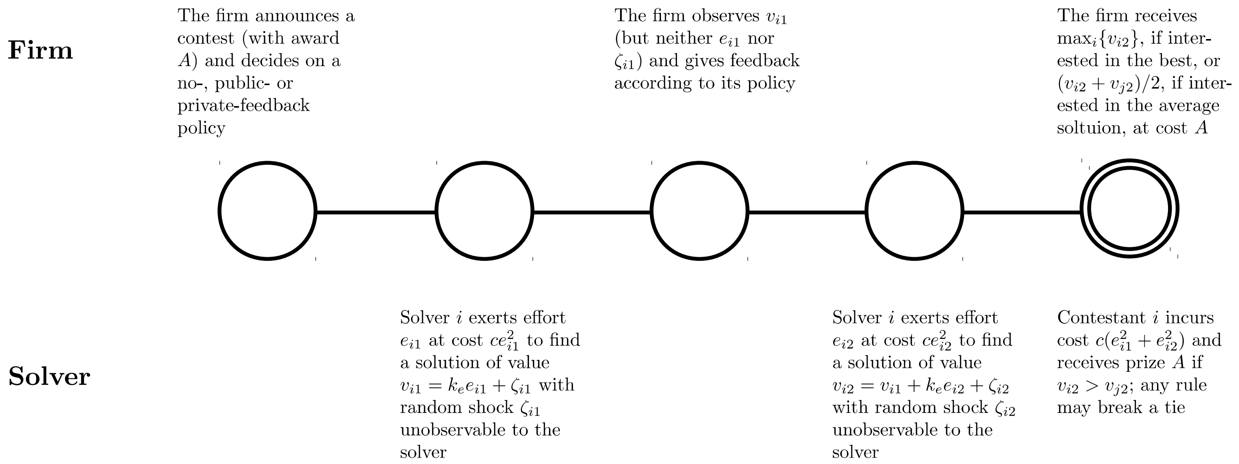

on the effort of solvers, the "inspiration" factor is always in play—and thus the contest outcome depends on an element of randomness.

Stochasticity in effort consequences implies that neither a solver's performance nor his solution efforts are verifiable. As a consequence, the firm can neither write a contract on performance (because its preference structure is non-verifiable) nor contract for a solver's effort (because the solver's effort is non-verifiable). The firm must therefore resort to alternative mechanisms in order to create proper incentives for the solvers. Foremost among such mechanisms are "contests", in which solvers compete against each other by submitting rival solutions. As a means of dynamically influencing the solvers' effort provision in the course of a contest, the firm may (partially) resolve the solvers' uncertainty about their performance by transmitting *interim performance feedback*. Such feedback can come in three different forms: no feedback, public feedback (each solver can observe the feedback given to all solvers for all solutions), or private feedback (each solver can only observe feedback for his own solution).

As do the different examples described in the Introduction, our model must reflect that different firms pursue different goals. A firm may be exclusively interested in the best solution and not care about any of the other solutions. Alternatively, the firm may be interested in using the contest to raise the average solution quality of all competing solvers.

Formal Model Description. In order to create a parsimonious model that nonetheless captures the essence of the scenario just outlined, we consider a firm that is hosting a dynamic innovation contest over two rounds, $t \in \{1,2\}$, with two risk-neutral solvers, i and j.[4] The primitives of the contest are common knowledge; its basic structure is depicted in Figure 2.

In a first step, the firm publicly announces the contest, the fixed award A for which the two solvers compete, and its feedback policy. In order to concentrate on the role of feedback (and to limit technical complexity), we treat $A > 0$ as a parameter. Our decision variable for the firm at this stage is whether and how to give feedback. The firm may choose to give no feedback at all; to offer public feedback (i.e., both solvers receive the same information about their own and their competitor's performance); or to provide private feedback (i.e., solver i receives feedback on his own performance but not on the performance of solver j, and vice versa).

Next, solver i expends effort $e_{i1} \geq 0$ at private cost ce_{i1}^2, with $c > 0$. He finds an initial solution of value $v_{i1} = k_e e_{i1} + \zeta_{i1}$, where $k_e > 0$ is the (commonly known) sensitivity of effort and ζ_{i1} is a random shock that follows a (commonly known) uniform distribution, $\zeta_{i1} \sim \text{Uniform}(-a/2, a/2)$ with $a > 0$.

After the first round, the firm perfectly observes v_{i1}. However, solver i's effort is unobservable to the firm (and the other solver); hence the firm cannot determine whether a high solution value stems from a high effort, a high random shock, or both. In contrast, solver i knows how much effort he has invested; however, since he cannot observe the realization of ζ_{i1}, he is uncertain about the true performance of his solution. To provide solver i with additional information regarding his performance, the firm provides interim performance feedback according to its feedback policy—which for now need not be truthful. Formally, let Σ be an arbitrary measurable set that contains all possible signals that the firm can transmit as feedback, and denote by $s_i \in \Sigma$ ($s_j \in \Sigma$) the feedback about solver i's (j's) absolute performance. Depending on the firm's feedback policy, solver i observes (i) neither s_i nor s_j (no feedback); (ii) s_i and s_j (public feedback), or (iii) only s_i (private feedback).

Upon observing the firm's feedback, solver i updates his belief about the realization of the first-round performances $v_1 = (v_{i1}, v_{j1})$ according to Bayesian rationality. Then, solver i expends additional solution effort $e_{i2} \geq 0$ and submits his final solution $v_{i2} = v_{i1} + k_e e_{i2} + \zeta_{i2}$, where ζ_{i2} is again a random shock that follows the same distributional assumptions as in the first round. Random shocks are independent and identically distributed across solvers and rounds. For notational simplicity we define $\Delta\zeta_t = \zeta_{it} - \zeta_{jt}$ as the difference of the random shocks in round t with associated probability density function $g_{\Delta\zeta_t}$.

[4] For notational simplicity, we explicitly define only the parameters for solver i; an identical set of parameters applies to solver j.

Finally, after receiving the final solutions $v_2 = (v_{i2}, v_{j2})$, the firm announces the winner of the contest by choosing the solution with the highest value. Thus solver i wins if $v_{i2} > v_{j2}$ (ties can be broken by invoking any rule).

Model Implications. The firm will naturally seek to employ the feedback strategy that maximizes its expected profits. The two different profit functions are $\Pi_{\text{best}} = \mathbb{E}[\max\{v_{i2}, v_{j2}\}] - A$ if the firm is interested in the performance of the *best* solution only, and $\Pi_{\text{avg}} = \mathbb{E}[v_{i2} + v_{j2}]/2 - A$ if the firm wishes to maximize the *average* performance of both solvers.

Whereas the firm—whatever its profit function—is interested in the solvers' absolute performance, each solver is interested only in his relative performance. More precisely, a solver only cares about whether or not he wins the contest. The utility that solver i receives from winning the contest is $A - \sum_t ce_{it}^2$; losing the contest yields a utility of $-\sum_t ce_{it}^2$. Hence solver i's expected utility of participating in the contest is $u_i = A \cdot \mathbb{P}(v_{i2} > v_{j2}) - \sum_t ce_{it}^2$, and he chooses how much effort to invest in the contest by maximizing his expected utility.

We are interested in Perfect Bayesian Equilibria (PBE) of the contest. To avoid unnecessary technical complications during the analysis, we assume that $\kappa \equiv (a^2 c)/(A k_e^2) > 1$. For technical reasons, similar assumptions on the contest's inherent performance uncertainty have become customary in virtually the entire literature on contests (see, e.g., Nalebuff and Stiglitz 1983, Aoyagi 2010, Ederer 2010). Clearly, κ increases in the variance of the random noise and the costs of effort, and it decreases in the size of the award and the effort sensitivity. Thus, with a relatively higher κ, improvement effort is relatively more expensive and the solution performance becomes relatively more stochastic.

4. Feedback: Commitment and Truthfulness

The goal of any interim performance feedback must be to influence dynamically the solvers' effort choices, for otherwise feedback would be futile. In this section, we establish the key properties that any feedback policy must possess in order to actually fulfill this goal.

As shown in Figure 2, the firm announces at the outset of the contest whether it provides feedback as well as whether the feedback content is public or remains private. Yet the firm does not necessarily announce how a solver's performance will be mapped into a specific feedback instance; that is, the firm is free to choose its *feedback mapping* after observing the solvers' first-round performance. Formally, a feedback mapping is a tuple (r, Σ) that maps solver i's first-round performance v_{i1} into a set of possible feedback signals Σ according to the relation $r \colon \mathbb{R} \to \Sigma$. Put differently, given the realization of v_{i1}, the firm determines a feedback signal s_i that is drawn from the probability density function $r(s_i|v_{i1})$. Because the firm can choose the feedback mapping *after* observing the solvers' performance, a strong incentive misalignment arises between the firm and the solvers. On the one hand, each solver appreciates true performance feedback. On the other hand, the firm wishes to use the feedback as a way of maximizing the solvers' second-round effort irrespective of first-round outcomes. This tension has a detrimental effect on the value of feedback in the unrestricted game.

Theorem 1 (Feedback Commitment and Truthfulness). *The firm can influence the solvers' second-round effort provision with interim performance feedback if and only if* (i) *the firm pre-commits to a feedback mapping* (r, Σ) *before the contest starts and* (ii) *the feedback mapping is indicative of a solver's true performance—that is,* $r(s_i|v_{i1}) \neq r(s_i|v'_{i1})$ *for some* s_i*,* v_{i1}*, and* v'_{i1}*.*

The essence of Theorem 1 is that, in order to provide impactful feedback, the firm must credibly commit to a feedback mapping at the outset of the contest and, moreover, this feedback mapping must reflect the solvers' actual performance. The intuition for this result is instructive. Suppose that the firm does *not* credibly commit to a certain feedback mapping. In this case, the firm has a strong incentive to manipulate feedback with the goal of boosting the solution efforts of the solvers. In fact, because the feedback signal does not affect the firm's profits directly but only indirectly (i.e., through the solvers' efforts), the firm always provides the feedback that maximizes second-round efforts *irrespective* of first-round outcomes. In equilibrium, each solver anticipates this manipulation and discards the received information as meaningless; thus each solver acts as if he never received any feedback. Mathematically, the unique equilibrium of the game depicted in Figure 2 is a so-called babbling equilibrium in the spirit of the well-known "cheap talk" results of Crawford and Sobel (1982). Therefore, the firm can avoid meaningless information transmission only if it commits to its feedback mapping at the beginning of the contest.

This simple game-theoretic property has deep repercussions for practice. First of all, the firm needs to devise methods for committing itself. One such method can be rooted in technology. Kaggle and the Netflix Prize, for example, distribute a part of their data set; this allows solvers to try out their algorithms on the data and receive automated performance information. Another method is to rely on the firm's reputation. It is clear that if the firm does not have a reputation for sticking to its feedback mapping, then the solvers will discount the feedback signals. This implies that firms should build a reputation for providing the feedback mapping that they announced. Deviating from this reputation will destroy the firm's credibility and render feedback moot.

The second major implication of Theorem 1 is that the feedback signal must be not only pre-committed, but also indicative of the truth. If feedback is not correlated with acutal performance then solvers will discount it, and again the situation would decay into a babbling equilibrium. Consider the Kaggle example. Because the firm wants the solvers to adapt the algorithm to a real situation, it cannot distort the data set without inflicting harm on itself. In short, the feedback mechanism must provide *commitment* and entail *truthfulness*.

5. Solution Efforts of the Solvers

Given the results of the previous section, we now turn our attention to the question of how the firm can influence the solvers' solution efforts by providing pre-committed and truthful feedback. Specifically, from here on we assume that if the firm provides feedback, it does so by truthfully revealing a solver's absolute performance; that is, $r\colon v_{i1} \mapsto v_{i1}$.

In this section, we focus on the solvers' solution efforts under each feedback policy. We can do so without specifying the firm's objectives because, given a particular feedback policy, the solvers' strategies are independent of the firm's goals. That is, each solver tries to win the contest, regardless of whether the firm aims to improve average performance or rather to attain the best performance.

As a benchmark, we characterize the solvers' equilibrium efforts in the absence of feedback (Section 5.1). We then examine how the provision of public feedback affects the solvers' solution efforts (Section 5.2). Finally, we determine equilibrium solution efforts under a private-feedback policy (Section 5.3). Although we discuss a few initial managerial implications, we leave the answer to our research question—of when and how feedback should be provided—to Section 6, which compares the different feedback policies in terms of the firm's profits.

5.1. No Feedback

In the benchmark case of no feedback, the firm does not provide any interim performance information to the solvers. Thus, the solvers' two-stage effort choice problem reduces to a simultaneous single-stage utility maximization problem. Mathematically, solver i chooses his solution efforts to maximize his expected utility $u_i = A \cdot \mathbb{P}(v_{i2} > v_{j2}) - \sum_t c e_{it}^2$.

Proposition 1 (No-Feedback Policy). *The unique PBE under a no-feedback policy is symmetric with*

$$e_1^{\text{no}} = e_2^{\text{no}} = \frac{Ak_e}{3ac}. \tag{1}$$

Proposition 1 parallels previous results of Taylor (1995) and Fullerton and McAfee (1999). Given that neither solver receives any interim performance information after the first round and that the costs of effort are convex, solution efforts are identical across rounds. Moreover, since solvers are symmetric at the outset of the contest, it follows that they always choose the same effort in equilibrium; in other words, no solver leapfrogs his competitor by investing more effort. So, under a no-feedback policy, it is only the inherent performance uncertainty that ultimately determines the contest winner.

