Innovative Products and their Impact on Operations

Habilitationsschrift

vorgelegt
an der Fakultät für Betriebswirtschaftslehre
der Universität Mannheim
von

Dr. Jochen Schlapp
Wilhelm-Leuschner-Str. 241
64347 Griesheim

Mannheim, Februar 2018

Table of Contents

List of Figures

List of Tables

[This page intentionally left blank.]

Chapter I

Introduction

Operations is the part of an organization that is responsible for creating and/or delivering the organization's products and services (