At this point, it is instructive to examine how our key contextual parameters affect a solver's solution efforts. As we would expect, solution efforts are increasing in the size of the award, A, and in the effort sensitivity, k_e; they are decreasing in the costs of effort, c, and in the involved uncertainty, a. Thus, a solver exerts relatively more effort if effort becomes relatively more rewarding (i.e., A/c increases) and if effort becomes relatively more important (i.e., k_e/a increases).

5.2. Public Feedback

Next we study the implications of truthful public feedback. In this case, after submitting his initial solution, each solver learns his own first-round performance as well as his competitor's performance. Thus public feedback perfectly reveals the solvers' first-round performance difference before the start of the second round. As a result, solvers are no longer symmetric in the second round of the contest.

PROPOSITION 2 (PUBLIC-FEEDBACK POLICY). *The unique PBE under a public-feedback policy is symmetric with*

$$e_1^{\text{pub}} = \mathbb{E}_{\Delta\zeta_1}\left[e_2^{\text{pub}}(\Delta\zeta_1)\right] = \frac{Ak_e}{3ac}, \tag{2}$$

$$e_2^{\text{pub}}(\Delta\zeta_1) = \frac{Ak_e}{2a^2c}\cdot(a - |\Delta\zeta_1|)\,. \tag{3}$$

Public feedback informs solver i about his *absolute* and *relative* first-round performance. In line with previous work (Aoyagi 2010, Ederer 2010), Proposition 2 states that each solver cares only about his relative performance and completely disregards the absolute performance information. This focus on relative performance has far-reaching implications for the solvers' solution efforts. The closer are the solvers' first-round performances, the greater are the second-round efforts. So the more competitive the contest is in the second round, the harder the solvers fight to win the contest. In contrast, if the first-round performance difference is substantial, then solvers reduce their solution efforts because the contest is de facto decided. In the extreme case, if $|\Delta\zeta_1| = a$,[5] the follower can no longer win with positive probability; hence no solver expends any solution effort in the second round.

Another interesting result of Proposition 2 is that, surprisingly, despite being asymmetric in the second round, both solvers expend the same amount of effort. In other words, the first-round leader

[5] An event of probability 0.

pursues a simple blocking strategy; that is, he tries to keep the follower at distance but without trying to increase the performance gap. At the same time, the follower tries to not fall farther behind but without attempting to close the gap. However, it is noteworthy that these dynamics do not ensure that the first-round leader will win the contest. The follower might be fortunate in the second round (i.e., he might receive a large positive second-round shock), which would enable him to overtake the interim leader.

Finally, we note that the ex ante expected effort in both rounds is identical. This result follows from the fact that a solver's performance increases linearly in effort while the costs of effort are convex increasing. In expectation, then, it is optimal for a solver to split his solution efforts equally between the two rounds.

5.3. Private Feedback

We have just shown that, under a public-feedback scenario, solvers set their second-round solution efforts as a function of their *relative* first-round performance. Such a policy is no longer possible under a private-feedback scenario because each solver receives information only about his own performance. Thus, it is only the *absolute* performance information that can affect a solver's solution effort.

The absence of relative performance information fundamentally affects the contest's information structure. Whereas solvers always possess symmetric and consistent beliefs under no and public feedback, private feedback introduces an asymmetric and inconsistent belief structure. In other words, the solvers' assessments of their chances to win need not be "coherent" under private feedback. Suppose, for example, that each solver receives the information that he performed extremely well in the first round. Then both solvers believe that their respective chances of winning are much greater than 1/2, although in reality they are merely 1/2. Moreover, solvers are never entirely certain whether they are ahead or behind their competitor—in contrast with the public-feedback scenario. It is this asymmetric belief structure that results in serious mathematical complications and that enforces asymmetric equilibrium solution efforts. Proposition 3, which summarizes the solvers' equilibrium solution efforts under private feedback, clearly mirrors this reality, as the second-round equilibrium can only be characterized in closed form using inverse functions.

Proposition 3 (Private-Feedback Policy). *Let* $v^{\text{pri}}(\zeta_{i1}) = \zeta_{i1} + k_e e_2^{\text{pri}}(\zeta_{i1})$. *Then the unique PBE under a private-feedback policy is symmetric and given by*

$$e_1^{\text{pri}} = \mathbb{E}_{\zeta_{i1}}\left[e_2^{\text{pri}}(\zeta_{i1})\right], \tag{4}$$

$$\zeta_{i1} = v^{-1}(v^{\text{pri}}) = \begin{cases} \gamma_1 e^{\frac{3}{2a\kappa}v^{\text{pri}}} + \gamma_2 e^{\frac{1}{2a\kappa}v^{\text{pri}}} - a\left(\frac{1}{6} + 2\kappa\right) & \text{if } v_o - a \le v^{\text{pri}} < v_u \\ \gamma_3 e^{\frac{1}{a\kappa}v^{\text{pri}}} - a\kappa & \text{if } v_u \le v^{\text{pri}} \le v_o \\ \gamma_4 e^{\frac{3}{2a\kappa}v^{\text{pri}}} + \gamma_5 e^{\frac{1}{2a\kappa}v^{\text{pri}}} + a\left(\frac{1}{6} - 2\kappa\right) & \text{if } v_o < v^{\text{pri}} \le v_u + a, \end{cases} \tag{5}$$

with the constants defined by $v_u = 2a\kappa\ln(x)$, $v_o = 2a\kappa\ln(y)$, $\gamma_1 = pn^2(x - ny)/(3x^2o)$, $\gamma_2 = pxy(n^3x + y)/(x^2o)$, $\gamma_3 = p(n^2x^2 + y^2)/(2x^2o)$, $\gamma_4 = -p(x - ny)/(3nx^2o)$, *and* $\gamma_5 = pxy(n^3x + y)/(nx^2o)$, *where* $m = (1 - 6\kappa)/(1 + 6\kappa)$, $n = e^{1/(2\kappa)}$, $o = 3y^2 - n^2x^2 + 4n^3xy$, $p = a(1 + 6\kappa)$, *and* $x \in \left[e^{-1/(4\kappa)}, e^{-1/(4\kappa)\cdot(1-1/\kappa)}\right]$ *and* $y \in \left[e^{1/(4\kappa)}, e^{1/(4\kappa)\cdot(1+1/\kappa)}\right]$ *are the unique solutions to the following system of equations*

$$mn^2x^4 - 4mn^3x^3y - 3(m + n^2)x^2y^2 - 4n^{-1}xy^3 + y^4 = 0 \tag{6}$$

$$\frac{1 - 6\kappa^2}{\kappa(1 + 6\kappa)} + m\ln(y) - \ln(x) + \frac{n^2x^4 + 8n^3x^3y + 9(1 + n^2)x^2y^2 + 8n^{-1}xy^3 + y^4}{6x^2(3y^2 - n^2x^2 + 4n^3xy)} = 0. \tag{7}$$

Before discussing the managerial implications of Proposition 3, it is worth commenting on the game-theoretic characteristics of the private-feedback policy. The key distinguishing feature of private feedback (compared to no- and public-feedback) is that it introduces an "undetermined asymmetry" owing to the asymmetric belief structure. It is well established that games with asymmetric belief structures are notoriously difficult to solve, so reviewing our solution methodology here may prove useful for future studies facing similar issues.[6] First, from the first-order conditions, the equilibrium effort function $e_2^{\text{pri}}(\zeta_{i1})$ has to solve the following integral equation:

$$Ak_e \cdot \mathbb{E}_{\zeta_{j1}}[g_{\Delta\zeta_2}(\zeta_{i1} + k_e e_2^{\text{pri}}(\zeta_{i1}) - \zeta_{j1} - k_e e_2^{\text{pri}}(\zeta_{j1}))|\zeta_{i1}] = 2ce_2^{\text{pri}}(\zeta_{i1}). \tag{8}$$

We can readily verify that $e_2^{\text{pri}}(\zeta_{i1})$ cannot be characterized in closed form (this argument is confirmed ex post by Equation (5)). To circumvent this issue we introduce the auxiliary function $v^{\text{pri}}(\zeta_{i1}) = \zeta_{i1} + k_e e_2^{\text{pri}}(\zeta_{i1})$, which summarizes the total effect of ζ_{i1} on solver i's performance. The key step in proving Proposition 3 is to substitute $v^{\text{pri}}(\zeta_{i1})$ into (8) and then to transform the integral equation into a system of three Bernoulli equations in the variable $v^{-1}(v^{\text{pri}})$.[7] The solution to this system of Bernoulli equations remains implicit, but we can show that the desired solution is identical to the solution of a system of "equations of damped vibrations". This second key step

[6] For full mathematical details, we relegate the interested reader to the Appendix.

[7] For general polynomial cost functions ce_{i2}^n with $n > 1$, Equation (8) can be transformed into a system of Abel equations of the second kind. In mathematical theory, it is well established that general Abel equations do not have known solutions. In the case of $n = 2$, however, these Abel equations can be further reduced to a system of Bernoulli equations.

Figure 3 Equilibrium Second-Round Effort under Private Feedback

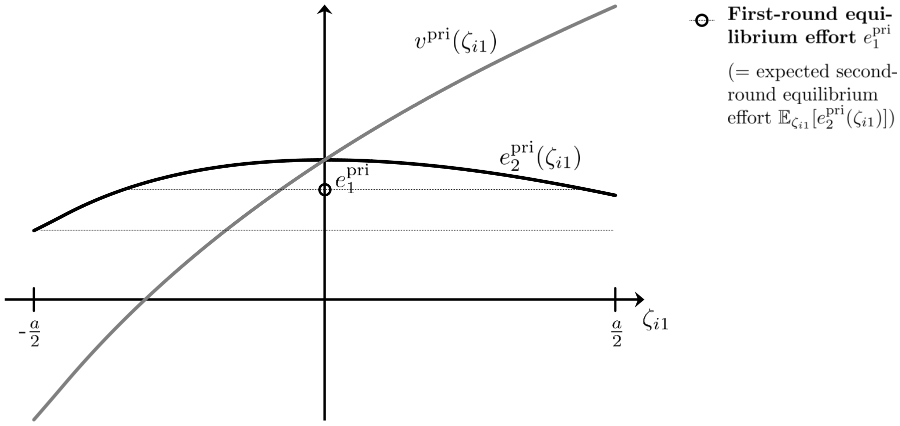

Note. The functions are based on the following set of parameters: $A=1$, $a=1$, $k_e=1$, $c=1.01$.

enables us to rely on standard mathematical tools and thus finally present $v^{-1}(v^{\text{pri}})$ in closed form (see e.g. Polyanin and Zaitsev 2003).

The formulation in Proposition 3 is somewhat unwieldly, but in Corollary 1 we provide an approximation for a solver's second-round effort function $e_2^{\text{pri}}(\zeta_{i1})$ that is much more tractable than the original. Our numerical analyses suggest that Corollary 1 provides an exceptionally good approximation even for low κ, which makes it a good starting point for reflecting on Proposition 3.

COROLLARY 1. *Define* $\widetilde{\gamma}_3 = a(1+6\kappa)e^{(\kappa-1)/(2\kappa^2)}/(2(1+2e^{1/\kappa}))$, *and let*

$$\widetilde{e}_2(\zeta_{i1}) = -\frac{\zeta_{i1}}{k_e} + \frac{a\kappa}{k_e}\ln(\zeta_{i1}+a\kappa) - \frac{a\kappa}{k_e}\ln(\widetilde{\gamma}_3). \tag{9}$$

Then $\lim_{\kappa\to\infty} e_2^{\text{pri}}(\zeta_{i1}) - \widetilde{e}_2(\zeta_{i1}) = 0$ *for all* ζ_{i1}.

Figure 3 visualizes the equilibrium effort functions e_1^{pri} and $e_2^{\text{pri}}(\zeta_{i1})$ as well as the value function $v^{\text{pri}}(\zeta_{i1}) = \zeta_{i1} + k_e e_2^{\text{pri}}(\zeta_{i1})$ for different first-round shocks. The graph makes salient that Proposition 3 provides striking managerial insights for those staging innovation contests.

First, as before, each solver splits his expected solution effort equally between the two rounds. That is: in expectation, the first and the second round contribute equally to a solver's overall performance. Second—and intriguingly—effort in the second round, $e_2^{\text{pri}}(\zeta_{i1})$, does not increase monotonically in ζ_{i1}. In fact, $e_2^{\text{pri}}(\zeta_{i1})$ has an inverted U-shape; it increases in ζ_{i1} for $\zeta_{i1} \leq 0$ but decreases in ζ_{i1} for $\zeta_{i1} > 0$. Thus solvers with a moderate first-round performance (i.e., $\zeta_{i1} = 0$) exert substantial effort in the second round, whereas solvers with a very high or very low first-round performance reduce their second-round efforts. The reason is that a moderately performing solver perceives the contest as being highly competitive whereas exceptionally good- or ill-performing solvers perceive the contest as de facto decided. Moreover, in contrast with the public-feedback scenario, with private feedback the bad solvers reduce their efforts to a greater extent than do good solvers; that is, $e_2^{\text{pri}}(-\zeta_{i1}) < e_2^{\text{pri}}(\zeta_{i1})$ for all $\zeta_{i1} > 0$ (observe the asymmetry in Figure 3). This finding is caused by the absence of relative performance information. A solver with a high first-round shock can never be certain that he is ahead of his competitor. As a result, he invests more effort to maintain his chances of winning in case the competitor is equally strong—even though that is an unlikely event. In other words, private feedback induces well-performing solvers to invest relatively more effort; it makes them relatively more risk averse.

But does this mean that less fortunate players can leapfrog better solvers through effort provision in the second round? The answer is No. To see this, observe that $v^{\text{pri}}(\zeta_{i1})$ (i.e., the total contribution of ζ_{i1} to a solver's performance) increases in ζ_{i1}. Clearly, the more fortunate a solver is in the first round (i.e., the higher his shock ζ_{i1}), the better he performs in the contest. More interestingly, this intuitive result also sheds light on the solvers' strategic behavior. In equilibrium, no solver ever allows a less fortunate solver (i.e., a solver with a lower first-round shock) to overtake him in the second round by exerting enough effort. Thus, once a solver has fallen behind his competitor after the first round, he needs a good random shock in the second round in order to win the contest.

6. The Optimal Feedback Policy

Having characterized the solvers' equilibrium solution efforts under the different feedback policies, we are now ready to answer our research question: Which feedback policy is the best for which kind of objective (refer to Figure 1)? To structure the analysis, we first discuss the optimal feedback policy for maximizing average performance (Section 6.1); we then shift our focus to maximizing the performance of the best solution (Section 6.2).

6.1. Maximizing Solvers' Average Performance

Since the firm must set the feedback policy at the outset of the contest and since solvers are ex ante symmetric, it follows that $\Pi_{\text{avg}} = \mathbb{E}[v_{i2} + v_{j2}]/2 - A = \mathbb{E}\left[\sum_{i,t} e_{it}\right]/2 - A = \mathbb{E}\left[\sum_t e_{it}\right] - A$. That is, maximizing average performance is equivalent to maximizing the sum of one solver's (ex ante) expected first- and second-round efforts. Theorem 2A compares the (ex ante) expected first- and second-round effort choices of a solver as well as the firm's expected profits for no, public, and private feedback.

Theorem 2A (The Optimal Feedback Policy for Average Performance). *The following statements hold:*

(i) $e_1^{\text{pri}} < e_1^{\text{pub}} = e_1^{\text{no}}$;

(ii) $\mathbb{E}_{\zeta_{i1}}\left[e_2^{\text{pri}}(\zeta_{i1})\right] < \mathbb{E}_{\Delta\zeta_1}\left[e_2^{\text{pub}}(\Delta\zeta_1)\right] = e_2^{\text{no}}$;

(iii) $\Pi_{\text{avg}}^{\text{pri}} < \Pi_{\text{avg}}^{\text{pub}} = \Pi_{\text{avg}}^{\text{no}}$.

The first noteworthy result of Theorem 2A is that, in each round, the ex ante expected effort of each solver is identical under a no-feedback and a public-feedback policy. Toward the end of better understanding this result, we remark that *public* feedback can have two opposing effects on a solver's effort choice in the second round. On the one hand, if the revealed first-round performance difference is relatively low ($|\Delta\zeta_1| < a/3$) then each solver understands that the contest is highly competitive and therefore is motivated to expend more effort than under a no-feedback policy. On the other hand, if the performance difference is relatively large ($|\Delta\zeta_1| > a/3$) then solvers are discouraged from investing effort because they believe that the contest is almost decided. In equilibrium, these countervailing effects of motivation and de-motivation offset each other owing to the symmetry of the random shock difference $\Delta\zeta_t$; thus, $\mathbb{E}_{\Delta\zeta_1}\left[e_2^{\text{pub}}(\Delta\zeta_1)\right] = e_2^{\text{no}}$. It is clear that, when deciding on his first-round solution effort, each solver anticipates this balance between the motivation and de-motivation effects and therefore chooses the same effort as under a no-feedback policy: $e_1^{\text{pub}} = e_1^{\text{no}}$.

In contrast, the announcement of *private* feedback reduces the willingness of solvers to expend solution effort as compared with both the no-feedback and public-feedback cases. Two different effects are responsible for this result. First, much as under a public-feedback policy, private feedback can motivate a solver to expend more effort than in the no-feedback case whenever he showed a middling performance in the first round.[8] However, this motivation effect is much less pronounced for private than for public feedback. To see why, recall that the motivation effect of public feedback is strongest when the firm communicates a small performance difference. Under private feedback, the firm never releases relative performance information and so each solver can only (and will) form a belief about the performance difference. Yet given the inherent randomness of performance, each solver knows that it is very unlikely his competitor has achieved a similar performance. As a result, solvers respond only moderately to the motivation effect of private feedback.

Second, and even worse, private feedback has a strong de-motivating effect on relatively ill-performing solvers. As Figure 3 illustrates, solvers with a bad first-round performance exert less effort in the second round than do solvers with a good first-round showing. Put differently, the anticipated performance gap between bad and good solvers widens in the second round because of

[8] This happens if and only if $-a\kappa(1+W_0(-\gamma_3 e^{-1+1/(3\kappa^2)}/(a\kappa))) < \zeta_{i1} < -a\kappa(1+W_{-1}(-\gamma_3 e^{-1+1/(3\kappa^2)}/(a\kappa)))$, where W_0 (resp. W_{-1}) is the upper (resp. lower) branch of the Lambert W function.

their asymmetric effort choices. As a result, a phenomenon arises that is not present under a public-feedback regime: a solver with a relatively bad first-round performance realizes that he may face a competitor that he can never beat. Hence the set of potential competitors against whom he can win shrinks and so the solver starts to shirk. In short, private feedback reduces the competitiveness of the contest, which in turn induces the solvers to reduce their effort.

This phenomenon also has a strong effect on a solver's first-round effort provision. Since effort in the second round is reduced, he is careful to refrain from wasting too much effort in the first round; that is why $e_1^{\text{pri}} < e_1^{\text{pub}}$. Thus the mere pre-announcement of *private* interim performance feedback has a negative effect on the solvers' expected behavior. This "strategic" effect is absent when announcing a public-feedback contest.

In sum: since maximizing the solvers' average performance is equivalent to maximizing the solvers' average effort provision, it follows that a private feedback policy always generates the lowest expected profits for the firm.

6.2. The Quest for the Best Solution

In practice, most innovation contests are designed to elicit one exceptional idea that promises significant value upon implementation. In this case, the firm focuses not on maximizing the solvers' average performance but rather on maximizing the performance of the best solution; that is, the firm maximizes $\Pi_{\text{best}} = \mathbb{E}[\max\{v_{i2}, v_{j2}\}] - A$. Theorem 2B reveals that, for certain types of innovation contests, private feedback is the optimal policy.

Theorem 2B (The Optimal Feedback Policy for Best Performance). *The following statements hold:*

(i) $\Pi_{\text{best}}^{\text{pub}} = \Pi_{\text{best}}^{\text{no}}$*;*

(ii) *There exists a* $\underline{\kappa} > 1$ *such that* $\Pi_{\text{best}}^{\text{pub}} > \Pi_{\text{best}}^{\text{pri}}$ *for all* $\kappa < \underline{\kappa}$*;*

(iii) *There exists a* $\overline{\kappa} < \infty$ *such that* $\Pi_{\text{best}}^{\text{pri}} > \Pi_{\text{best}}^{\text{pub}}$ *for all* $\kappa > \overline{\kappa}$*.*

Irrespective of whether the firm is interested in the solvers' average or best performance, employing a public-feedback policy generates the same expected profits as does a no-feedback policy. This result stems from the identity of expected effort under these two feedback policies (as established in Theorem 2A).

The key result of Theorem 2B is that public feedback dominates private feedback if $\kappa < \underline{\kappa}$ whereas private feedback is optimal if $\kappa > \overline{\kappa}$.[9] To better understand this result, recall that $\kappa = (a^2 c)/(A k_e^2)$;

[9] Unfortunately, the complexity of the equilibrium emerging under a private-feedback policy makes it hard to find the dominant strategy for medium κ. Investigating this question numerically, however, reveals that $\underline{\kappa} = \overline{\kappa}$; thus there is a unique threshold for the performance uncertainty κ above which a private-feedback policy maximizes the firm's expected profits.

Figure 4 Ex ante Expected Second-Round Efforts under Private and Public Feedback

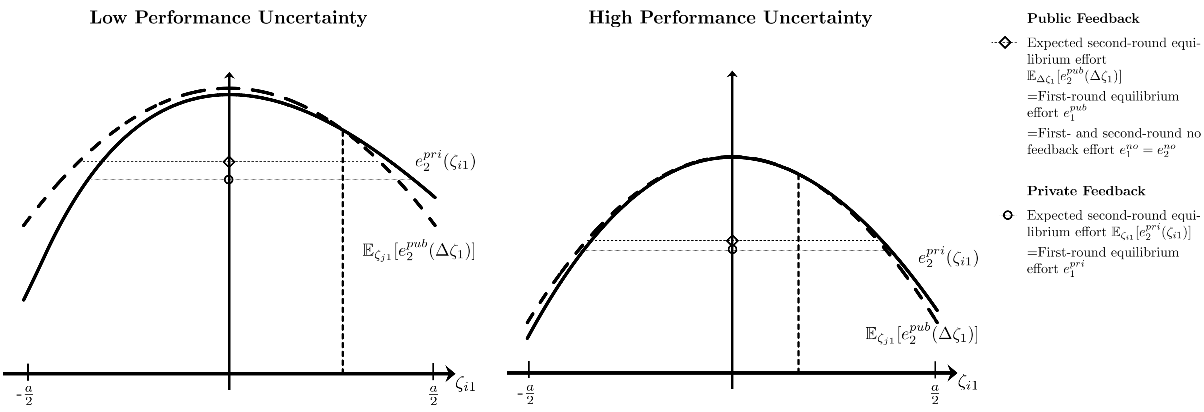

Note. The graphs compare solver i's (expected) equilibrium second-round effort conditional on ζ_{i1} under private feedback (solid line; $e_2^{\text{pri}}(\zeta_{i1})$) and under public feedback (dashed line; $\mathbb{E}_{\zeta_{j1}}[e_2^{\text{pub}}(\Delta\zeta_1)]$. In the left panel, the performance uncertainty is low ($\kappa = 1.01$); in the right panel, the performance uncertainty is high ($\kappa = 4$). The vertical dotted line marks the unique intersection point of the two curves. The parameters employed are: $A = 1$, $a = 1$, and $k_e = 1$; $c = 1.01$ (left panel), $c = 4$ (right panel).

that is, κ increases in the variance of the random noise, a, and the cost of effort, c, and κ decreases in the announced award, A, and the effort sensitivity, k_e. Reformulating $\kappa = (a/k_e)^2/(A/c)$ helps us interpret κ. The numerator of the reformulation is a measure of how uncertain the contest is; if a is large and k_e is low, then effort does not play a large role in winning the contest and hence uncertainty dominates. The denominator is a measure of profitability; if the prize is large and the cost is low, then profitability is high. Overall, then κ is a measure of how uncertain one unit of gain is for each of the solvers and thus it is a normalized measure of contest uncertainty. Hence, we denote κ as "performance uncertainty".[10] Then Theorem 2B implies that, for innovation contests in which effort is relatively more important than uncertainty (i.e., when $\kappa < \underline{\kappa}$), public feedback is optimal. In contrast, for innovation contests with substantial performance uncertainty ($\kappa > \overline{\kappa}$), private feedback outperforms public feedback.

But what is the intuition behind this result—particularly since, according to parts (i) and (ii) of Theorem 2A, private feedback induces lower average performance than public feedback (or no feedback)? We can answer this question by considering Figure 4, which compares solver i's expected second-round effort conditional on his first-round shock, ζ_{i1}, under the private-feedback (solid line) and public-feedback (dashed line) scenarios.[11] The figure's left (resp. right) panel shows

[10] An even simpler way to interpret the meaning of κ is based on recognizing that the quadratic numerator easily dominates the denominator in many situations, from which it follows that κ can be viewed simply as a measure of contest uncertainty.

[11] Under private feedback, solver i's expected second-round effort conditional on ζ_{i1} is simply $e_2^{\text{pri}}(\zeta_{i1})$, as stated in Proposition 3. Under public feedback, solver i's expected second-round effort conditional on ζ_{i1} is given by

the functions for low (resp. high) κ. There are two important observations to be made here. First, comparing the solid and the dashed lines plotted in each panel shows that public feedback induces a larger effort than does private feedback for most first-round shocks. This result is consistent with our finding that private feedback induces a lower ex ante expected effort than does public feedback. Moreover, comparing the two panels reveals that, for low performance uncertainty κ, the reduction in average effort under private feedback is much greater than for high performance uncertainty. Second, for top-performing solvers (i.e., solvers with a high first-round shock), private feedback increases effort provision: the solid line surpasses the dashed line for sufficiently high ζ_{i1}. This result reflects the need of top performers to protect their good position more fiercely under private than under public feedback owing to the lack of relative performance information. Additionally, Figure 4 also shows that the fraction of performers for which private feedback increases expended effort is small under low performance uncertainty but is substantial under high performance uncertainty.

Of course, it is exactly these top performers in whom the firm is interested when maximizing the performance of the *best* solution. So when using private feedback, the firm faces a non trivial trade-off. On the one hand, private feedback reduces the solvers' average effort provision; on the other hand, it fosters higher effort from the best solvers. Thus the optimal feedback policy is the one that best balances the average effort provision with the likelihood that a top-performing solver participates in the contest. Consider the left panel of Figure 4. For low κ, the decrease in average effort under private feedback is relatively pronounced and the likelihood of a top-performing solver (i.e., a solver with a first-round shock to the right of the dotted vertical line) participating in the contest is relatively low. As a result, public feedback dominates private feedback. In contrast, the right panel reveals that the reduction in average effort is much less pronounced for high κ. Furthermore, the chances that a solver exerts more effort under private than public feedback are much greater (i.e., the solid line crosses over the dashed line much more to the left). So in this case, private feedback is the optimal feedback policy.

Our finding that the optimal feedback policy is tightly linked to the relative importance of effort and uncertainty yields two immediate managerial implications. First, when setting up an innovation contest, it is crucial for the firm to identify the extent to which a solver's performance depends on stochasticity. For instance, there is seldom much uncertainty involved with contests that seek to foster incremental innovation. For such contests, the hosting firm should provide public feedback. In contrast, private feedback is the preferred choice for ideation contests that aim to develop

$\mathbb{E}_{\zeta_{j1}}[e_2^{\text{pub}}(\Delta\zeta_1)]$ for $e_2^{\text{pub}}(\Delta\zeta_1)$ as in Proposition 2. Note that we take the expectation over ζ_{j1} in the public-feedback case in order to make it directly comparable with the private-feedback scenario.

novel concepts, new ideas, or breakthrough research with high levels of uncertainty. Second, a firm should not always strive to maximize the solvers' expected efforts. In particular, if the performance uncertainty is substantial then the firm is better-off sacrificing average effort levels and instead maximizing only the effort levels of the best solvers.

From the viewpoint of social welfare, Theorem 2B (in conjunction with Theorem 2A) constitutes another important result for contests with substantial inherent uncertainties: the private feedback policy is not only optimal for the firm but also *socially efficient*. More precisely, a private-feedback policy maximizes the firm's expected profits and also allows solvers to reduce their expected efforts, leaving them with a higher expected utility. Thus, both the firm and the solvers prefer private feedback over public feedback in settings of high performance uncertainty. Hence the announcement of private feedback enables a firm to ex ante minimize the frictions between its own interests and those of the solvers.

7. Conclusions

Contests are a frequently used mechanism for providing incentives when companies source innovation from the outside. In fact, companies sometimes even use them to provide incentives on the inside. Prize competitions organized via the Internet are contests, as are many other efforts to procure innovative and tailored parts. Feedback has been extensively used in practice to improve both the efficiency and the efficacy of contests. However, our understanding of when and how to provide which kind of feedback—and of when to refrain from giving feedback—is limited. The aim of this paper is to begin building a more comprehensive understanding of feedback in contests.

As a first result, we have rigorously established the importance of commitment and truthfulness in feedback. Whatever its feedback policy, a firm that holds a contest must announce at the outset a feedback mapping (a translation of first-round performance to feedback signal) that transmits truthful information, and then it must adhere to that mapping. If these criteria are not met, then offering feedback is futile because it will simply be ignored in equilibrium. As a consequence, companies must either establish technical or organizational means to ensure commitment and truthfulness. An example of technical means is the establishment of clear and measurable success criteria that solvers can verify—for example, by testing their algorithms on a real data set. An example of organizational means is reputation building. A firm that is perceived as being trustworthy and fair will benefit from its reputation in terms of being able to influence the solvers; any diminution in its reputation would eliminate that benefit. This explains why, in practice, truthful pre-committed feedback is frequently observed as part of a trustful buyer–supplier relation.

Figure 5 Summary of Feedback Benefits

	Low performance uncertainty	**High performance uncertainty**
Best solution	no/public feedback	private feedback
Average solution	no/public feedback	no/public feedback

The main contribution of this paper is to extend current research on feedback in contests—which has concentrated on *public* feedback to raise *average* solution effort—along two dimensions. First, we study feedback in contests where the host is interested in the best performance only. Second, and most importantly, we introduce private feedback. By building such a general framework for analyzing feedback in contests, we can exhaustively identify the settings in which any of a wealth of practically relevant feedback policies is ideal. It is remarkable that, when deciding on their feedback policy, companies need to focus on only two simple dimensions: the contest's objective (average versus best performance) and the solvers' uncertainty about outcomes. If the firm is concerned about the solvers' average performance, then either no feedback or public feedback is the preferred option. The same preference obtains if the company seeks to improve the best performance when performance uncertainty is low. However, if the company wants the best possible performance but performance uncertainty is high, then private feedback outperforms public feedback as well as no feedback. The matrix presented in Figure 5 summarizes these results.

Our findings have immediate managerial implications. Contest holders that aim to raise the overall effort level among all solvers—as in the SAP Google Glass Challenge or the Microsoft Imagine Cup—should refrain from giving private feedback; if performance information is released, it should be made public. Incremental algorithmic innovation contests likewise do not benefit when private feedback is provided; for such contests, the relatively low performance uncertainty makes public feedback the preferred policy. In contrast, contests looking for breakthrough innovation (e.g., completely new algorithmic solutions, novel engineering concepts, any problem that requires the exploration of uncharted territories) should solely rely on private feedback.

Beyond these managerial implications, an intriguing aspect of the private-feedback case is its social efficiency. Not only does private feedback yield the best results from the firm's perspective (whenever mandated), it also lowers the solvers' expected effort and therefore maximizes social welfare.

Finally, this paper's contribution applies also to traditional procurement efforts. We contextualized contests mainly by examples from the information industry, but contests are by no means niche mechanisms to source algorithmic innovation. Many procurement efforts by traditional companies share the features of a contest. Consider a typical automotive manufacturer. For all parts that require adaptation and development (e.g., headlamps), the manufacturer invites a set of potential suppliers to create concept designs of what they think best suits the new car model under development. Often such concepts include technological innovations; in the case of headlamps, recent innovations include laser-based light and "curve adaptive" lighting. The car manufacturer provides feedback on each design—usually to the inventor only—and then asks suppliers to flesh out their inventions somewhat further, before it decides to engage in detailed development with one supplier. All of the suppliers must bear (at least partially) the cost of their efforts regardless of the contest outcome. It is most surprising that the extensive academic literature on efficient procurement settings (e.g., Engelbrecht-Wiggans and Katok 2006, Kostamis et al. 2009, Beil 2010, Wan et al. 2012, Tunca et al. 2014, Gupta et al. 2015) has not considered contests as sourcing mechanisms and has focused instead on auctions. This neglect is astonishing given that, for novel and innovative goods: (i) suppliers endogenously determine their performance by investing in development efforts; and (ii) contests are an efficient incentive mechanism for steering those efforts.

Previous research on contests has failed to comprehensively explore the repercussions of feedback. The aim and principal contribution of this paper is to fill that important gap. It is only by incorporating all forms of feedback into the analysis that managers can have a reasonable hope of making the contest mechanism—a method often relied upon in practice to source innovation—more efficient and effective.

Appendix. Proofs.

Proof of Theorem 1. As a preliminary step, note that solver i's expected utility $u_i = A \cdot \mathbb{P}(v_{i2} > v_{j2}) - \sum_t ce_{it}^2 < A - ce_{it}^2$ for all t. Since solver i only participates in the contest if his expected utility is non-negative, it follows that $e_{it} \leq \sqrt{A/c}$ for all t. Thus, the equilibrium emerging in the original contest is identical to the equilibrium of a contest where solver i's effort is restricted to the following compact set: $e_{it} \in [0, \sqrt{A/c}]$. With this equivalence result, three steps remain to prove the Theorem. First, we state and prove an important lemma that simplifies the subsequent argument. Second, we show that without commitment to a feedback mapping, feedback does not convey any information. Last, we establish that feedback must be indicative of a solver's true performance.

LEMMA A1. *For any transmitted feedback, there exists a unique pure-strategy second-round equilibrium.*

Proof of Lemma A1. We first prove the existence of a pure-strategy Nash equilibrium for any feedback, and then proceed to show the uniqueness of this equilibrium.

Step 1: Existence. Theorem 1.2 in Fudenberg and Tirole (1991, p. 34) ensures the existence of a pure-strategy Nash equilibrium if (i) each solver's strategy space is a nonempty compact convex subset of the Euclidian space, and (ii) each solver's expected utility, u_{i2}, is continuous in e_2 and quasi-concave in e_{i2}. The first condition is satisfied by our previous argument that the contest is equivalent to a contest with $e_{i2} \in [0, \sqrt{A/c}]$. Moreover, u_{i2} is clearly continuous in e_2. To verify the quasi-concavity of u_{i2} in e_{i2}, we now show that u_{i2} is strictly concave in e_{i2}. Given an arbitrary feedback $f_i \in \{\emptyset, s_i, (s_i, s_j)\}$, solver i's expected second-round utility is $u_{i2}(e_{i2}, e_{j2}|f_i) = A\mathbb{E}_{\zeta_1}\left[G_{\Delta\zeta_2}(\zeta_{i1} + k_e(e_{i1} + e_{i2}) - \zeta_{j1} - k_e(e_{j1} + e_{j2}))|f_i\right] - ce_{i2}^2$. Note that u_{i2} is twice continuously differentiable in e_{i2}, with the second-order partial derivative being

$$\frac{\partial^2 u_{i2}}{\partial e_{i2}^2} = Ak_e \frac{\partial}{\partial e_{i2}} \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(\zeta_{i1} + k_e(e_{i1} + e_{i2}) - \zeta_{j1} - k_e(e_{j1} + e_{j2}))|f_i\right] - 2c. \tag{10}$$

Since $\sup\left\{\partial\mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(\zeta_{i1} + k_e(e_{i1}^* + e_{i2}) - \zeta_{j1} - k_e(e_{j1}^* + e_{j2}))|f_i\right]/\partial e_{i2}\right\} = k_e/a^2$ and $a^2c > Ak_e^2$, it follows that $\partial^2 u_{i2}/\partial e_{i2}^2 \le Ak_e^2/a^2 - 2c < 0$.

Step 2: Uniqueness. We make use of Theorem 2 in Rosen (1965), which asserts that the Nash equilibrium is unique if $\sigma(e_2, r) \equiv r_i u_{i2}(e_2) + r_j u_{j2}(e_2)$ is diagonally strictly concave for some $r_i, r_j > 0$. According to Theorem 6 in Rosen (1965), this is true if $H(e_2, r) + H^T(e_2, r)$ is negative definite for any e_2 and some $r = (r_i, r_j) > 0$, where $H(e_2, r)$ is the Jacobian with respect to e_2 of the pseudogradient $h(e_2, r) = [r_i \partial u_{i2}/\partial e_{i2}, r_j \partial u_{j2}/\partial e_{j2}]^T$. Note that

$$H(e_2, r) + H^T(e_2, r) = \begin{bmatrix} 2r_i \frac{\partial^2 u_{i2}}{\partial e_{i2}^2} & r_i \frac{\partial^2 u_{i2}}{\partial e_{i2}\partial e_{j2}} + r_j \frac{\partial^2 u_{j2}}{\partial e_{j2}\partial e_{i2}} \\ r_i \frac{\partial^2 u_{i2}}{\partial e_{i2}\partial e_{j2}} + r_j \frac{\partial^2 u_{j2}}{\partial e_{j2}\partial e_{i2}} & 2r_j \frac{\partial^2 u_{j2}}{\partial e_{j2}^2} \end{bmatrix}, \tag{11}$$

which is negative definite if $\det\left(H(e_2, r) + H^T(e_2, r)\right) \ge 0$ for any e_2. It is readily seen that $\partial^2 u_{i2}/\partial e_{i2}^2 + 2c = -\partial^2 u_{i2}/\partial e_{i2}\partial e_{j2}$, and $\inf\left\{\partial\mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(\zeta_{i1} + k_e(e_{i1}^* + e_{i2}) - \zeta_{j1} - k_e(e_{j1}^* + e_{j2}))|f_i\right]/\partial e_{i2}\right\} = -k_e/a^2$. With $r_i = r_j$, it follows that $\det\left(H(e_2, r) + H^T(e_2, r)\right) \ge 0$ if $r_i > (Ak_e^2)^2/(8a^2c(a^2c - Ak_e^2)) > 0$, which concludes the proof. □

We are now ready to establish Theorem 1 in two separate steps.

Step 1: Commitment. The key step is to show that without a credible commitment to a specific feedback mapping, the firm's feedback is always independent of the actual value of the solvers' first-round submissions. We show this for the case where the firm is solely interested in the best solution. An identical argument applies to the average performance case as well. Clearly, after observing v_1, the firm chooses the feedback $s = (s_i, s_j)$ that maximizes expected profits; i.e.,

$$s^* \in \arg\max_s \mathbb{E}\left[\max\{v_{i1} + \zeta_{i2} + k_e e_{i2}(f_i), v_{j1} + \zeta_{j2} + k_e e_{j2}(f_j)\}|v_1\right]. \tag{12}$$

Trivially, under a no feedback policy, solvers do not observe s, i.e., $f_i = f_j = \emptyset$, and therefore, no information is transmitted. We now demonstrate that the same is true for public and private feedback.

Public feedback: Suppose that the firm gives public feedback; i.e., $f_i = f_j = (s_i, s_j)$. Also assume that solver i believes that the firm's feedback contains some information about v_1. (Otherwise, if solver i believes that feedback does not convey any information, then feedback would be meaningless.) Then, upon observing f_i, solver i chooses the effort that satisfies his necessary and sufficient first-order optimality condition:

$$\frac{\partial u_{i2}}{\partial e_{i2}} = Ak_e \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(\zeta_{i1} + k_e(e^*_{i1} + e_{i2}) - \zeta_{j1} - k_e(e^*_{j1} + e_{j2}))|f_i\right] - 2ce_{i2} = 0. \tag{13}$$

By the symmetry of $g_{\Delta\zeta_2}$ around zero, it follows that both solvers choose the same second-round effort, $e_{i2}(s_i, s_j) = e_{j2}(s_i, s_j)$. Therefore, we can rewrite the firm's problem (12) as follows:

$$s^* \in \arg\max_s \mathbb{E}\left[k_e e_{i2}(s) + \max\{v_{i1} + \zeta_{i2}, v_{j1} + \zeta_{j2}\}|v_1\right] \tag{14}$$

$$= \arg\max_s k_e e_{i2}(s) + \mathbb{E}\left[\max\{v_{i1} + \zeta_{i2}, v_{j1} + \zeta_{j2}\}|v_1\right] \tag{15}$$

$$= \arg\max_s k_e e_{i2}(s). \tag{16}$$

By (16), it can be readily seen that the firm chooses the feedback that maximizes the solvers' second-round effort, irrespective of the realization of v_1; thereby implying that, in equilibrium, feedback does not convey any information about v_1.

Private feedback: Suppose that the firm gives private feedback; i.e., $f_i = s_i$. Also, assume again that solver i believes that the firm's feedback contains some information about v_{i1}. Since s_i never conveys information about v_{j1}, solver i's effort choice is independent of v_{j1}. This implies that the firm's problem (12) decomposes into two separate problems for each solver respectively; i.e.,

$$s_i^* \in \arg\max_{s_i} \mathbb{E}\left[k_e e_{i2}(s_i) + v_{i1} + \zeta_{i2}|v_{i1}\right] \tag{17}$$

$$= \arg\max_{s_i} k_e e_{i2}(s_i). \tag{18}$$

Again, (18) reveals that, in equilibrium, the firm chooses the feedback s_i to maximize solver i's second-round effort, without revealing any information about v_{i1}.

Step 2: Truthfulness. Suppose that the firm is committed to a feedback mapping (r, Σ), but this mapping is not indicative of a solver's true performance; i.e., $r(s_i|v_{i1}) = r(s_i|v'_{i1})$ for all s_i, v_{i1}, and v'_{i1}. Then, by Bayes' rule, a solver's posterior belief about v_{i1} is given by $g(v_{i1}|s_i) = r(s_i|v_{i1})g(v_{i1})/\int r(s_i|v'_{i1})g(v'_{i1})dv'_{i1} = r(s_i|v_{i1})g(v_{i1})/r(s_i|v_{i1}) = g(v_{i1})$; i.e., no belief updating is possible. Thus, feedback does not influence a solver's effort choice. In contrast, if $r(s_i|v_{i1}) \neq r(s_i|v'_{i1})$ for some s_i, v_{i1}, and v'_{i1}, then posterior beliefs differ from prior beliefs, and therefore, a solver adjusts his effort choice according to the received feedback. □

Proof of Proposition 1. Without feedback between round one and two, solver i's optimization problem is equivalent to a utility maximization problem where he simultaneously decides on both effort levels, e_{i1} and e_{i2}. Moreover, since performance is linear in effort, while the costs are strictly convex, equilibrium effort levels must be the same in both rounds. Thus, in equilibrium, $e_{i1} = e_{i2} = e_i^{\text{no}}$, and solver i's equilibrium effort has to solve

$$e_i^{\text{no}} \in \arg\max_{e_i} A \cdot \mathbb{E}_{\zeta_1}[G_{\Delta\zeta_2}(2k_e e_i + \zeta_{i1} - 2k_e e_j - \zeta_{j1})] - 2ce_i^2, \tag{19}$$

and the corresponding necessary and sufficient first-order optimality condition is given by

$$2Ak_e \cdot \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(2k_e e_i^{\text{no}} + \zeta_{i1} - 2k_e e_j^{\text{no}} - \zeta_{j1})\right] = 4ce_i^{\text{no}}. \tag{20}$$

By the symmetry of $g_{\Delta\zeta_2}$ around zero, it follows readily that the unique solution to the solvers' optimality conditions is symmetric; i.e., $e_i^{\text{no}} = e_j^{\text{no}}$. Inserting this information in (20) yields $e_i^{\text{no}} = Ak_e/(2c) \cdot \mathbb{E}_{\zeta_1}[g_{\Delta\zeta_2}(\zeta_{i1} - \zeta_{j1})] = Ak_e/(3ac)$. □

Proof of Proposition 2. Given truthful public feedback, the solvers perfectly learn v_1 after round one. As such, solver i's second-round equilibrium effort solves

$$e_{i2}^{\text{pub}} \in \arg\max_{e_{i2}} A \cdot G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2}) - ce_{i2}^2, \tag{21}$$

and the corresponding necessary and sufficient first-order optimality condition is given by

$$Ak_e \cdot g_{\Delta\zeta_2}(v_{i1} + k_e e_{i2}^{\text{pub}} - v_{j1} - k_e e_{j2}^{\text{pub}}) = 2ce_{i2}^{\text{pub}}. \tag{22}$$

By the symmetry of $g_{\Delta\zeta_2}$ around zero, it follows immediately that the unique second-round equilibrium is symmetric, $e_{i2}^{\text{pub}} = e_{j2}^{\text{pub}}$.

In the first round, solver i's equilibrium effort has to solve

$$e_{i1}^{\text{pub}} \in \arg\max_{e_{i1}} A \cdot \mathbb{E}_{\zeta_1}\left[G_{\Delta\zeta_2}(k_e(e_{i1} + e_{i2}^{\text{pub}}) + \zeta_{i1} - k_e(e_{j1} + e_{j2}^{\text{pub}}) - \zeta_{j1})\right] - ce_{i1}^2 - \mathbb{E}_{\zeta_1}\left[c(e_{i2}^{\text{pub}})^2\right] \tag{23}$$

$$= A \cdot \mathbb{E}_{\zeta_1}[G_{\Delta\zeta_2}(k_e e_{i1} + \zeta_{i1} - k_e e_{j1} - \zeta_{j1})] - ce_{i1}^2 - \mathbb{E}_{\zeta_1}\left[c(e_{i2}^{\text{pub}})^2\right], \tag{24}$$

and the corresponding necessary first-order optimality condition is given by

$$Ak_e \cdot \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(k_e e_{i1}^{\text{pub}} + \zeta_{i1} - k_e e_{j1}^{\text{pub}} - \zeta_{j1})\right] - 2ce_{i1}^{\text{pub}} - \frac{\partial}{\partial e_{i1}}\mathbb{E}_{\zeta_1}\left[c(e_{i2}^{\text{pub}})^2\right] = 0. \tag{25}$$

Note that the first two terms in (25) capture the direct effect of e_{i1}^{pub} on solver i's expected utility, whereas the third term captures the indirect effect of e_{i1}^{pub} on his own second-round effort e_{i2}^{pub}. In equilibrium, this indirect effect must be zero. To see this, note that (25) reveals that e_{i1}^{pub} has no strategic effect on e_{j2}^{pub}. By the symmetry of the second-round equilibrium, this implies that, in equilibrium, the strategic effect of e_{i1}^{pub} on e_{i2}^{pub} has to be zero as well. Yet, this is true if and only if $e_{i1}^{\text{pub}} = e_{j1}^{\text{pub}}$; i.e., first-round equilibrium efforts are symmetric. Inserting this information in (22) and (25) shows that the unique PBE under public feedback is given by $e_{i2}^{\text{pub}} = Ak_e/(2c) \cdot g_{\Delta\zeta_2}(\zeta_{i1} - \zeta_{j1}) = Ak_e/(2a^2c) \cdot (a - |\zeta_{i1} - \zeta_{j1}|)$, and $e_{i1}^{\text{pub}} = Ak_e/(2c) \cdot \mathbb{E}_{\zeta_1}[g_{\Delta\zeta_2}(\zeta_{i1} - \zeta_{j1})] = \mathbb{E}_{\zeta_1}\left[e_{i2}^{pub}\right] = Ak_e/(3ac)$. □

Proof of Proposition 3. In a first step, we establish uniqueness of the first-round equilibrium. Given truthful private feedback, solver i perfectly learns v_{i1} after round one, but receives no additional information on v_{j1}. As such, solver i's second-round equilibrium effort solves

$$e_{i2}^{\text{pri}} \in \arg\max_{e_{i2}} A \cdot \mathbb{E}_{v_{j1}}[G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2})|v_{i1}] - ce_{i2}^2, \tag{26}$$

and the corresponding necessary and sufficient first-order optimality condition is given by

$$Ak_e \cdot \mathbb{E}_{v_{j1}}[g_{\Delta\zeta_2}(v_{i1} + k_e e_{i2}^{\text{pri}} - v_{j1} - k_e e_{j2}^{\text{pri}})|v_{i1}] = 2ce_{i2}^{\text{pri}}. \tag{27}$$

Lemma A1 ensures that the second-round equilibrium defined by (27) is unique. In the first round, solver i's equilibrium effort has to solve

$$e_{i1}^{\text{pri}} \in \arg\max_{e_{i1}} A \cdot \mathbb{E}_{\zeta_1}\left[G_{\Delta\zeta_2}(k_e(e_{i1}+e_{i2}^{\text{pri}})+\zeta_{i1}-k_e(e_{j1}+e_{j2}^{\text{pri}})-\zeta_{j1})\right]-ce_{i1}^2-\mathbb{E}_{\zeta_1}\left[c(e_{i2}^{\text{pri}})^2\right], \quad (28)$$

and the corresponding necessary first-order optimality condition is given by

$$Ak_e \cdot \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(k_e(e_{i1}^{\text{pri}}+e_{i2}^{\text{pri}})+\zeta_{i1}-k_e(e_{j1}^{\text{pri}}+e_{j2}^{\text{pri}})-\zeta_{j1})\cdot\left(1+\frac{\partial e_{i2}^{\text{pri}}}{\partial e_{i1}}-\frac{\partial e_{j2}^{\text{pri}}}{\partial e_{i1}}\right)\right]-2ce_{i1}^{\text{pri}}- \mathbb{E}_{\zeta_1}\left[2ce_{i2}^{\text{pri}}\cdot\frac{\partial e_{i2}^{\text{pri}}}{\partial e_{i1}}\right]=0. \quad (29)$$

Clearly, solver j's second-round effort cannot be influenced by solver i's first-round effort, because solver j does not receive any information on v_{i1}. Therefore, $\partial e_{j2}^{\text{pri}}/\partial e_{i1}=0$. Rewriting (29) yields

$$\begin{aligned}&Ak_e\cdot\mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(k_e(e_{i1}^{\text{pri}}+e_{i2}^{\text{pri}})+\zeta_{i1}-k_e(e_{j1}^{\text{pri}}+e_{j2}^{\text{pri}})-\zeta_{j1})\right]-2ce_{i1}^{\text{pri}}\\&\quad+\mathbb{E}_{\zeta_1}\left[\left(Ak_eg_{\Delta\zeta_2}(k_e(e_{i1}^{\text{pri}}+e_{i2}^{\text{pri}})+\zeta_{i1}-k_e(e_{j1}^{\text{pri}}+e_{j2}^{\text{pri}})-\zeta_{j1})-2ce_{i2}^{\text{pri}}\right)\cdot\frac{\partial e_{i2}^{\text{pri}}}{\partial e_{i1}}\right]=0,\end{aligned}$$

where the third term is zero because

$$\begin{aligned}&\mathbb{E}_{\zeta_1}\left[\left(Ak_eg_{\Delta\zeta_2}(k_e(e_{i1}^{\text{pri}}+e_{i2}^{\text{pri}})+\zeta_{i1}-k_e(e_{j1}^{\text{pri}}+e_{j2}^{\text{pri}})-\zeta_{j1})-2ce_{i2}^{\text{pri}}\right)\cdot\frac{\partial e_{i2}^{\text{pri}}}{\partial e_{i1}}\right]\\&=\mathbb{E}_{v_{i1}}\left[\mathbb{E}_{v_{j1}}\left[\left(Ak_eg_{\Delta\zeta_2}(k_e(e_{i1}^{\text{pri}}+e_{i2}^{\text{pri}})+\zeta_{i1}-k_e(e_{j1}^{\text{pri}}+e_{j2}^{\text{pri}})-\zeta_{j1})-2ce_{i2}^{\text{pri}}\right)\cdot\frac{\partial e_{i2}^{\text{pri}}}{\partial e_{i1}}\,\middle|\,v_{i1}\right]\right]\\&=\mathbb{E}_{v_{i1}}\left[\frac{\partial e_{i2}^{\text{pri}}}{\partial e_{i1}}\cdot\mathbb{E}_{v_{j1}}\left[Ak_eg_{\Delta\zeta_2}(k_e(e_{i1}^{\text{pri}}+e_{i2}^{\text{pri}})+\zeta_{i1}-k_e(e_{j1}^{\text{pri}}+e_{j2}^{\text{pri}})-\zeta_{j1})-2ce_{i2}^{\text{pri}}\cdot\,|\,v_{i1}\right]\right]\\&=\mathbb{E}_{v_{i1}}\left[\frac{\partial e_{i2}^{\text{pri}}}{\partial e_{i1}}\cdot\left(Ak_e\cdot\mathbb{E}_{v_{j1}}\left[g_{\Delta\zeta_2}(v_{i1}+k_ee_{i2}^{\text{pri}}-v_{j1}-k_ee_{j2}^{\text{pri}}))|v_{i1}\right]-2ce_{i2}^{\text{pri}}\right)\right]\\&=0.\end{aligned}$$

The first equality follows from the law of iterated expectations, the second equality is true because solver i's second-round effort choice is independent of v_{j1}, the third equality follows from rearranging terms, and the last equality follows from solver i's second-round optimality condition (27). Thus, the first-order optimality condition of solver i is

$$Ak_e\cdot\mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(k_e(e_{i1}^{\text{pri}}+e_{i2}^{\text{pri}})+\zeta_{i1}-k_e(e_{j1}^{\text{pri}}+e_{j2}^{\text{pri}})-\zeta_{j1})\right]-2ce_{i1}^{\text{pri}}=0,$$

and by the symmetry of $g_{\Delta\zeta_2}$ around zero, it follows readily that the unique solution to the solvers' optimality conditions is symmetric, $e_{i1}^{\text{pri}}=e_{j1}^{\text{pri}}$, and $e_1^{\text{pri}}=\mathbb{E}_{\zeta_{i1}}\left[e_{i2}^{\text{pri}}\right]$.

We now proceed with deriving the solvers' second-round equilibrium effort. We conjecture that the unique second-round equilibrium is symmetric in the sense that $e_{i2}^{\text{pri}}=e_2^{\text{pri}}(\zeta_{i1})$ and $e_{j2}^{\text{pri}}=e_2^{\text{pri}}(\zeta_{j1})$, and that $v^{\text{pri}}(\zeta_{i1})$ increases in ζ_{i1}. We will demonstrate in retrospective that these claims are true. Together with (27), the above properties imply that the equilibrium effort function, $e_2^{\text{pri}}(\cdot)$, solves the following integral equation:

$$Ak_e\cdot\mathbb{E}_{\zeta_{j1}}[g_{\Delta\zeta_2}(\zeta_{i1}+k_ee_2^{\text{pri}}(\zeta_{i1})-\zeta_{j1}-k_ee_2^{\text{pri}}(\zeta_{j1}))|\zeta_{i1}]=2ce_2^{\text{pri}}(\zeta_{i1}), \quad (30)$$

or equivalently,

$$Ak_e^2 \cdot \mathbb{E}_{\zeta_{j1}}[g_{\Delta\zeta_2}(v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}))|\zeta_{i1}] = 2c\left(v^{\text{pri}}(\zeta_{i1}) - \zeta_{i1}\right). \tag{31}$$

Because $g_{\Delta\zeta_2}(v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}))$ is positive only if $v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}) \in [-a, a]$, we need to consider three different cases. (1) If $v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}) \in [-a, a]$ for all ζ_{j1}, then $\zeta_{i1} \in [\zeta_u, \zeta_o]$. In this case, (31) is given by

$$\int_{-\frac{a}{2}}^{\zeta_{i1}} (a - v^{\text{pri}}(\zeta_{i1}) + v^{\text{pri}}(\zeta_{j1}))d\zeta_{j1} + \int_{\zeta_{i1}}^{\frac{a}{2}} (a + v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}))d\zeta_{j1} = 2a\kappa(v^{\text{pri}}(\zeta_{i1}) - \zeta_{i1}), \tag{32}$$

and differentiating both sides with respect to ζ_{i1} leads to the following first-order ordinary differential equation:

$$\left(v^{\text{pri}}(\zeta_{i1})\right)' = \frac{a\kappa}{\zeta_{i1} + a\kappa}, \tag{33}$$

with the canonical solution

$$v^{-1}(v^{\text{pri}}) = \gamma_3 e^{\frac{1}{a\kappa}v^{\text{pri}}} - a\kappa. \tag{34}$$

Given (34), it is easy to verify that $v^{\text{pri}}(a/2) - v^{\text{pri}}(-a/2) = a\kappa \ln((2\kappa + 1)/(2\kappa - 1)) > a$, implying that $[\zeta_u, \zeta_o] \subset [-a/2, a/2]$.

(2) For $\zeta_{i1} \in [-a/2, \zeta_u]$, (31) becomes

$$\int_{-\frac{a}{2}}^{\zeta_{i1}} (a - v^{\text{pri}}(\zeta_{i1}) + v^{\text{pri}}(\zeta_{j1}))d\zeta_{j1} + \int_{\zeta_{i1}}^{v^{-1}(v^{\text{pri}}(\zeta_{i1})+a)} (a + v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}))d\zeta_{j1} = 2a\kappa(v^{\text{pri}}(\zeta_{i1}) - \zeta_{i1}), \tag{35}$$

and differentiating both sides with respect to ζ_{i1} leads to the following differential equation:

$$\left(v^{\text{pri}}(\zeta_{i1})\right)' \cdot \left[2\zeta_{i1} + 2a\kappa + \frac{a}{2} - v^{-1}(v^{\text{pri}}(\zeta_{i1}) + a)\right] - 2a\kappa = 0, \tag{36}$$

which is a Bernoulli equation in $v^{-1}(\cdot)$. The implicit solution to (36) is

$$v^{-1}(v^{\text{pri}}) = Ce^{\frac{1}{a\kappa}v^{\text{pri}}} - a\left(\kappa + \frac{1}{4}\right) - \frac{1}{2a\kappa}e^{\frac{1}{a\kappa}v^{\text{pri}}} \int e^{-\frac{1}{a\kappa}v^{\text{pri}}} v^{-1}(v^{\text{pri}} + a)dv^{\text{pri}}. \tag{37}$$

(3) In a similar vein, we can show that for $\zeta_{i1} \in [\zeta_o, a/2]$, the implicit solution to (31) is given by

$$v^{-1}(v^{\text{pri}}) = C'e^{\frac{1}{a\kappa}v^{\text{pri}}} - a\left(\kappa - \frac{1}{4}\right) - \frac{1}{2a\kappa}e^{\frac{1}{a\kappa}v^{\text{pri}}} \int e^{-\frac{1}{a\kappa}v^{\text{pri}}} v^{-1}(v^{\text{pri}} - a)dv^{\text{pri}}. \tag{38}$$

Combining (37) and (38) allows us to derive closed-form solutions. With (38), (37) becomes

$$v^{-1}(v^{\text{pri}}) = Ce^{\frac{1}{a\kappa}v^{\text{pri}}} - a\left(\frac{3}{2}\kappa + \frac{1}{8}\right) - C'\frac{1}{2a\kappa}e^{\frac{1}{\kappa}}e^{\frac{1}{a\kappa}v^{\text{pri}}}v^{\text{pri}} + \left(\frac{1}{2a\kappa}\right)^2 e^{\frac{1}{a\kappa}v^{\text{pri}}} \iint e^{-\frac{1}{a\kappa}v^{\text{pri}}} v^{-1}(v^{\text{pri}})d(v^{\text{pri}} + a)dv^{\text{pri}}. \tag{39}$$

Note that $v^{-1}(v^{\text{pri}}) - 2a\kappa(v^{-1}(v^{\text{pri}}))' + (a\kappa)^2(v^{-1}(v^{\text{pri}}))'' = v^{-1}(v^{\text{pri}})/4 - a(3\kappa/2 + 1/8)$. This is an equation of damped vibrations with canonical solution

$$v^{-1}(v^{\text{pri}}) = \gamma_1 e^{\frac{3}{2a\kappa}v^{\text{pri}}} + \gamma_2 e^{\frac{1}{2a\kappa}v^{\text{pri}}} - a\left(\frac{1}{6} + 2\kappa\right). \tag{40}$$

In an identical way, we can derive the canonical solution for $\zeta_{i1} \in [\zeta_o, a/2]$:

$$v^{-1}(v^{\text{pri}}) = \gamma_4 e^{\frac{3}{2a\kappa}v^{\text{pri}}} + \gamma_5 e^{\frac{1}{2a\kappa}v^{\text{pri}}} + a\left(\frac{1}{6} - 2\kappa\right). \tag{41}$$

Thus, the canonical solution to (31) is given by

$$v^{-1}(v^{\text{pri}}) = \begin{cases} \gamma_1 e^{\frac{3}{2a\kappa}v^{\text{pri}}} + \gamma_2 e^{\frac{1}{2a\kappa}v^{\text{pri}}} - a\left(\frac{1}{6}+2\kappa\right) & \text{if } v_o - a \le v^{\text{pri}} < v_u \\ \gamma_3 e^{\frac{1}{a\kappa}v^{\text{pri}}} - a\kappa & \text{if } v_u \le v^{\text{pri}} \le v_o \\ \gamma_4 e^{\frac{3}{2a\kappa}v^{\text{pri}}} + \gamma_5 e^{\frac{1}{2a\kappa}v^{\text{pri}}} + a\left(\frac{1}{6}-2\kappa\right) & \text{if } v_o < v^{\text{pri}} \le v_u + a, \end{cases} \tag{42}$$

where $v_u = v^{\text{pri}}(\zeta_u)$, $v_o = v^{\text{pri}}(\zeta_o)$, $v_o - a = v^{\text{pri}}(-a/2)$, and $v_u + a = v^{\text{pri}}(a/2)$. It remains to determine all integration constants γ_k, v_u and v_o. From (36), it follows readily that $\gamma_1 = -\gamma_4 n^3$ and $\gamma_2 = \gamma_5 n$. Moreover, (31) satisfies all requirements of the Implicit Function Theorem. Therefore, $v^{\text{pri}}(\zeta_{i1})$ is continuously differentiable, and so is $v^{-1}(v^{\text{pri}})$. From the continuity of $(v^{-1}(v^{\text{pri}}))'$, it follows that $\gamma_2 = 2\gamma_3 xy(n^3x+y)/(n^2x^2+y^2)$, and $3\gamma_1 = 2\gamma_3 n^2(x-ny)/(n^2x^2+y^2)$. Additionally, the continuity of $v^{-1}(v^{\text{pri}})$ implies that $\gamma_3 = 3a(\kappa + 1/6)(n^2x^2+y^2)/(x^2(3y^2-n^2x^2+4n^3xy))$, and $\gamma_3 = 3a(\kappa-1/6)n(n^2x^2+y^2)/(y^2(3n^3x^2-ny^2+4xy))$. Equating these two expressions leads to (6). Finally, the integral equation (31) becomes

$$\frac{a}{6\kappa} + a\left(\frac{1}{6}-\kappa\right)\ln(y) - a\left(\frac{1}{6}+\kappa\right)\ln(x) - a\kappa + \frac{2}{3}\gamma_1\left(\left(\frac{y}{n}\right)^3 - x^3\right) + \frac{1}{2}\gamma_3(y^2+x^2) = 0, \tag{43}$$

which is the same as (7). As a last step, we need to verify that our initial conjecture that $v^{\text{pri}}(\zeta_{i1})$ increases in ζ_{i1} is true. Note that because $0 < x < y$, and $n > 1$, we have $\gamma_1 < 0$, and $\gamma_2, \gamma_3, \gamma_4, \gamma_5 > 0$. Thus, it is obvious that $v^{-1}(v^{\text{pri}})$ increases in v^{pri} for $v^{\text{pri}} \ge v_u$. For $v^{\text{pri}} < v_u$, we have $(v^{-1}(v^{\text{pri}}))' > 0$ if $3\gamma_1 x^2 + \gamma_2 = 2\gamma_3 x > 0$, which is true. Therefore, $v^{\text{pri}}(\zeta_{i1})$ increases in ζ_{i1}, which concludes the proof.

At this point, it is worthwhile to briefly comment on our methodology for finding the solution to the integral equation (31). The crucial step is to transform the integral equation into an ordinary differential equation (ODE) by differentiating both sides of the equality with respect to ζ_{i1}. Clearly, the unique solution to (31) also solves the ODE. However, the ODE may have solutions that do not solve (31). The only way to circumvent this problem is to identify the canonical solution of the ODE, which is given by (42). Having the canonical solution allows us to conclude that the solution to (31) has the same structural form, and it remains to determine the integration constants appropriately by exploiting the properties of the original integral equation (31). As a result, it is guaranteed that the system of equations (6)-(7) admits a unique solution for $x \in \left[e^{-1/(4\kappa)}, e^{-1/(4\kappa)\cdot(1-1/\kappa)}\right]$ and $y \in \left[e^{1/(4\kappa)}, e^{1/(4\kappa)\cdot(1+1/\kappa)}\right]$. □

Proof of Corollary 1. Let $\widetilde{x} = e^{-(\kappa-1)/(4\kappa^2)} = n^{-(\kappa-1)/(2\kappa)}$ and $\widetilde{y} = e^{(\kappa+1)/(4\kappa^2)} = n^{(\kappa+1)/(2\kappa)}$. We now show that $(\widetilde{x}, \widetilde{y})$ is the solution to the system of equations (6)-(7) as $\kappa \to \infty$. Inserting $\widetilde{x}$ and $\widetilde{y}$ in (6) reveals that the left-hand side is equal to $-2n^{2/\kappa}\cdot(2+n^2+m(1+2n^2))$, which converges to zero as $\kappa \to \infty$, because $\lim_{\kappa\to\infty} n = 1$, and $\lim_{\kappa\to\infty} m = -1$. Similarly, the left-hand side of (7) is given by $(1-6\kappa^2)/(\kappa(1+6\kappa)) + (m+1)/(4\kappa) + (m-1)/(4\kappa^2) + 3(1+n^2)/(2(1+2n^2))$, which clearly converges to zero as $\kappa \to \infty$. Moreover, $\lim_{\kappa\to\infty} v_u = \lim_{\kappa\to\infty} 2a\kappa\ln(\widetilde{x}) = -a/2$, and $\lim_{\kappa\to\infty} v_o = \lim_{\kappa\to\infty} 2a\kappa\ln(\widetilde{y}) = a/2$. From (5), it follows that only the middle sector persists as $\kappa \to \infty$. Also, by inserting $\widetilde{x}$ and $\widetilde{y}$ in the formula for γ_3 in Proposition 3, it is easy to see that $\lim_{\kappa\to\infty} \widetilde{\gamma}_3 = \lim_{\kappa\to\infty} \gamma_3$. Taken together, these results imply that $\lim_{\kappa\to\infty} \widetilde{e}_2(\zeta_{i1}) = \lim_{\kappa\to\infty} e_2^{\text{pri}}(\zeta_{i1})$ for all ζ_{i1}. □

Proof of Theorem 2A. (i) From Propositions 1 and 2, it follows readily that $e_1^{\text{no}} = e_1^{\text{pub}}$. It remains to show that $e_1^{\text{no}} > e_1^{\text{pri}}$. Note that (27) reveals that $e_2^{\text{pri}}(\zeta_{i1}) < Ak_e/(2ac)$ for all ζ_{i1}. Thus, $0 \leq e_1^{\text{pri}} = \mathbb{E}_{\zeta_{i1}}\left[e_2^{\text{pri}}(\zeta_{i1})\right] < Ak_e/(2ac)$. It follows that $\lim_{a\to\infty} e_1^{\text{no}} = \lim_{a\to\infty} e_1^{\text{pri}} = 0$. Furthermore, $\partial e_1^{\text{no}}/\partial a = -e_1^{\text{no}}/a$, and $\partial e_1^{\text{pri}}/\partial a = -e_1^{\text{pri}}/a + \partial(a e_1^{\text{pri}})/\partial a > -e_1^{\text{pri}}/a$. As a result, $e_1^{\text{no}} = e_1^{\text{pri}} = 0$ for $a \to \infty$, but $\partial e_1^{\text{pri}}/\partial a > \partial e_1^{\text{no}}/\partial a$; i.e., e_1^{pri} decreases less steeply than e_1^{no}. This implies that $e_1^{\text{no}} > e_1^{\text{pri}}$ if a becomes an $\epsilon > 0$ smaller. But if $e_1^{\text{no}} > e_1^{\text{pri}}$, then e_1^{pri} decreases even less steeply compared to e_1^{no}. By an inductive argument, it follows that $e_1^{\text{no}} - e_1^{\text{pri}} > 0$, and this difference decreases in a.

(ii) The result is a direct consequence of (i) in combination with Propositions 1 - 3.

(iii) By (i) and (ii), it follows that $\Pi_{\text{avg}}^{\text{pub}} = \mathbb{E}\left[v_{i2}^{\text{pub}} + v_{j2}^{\text{pub}}\right]/2 - A = \left(e_1^{\text{pub}} + \mathbb{E}_{\Delta\zeta_1}\left[e_2^{\text{pub}}(\Delta\zeta_1)\right]\right) - A = (e_1^{\text{no}} + e_2^{\text{no}}) - A = \Pi_{\text{avg}}^{\text{no}} > \Pi_{\text{avg}}^{\text{pri}} = \left(e_1^{\text{pri}} + \mathbb{E}_{\zeta_{i1}}\left[e_2^{\text{pri}}(\zeta_{i1})\right]\right) - A.$ □

Proof of Theorem 2B. (i) The result follows directly from comparing the firm's expected profits under the two different feedback policies:

$$\begin{aligned}
\Pi_{\text{best}}^{\text{no}} &= \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2} + 2k_e e_i^{\text{no}}\}\right] - A = 2k_e e_1^{\text{no}} + \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2}\}\right] - A = a\left(\frac{2}{3\kappa} + \frac{7}{30}\right) - A \\
\Pi_{\text{best}}^{\text{pub}} &= \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2} + k_e e_1^{\text{pub}} + k_e e_2^{\text{pub}}(\zeta_1)\}\right] - A = k_e e_1^{\text{pub}} + \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2} + k_e e_2^{\text{pub}}(\zeta_1)\}\right] - A \\
&= k_e e_1^{\text{pub}} + \frac{1}{2}\cdot\mathbb{E}\left[\zeta_{i1} + \zeta_{i2} + 2k_e e_2^{\text{pub}}(\zeta_1) + \zeta_{j1} + \zeta_{j2} + |\zeta_{i1} + \zeta_{i2} - \zeta_{j1} - \zeta_{j2}|\right] - A \\
&= 2k_e e_1^{\text{pub}} + \frac{1}{2}\cdot\mathbb{E}\left[\zeta_{i1} + \zeta_{i2} + \zeta_{j1} + \zeta_{j2} + |\zeta_{i1} + \zeta_{i2} - \zeta_{j1} - \zeta_{j2}|\right] - A \\
&= 2k_e e_1^{\text{pub}} + \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2}\}\right] - A = 2k_e e_1^{\text{no}} + \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2}\}\right] - A = \Pi_{\text{best}}^{\text{no}},
\end{aligned}$$

where we made use of the well-known result that $\max\{a, b\} = (a + b + |a - b|)/2$.

(ii) Let $\kappa = 1$. Then, $\Pi_{\text{best}}^{\text{pri}} \approx 0.889a - A < 0.9a - A = \Pi_{\text{best}}^{\text{pub}}$, and by the continuity of $\Pi_{\text{best}}^{\text{pub}}$ and $\Pi_{\text{best}}^{\text{pri}}$, it follows that there exists a $\underline{\kappa} > 1$, such that $\Pi_{\text{best}}^{\text{pri}} < \Pi_{\text{best}}^{\text{pub}}$ for all $\kappa < \underline{\kappa}$.

(iii) The proof proceeds in two steps. First, we establish a lower bound for the firm's expected profits under a private feedback policy, $\underline{\Pi}_{\text{best}}^{\text{pri}} < \Pi_{\text{best}}^{\text{pri}}$, and show that $\underline{\Pi}_{\text{best}}^{\text{pri}} > \Pi_{\text{best}}^{\text{pub}}$ if γ_3 is sufficiently low. Last, we verify that there exists a $\overline{\kappa}$ such that γ_3 becomes sufficiently low for all $\kappa > \overline{\kappa}$.

Lower bound. The firm's expected profit is $\Pi_{\text{best}}^{\text{pri}} = k_e e_1^{\text{pri}} + \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2} + k_e e_2^{\text{pri}}(\zeta_{i1})\}\right] - A$. Clearly, for any effort function $\underline{e}_2(\zeta_{i1})$ with $\underline{e}_2(\zeta_{i1}) \leq e_2^{\text{pri}}(\zeta_{i1})$ for all ζ_{i1}, we have $\underline{\Pi}_{\text{best}}^{\text{pri}} = k_e e_1^{\text{pri}} + \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2} + k_e \underline{e}_2(\zeta_{i1})\}\right] - A \leq \Pi_{\text{best}}^{\text{pri}}$. In the remainder, we set $\underline{e}_2(\zeta_{i1}) \equiv -\frac{\zeta_{i1}}{k_e} + \frac{a\kappa}{k_e}\ln(\zeta_{i1} + a\kappa) - \frac{a\kappa}{k_e}\ln(\gamma_3)$. To see that this is indeed a lower bound on $e_2^{\text{pri}}(\zeta_{i1})$, note that $\underline{e}_2(\zeta_{i1})$ solves the integral equation (32) for all ζ_{i1}. By doing so, however, we ignore the fact that for some ζ_{i1} and ζ_{j1}, we have $\zeta_{i1} + k_e\underline{e}_2(\zeta_{i1}) - \zeta_{j1} + k_e\underline{e}_2(\zeta_{j1}) \notin [-a, a]$. This implies that the left-hand side of (32) is extended by negative terms compared to the correct solution outlined in Proposition 3. Now, since the left-hand side is smaller, it follows by equality that the right-hand side is smaller as well, thereby implying $\underline{e}_2(\zeta_{i1}) \leq e_2^{\text{pri}}(\zeta_{i1})$.

With $\underline{e}_2(\zeta_{i1})$, the firm's expected profit becomes $\underline{\Pi}_{\text{best}}^{\text{pri}} = a\cdot(-\kappa^3(\kappa^2 + \frac{1}{4})(e^{\frac{1}{\kappa}} - e^{-\frac{1}{\kappa}}) + \kappa^4(e^{\frac{1}{\kappa}} + e^{-\frac{1}{\kappa}}) + \kappa(\kappa^2 - 2\kappa + \frac{1}{4})\ln(\frac{2\kappa-1}{2\kappa+1}) + \frac{1}{2}\kappa\ln(a^4(2\kappa-1)(2\kappa+1)^3) + \frac{5}{6}\kappa^2 - 2\kappa(1 + \ln(2) + \ln(\gamma_3)) + \frac{1}{12})$, and $\underline{\Pi}_{\text{best}}^{\text{pri}} > \Pi_{\text{best}}^{\text{pub}}$ if and only if $\gamma_3 < \underline{\gamma}_3$, with

$$\underline{\gamma}_3 = \frac{a}{2}\cdot e^{-\frac{\kappa^2}{2}\left(\kappa^2+\frac{1}{4}\right)\left(e^{\frac{1}{\kappa}}-e^{-\frac{1}{\kappa}}\right)}\cdot e^{\frac{\kappa^3}{2}\left(e^{\frac{1}{\kappa}}+e^{-\frac{1}{\kappa}}\right)}\cdot\left(\frac{2\kappa-1}{2\kappa+1}\right)^{\frac{1}{2}\left(\kappa^2-2\kappa+\frac{1}{4}\right)}\cdot\sqrt[4]{(2\kappa-1)(2\kappa+1)^3}\cdot e^{\frac{5}{12}\kappa-1-\frac{3}{40\kappa}-\frac{1}{3\kappa^2}}. \tag{44}$$

Taking the limit. To test whether $\gamma_3 < \underline{\gamma}_3$, we will first derive an upper bound on γ_3, and then show that this upper bound is smaller than $\underline{\gamma}_3$. As a preliminary step, define the function

$$\Gamma(x,y) = \frac{2\gamma_3}{a(1+6\kappa)} = \frac{n^2x^2+y^2}{x^2(3y^2-n^2x^2+4n^3xy)}, \tag{45}$$

which decreases in x and y. Thus, since $x \in \left[e^{-1/(4\kappa)}, e^{-1/(4\kappa)\cdot(1-1/\kappa)}\right]$ and $y \in \left[e^{1/(4\kappa)}, e^{1/(4\kappa)\cdot(1+1/\kappa)}\right]$, it follows that $\Gamma(x,y) \in [\underline{\Gamma}, \overline{\Gamma}] = [e^{1/(2\kappa)}/((1+2e^{1/\kappa})e^{1/(2\kappa^2)}), e^{1/(2\kappa)}/(1+2e^{1/\kappa})]$. Given the monotonicity of $\Gamma(x,y)$, we can build the inverse function of $\Gamma(x,y)$ with respect to y:

$$y(x,\Gamma) = \frac{nx}{1-3\Gamma x^2} \cdot \left(2n^2\Gamma x^2 - \sqrt{4(1+n^4)\Gamma^2x^4 - (\Gamma x^2-1)^2}\right). \tag{46}$$

Inserting (46) in (6) and (7) allows us to eliminate y from the system of equations, and to represent it in variables x and Γ. Now, $\underline{\Pi}^{\text{pri}}_{\text{best}} \leq \Pi^{\text{pub}}_{\text{best}}$ if and only if the transformed system of equations has a solution for Γ in the interval $[\underline{\Gamma}_3, \overline{\Gamma}]$, and x arbitrary, where $\underline{\Gamma}_3 = 2\underline{\gamma}_3/(a(1+6\kappa))$. We proceed to show that for sufficiently large κ, such a solution does not exist.

Before doing so, we derive some important properties. Let $l_1(x,y)$ be the left-hand side of (6), and $l_2(x,y)$ be the left-hand side of (7). Straightforward differentiation verifies that there exists a $\overline{\kappa}$ such that for all $\kappa > \overline{\kappa}$, $l_1(x,y)$ increases in x and decreases in y, whereas $l_2(x,y)$ decreases in x and y. Furthermore, denote by $x_1(y)$ $(x_2(y))$ the solution to $l_1(x_1(y),y)=0$ $(l_2(x_2(y),y)=0)$ for any y. Applying the Implicit Function Theorem reveals that $x_1(y)$ increases in y, while $x_2(y)$ decreases in y for $\kappa > \overline{\kappa}$.

In a next step, we transfer these results to the transformed system of equations, which we denote by $l_1'(x,\Gamma)=0$ and $l_2'(x,\Gamma)=0$. Analogously to above, let $x_1'(\Gamma)$ $(x_2'(\Gamma))$ be the solution to $l_1'(x_1'(\Gamma),\Gamma)=0$ $(l_2'(x_2'(\Gamma),\Gamma)=0)$ for any Γ. Moreover, note that by the Inverse Function Theorem, $y(x,\Gamma)$ decreases in Γ, because $\Gamma(x,y)$ decreases in y. Therefore, by total differentiation, it follows that $\partial x_1'(\Gamma)/\partial\Gamma = \partial x_1(y)/\partial y \cdot \partial y(x,\Gamma)/\partial\Gamma < 0$, and $\partial x_2'(\Gamma)/\partial\Gamma = \partial x_2(y)/\partial y \cdot \partial y(x,\Gamma)/\partial\Gamma > 0$ for $\kappa > \overline{\kappa}$.

We are now well-equipped to complete the proof. We want to show that there exists a $\overline{\kappa}$ such that for all $\kappa > \overline{\kappa}$, the transformed system of equations $l_1'(x^*,\Gamma^*)=0$ and $l_2'(x^*,\Gamma^*)=0$ admits no solution with $\Gamma^* \in [\underline{\Gamma}_3, \overline{\Gamma}]$. We do so by verifying that for all $\kappa > \overline{\kappa}$, $l_2'(x,\Gamma) > 0$ for any x and $\Gamma \in [\underline{\Gamma}_3, \overline{\Gamma}]$. Note that for $\kappa > \overline{\kappa}$, $l_2'(x,\Gamma)$ decreases in x, and increases in Γ. This is true because $\partial l_2'(x,\Gamma)/\partial\Gamma = \partial l_2(x,y)/\partial y \cdot \partial y/\partial\Gamma > 0$, and, by the Implicit Function Theorem, $\partial l_2'(x,\Gamma)/\partial x = -(\partial l_2'(x,\Gamma)/\partial\Gamma)/(\partial x_2'(\Gamma)/\partial\Gamma) < 0$. Therefore, for $\kappa > \overline{\kappa}$, $l_2'(x,\Gamma) \geq l_2'(\overline{x}, \underline{\Gamma}_3)$, where $\overline{x} = e^{-1/(4\kappa)\cdot(1-1/\kappa)}$. It remains to demonstrate that $l_2'(\overline{x}, \underline{\Gamma}_3) > 0$, or equivalently, $l_2(\overline{x}, y(\overline{x}, \underline{\Gamma}_3)) > 0$ for $\kappa > \overline{\kappa}$. We will conclude this final step with the help of a two-step Taylor series expansion. As a starting point, we substitute $1/\kappa$ by z. This substitution allows us to develop the Taylor series at $\hat{z}=0$. Now, as a first step, the Taylor series of $y(\overline{x}, \underline{\Gamma}_3)$ at $\hat{z}=0$ is given by $y^{\text{Taylor}}(z) = 1 + z/4 - 3z^2/32 - 1177z^3/4480 - 611z^4/14336 + \mathcal{O}(z^5)$. In a second step, we can now derive the Taylor series of $l_2(\overline{x}, y(\overline{x}, \underline{\Gamma}_3)) = l_2(\overline{x}, y^{\text{Taylor}}(z))$ at $\hat{z}=0$. After resubstitution, this Taylor series becomes $l_2^{\text{Taylor}}(\overline{x}, y(\overline{x}, \underline{\Gamma}_3)) = 11/(3360\kappa^2) - 23/(960\kappa^3) + \mathcal{O}(1/\kappa^4)$. Obviously, since the first term is positive, we can conclude that there exists a $\overline{\kappa} < \infty$ such that $l_2^{\text{Taylor}}(\overline{x}, y(\overline{x}, \underline{\Gamma}_3)) > 0$ for all $\kappa > \overline{\kappa}$. $\square$

References

Ales, L., S.-H. Cho, E. Körpeoğlu. 2014. Innovation tournaments with multiple contributors. Working Paper.

Amann, E., W. Leininger. 1996. Asymmetric all-pay auctions with incomplete information: The two-player case. *Games and Economic Behavior* **14**(1) 1–18.

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

Beil, D. 2010. Procurement auctions. *INFORMS Tutorials in Operations Research*. INFORMS, 248–259.

Berger, J., D. Pope. 2011. Can losing lead to winning? *Management Science* **57**(5) 817–827.

Bockstedt, J., C. Druehl, A. Mishra. 2015. Problem-solving effort and success in innovation contests: The role of national wealth and national culture. *Journal of Operations Management* **36** 187–200.

Boudreau, K., N. Lacetera, K. Lakhani. 2011. Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science* **57**(5) 843–863.

Brookins, P., D. Ryvkin. 2014. An experimental study of bidding in contests of incomplete information. *Experimental Economics* **17**(2) 245–261.

Bull, C., A. Schotter, K. Weigelt. 1987. Tournaments and piece rates: An experimental study. *Journal of Political Economy* **95**(1) 1–33.

Casas-Arce, P., F. Asis Martinez-Jerez. 2009. Relative performance compensation, contests, and dynamic incentives. *Management Science* **55**(8) 1306–1320.

Che, Y.-K., I. Gale. 2003. Optimal design of research contests. *The American Economic Review* **93**(3) 646–671.

Chesbrough, H. W. 2003. *Open Innovation:* T*he New Imperative for Creating and Profiting from Technology*. Harvard Business School Press, Cambridge, MA.

Crawford, V. P., J. Sobel. 1982. Strategic information transmission. *Econometrica* **50**(6) 1431–1451.

Dechenaux, E., D. Kovenock, R. Sheremeta. 2014. A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics* 1–61Forthcoming.

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

Ederer, F., E. Fehr. 2009. Deception and incentives how dishonesty undermines effort provision. Working Paper.

Ederer, F., G. Manso. 2013. Is pay for performance detrimental to innovation? *Management Science* **59**(7) 1496–1513.

Engelbrecht-Wiggans, R., E. Katok. 2006. E-sourcing in procurement: Theory and behavior in reverse auctions with noncompetitive contracts. *Management Science* **52**(4) 581–596.

Erat, S., U. Gneezy. 2015. Incentives for creativity. *Experimental Economics* 1–12Forthcoming.

Erat, S., V. Krishnan. 2012. Managing delegated search over design spaces. *Management Science* **58**(3) 606–623.

Fonseca, M. 2009. An experimental investigation of asymmetric contests. *International Journal of Industrial Organization* **27**(5) 582–591.

Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.

Fullerton, R. L., R. P. McAfee. 1999. Auctioning entry into tournaments. *Journal of Political Economy* **107**(3) 573–605.

Gaba, A., I. Tsetlin, R. L. Winkler. 2004. Modifying variability and correlations in winner-take-all contests. *Operations Research* **52**(3) 384–395.

Garcia, S., A. Tor. 2009. The n-effect. *Psychological Science* **20**(7) 871–877.

Gavious, A., B. Moldovanu, A. Sela. 2002. Bid costs and endogenous bid caps. *The RAND Journal of Economics* **33**(4) 709–722.

Gershkov, A., M. Perry. 2009. Tournaments with midterm reviews. *Games and Economic Behavior* **66**(1) 162–190.

Glazer, A., R. Hassin. 1988. Optimal contests. *Economic Inquiry* **26**(1) 133–143.

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

Green, J. R., N. L. Stokey. 1983. A comparison of tournaments and contracts. *The Journal of Political Economy* **91**(3) 349–364.

Gupta, S., W. Chen, M. Dawande, G. Janakiraman. 2015. Optimal descending mechanisms for constrained procurement. *Production and Operations Management* Forthcoming.

Gürtler, O., C. Harbring. 2010. Feedback in tournaments under commitment problems: Experimental evidence. *Journal of Economics & Management Strategy* **19**(3) 771–810.

Gürtler, O., J. Münsure, P. Nieken. 2013. Information policy in tournaments with sabotage. *The Scandinavian Journal of Economics* **115**(3) 932–966.

Jeppesen, L., K. Lakhani. 2010. Marginality and problem-solving effectiveness in broadcast search. *Organization Science* **21**(5) 1016–1033.

Konrad, K. 2009. *Strategy and Dynamics in Contests*. Oxford University Press, Oxford.

Konrad, K. A., D. Kovenock. 2009. Multi-battle contests. *Games and Economic Behavior* **66**(1) 256–274.

Körpeoğlu, E., S.-H. Cho. 2015. Incentives in contests with heterogeneous solvers. Working Paper.

Kostamis, D., D. Beil, I. Duenyas. 2009. Total-cost procurement auctions: Impact of suppliers' cost adjustments on auction format choice. *Management Science* **55**(12) 1985–1999.

Krishna, V., J. Morgan. 1997. An analysis of the war of attrition and the all-pay auction. *Journal of Economic Theory* **72**(2) 343–362.

Kuhnen, C., A. Tymula. 2012. Feedback, self-esteem, and performance in organizations. *Management Science* **58**(1) 94–113.

Lazear, E. P., S. Rosen. 1981. Rank-order tournaments as optimum labor contracts. *The Journal of Political Economy* **89**(5) 841–864.

Lim, W., A. Matros, T. Turocy. 2014. Bounded rationality and group size in Tullock contests: Experimental evidence. *Journal of Economic Behavior & Organization* **99** 155–167.

Liu, T., J. Yang, L. Adamic, Y. Chen. 2014. Crowdsourcing with all-pay auctions: A field experiment on Taskcn. *Management Science* **60**(8) 2020–2037.

Ludwig, S., G. Lünser. 2012. Observing your competitor – The role of effort information in two-stage tournaments. *Journal of Economic Psychology* **33**(1) 166–182.

Manso, G. 2011. Motivating innovation. *Journal of Finance* **66**(5) 1823–1860.

Moldovanu, B., A. Sela. 2001. The optimal allocation of prizes in contests. *The American Economic Review* **91**(3) 542–558.

Moldovanu, B., A. Sela. 2006. Contest architecture. *Journal of Economic Theory* **126**(1) 70–96.

Nalebuff, B. J., J. E. Stiglitz. 1983. Prizes and incentives: Towards a general theory of compensation and competition. *The Bell Journal of Economics* **14**(1) 21–43.

Parreiras, S., A. Rubinchik. 2010. Contests with three or more heterogeneous agents. *Games and Economic Behavior* **68**(2) 703–715.

Polyanin, A. D., V. F. Zaitsev. 2003. *Handbook of Exact Solutions for Ordinary Differential Equations*. Chapman & Hall/CRC Press, Boca Raton.

Rosen, J. 1965. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica* **33**(3) 520–534.

Schotter, A., K. Weigelt. 1992. Asymmetric tournaments, equal opportunity laws, and affirmative action: Some experimental results. *The Quarterly Journal of Economics* **107**(2) 511–539.

Siegel, R. 2009. All-pay contests. *Econometrica* **77**(1) 71–92.

Siegel, R. 2010. Asymmetric contests with conditional investments. *The American Economic Review* **100**(5) 2230–2260.

Siegel, R. 2014a. Asymmetric all-pay auctions with interdependent valuations. *Journal of Economic Theory* **153**(0) 684–702.

Siegel, R. 2014b. Asymmetric contests with head starts and nonmonotonic costs. *American Economic Journal: Microeconomics* **6**(3) 59–105.

Taylor, C. R. 1995. Digging for golden carrots: An analysis of research tournaments. *The American Economic Review* **85**(4) 872–890.

Terwiesch, C., K. T. Ulrich. 2009. *Innovation Tournaments*. Harvard Business School Press.

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

Tsetlin, I., A. Gaba, R. L. Winkler. 2004. Strategic choice of variability in multiround contests and contests with handicaps. *Journal of Risk and Uncertainty* **29**(2) 143–158.

Tunca, T., J. Wu, F. Zhong. 2014. An empirical analysis of price, quality, and incumbency in procurement auctions. *Manufacturing & Service Operations Management* **16**(3) 346–364.

von Hippel, E. 2005. *Democratizing Innovation*. MIT Press, Cambridge, MA.

Wan, Z., D. Beil, E. Katok. 2012. When does it pay to delay supplier qualification? Theory and experiments. *Management Science* **58**(11) 2057–2075.

Wooten, J., K. Ulrich. 2014. Idea generation and the role of feedback: Evidence from field experiments with innovation tournaments. Working Paper.

Yildirim, H. 2005. Contests with multiple rounds. *Games and Economic Behavior* **51**(1) 213–227.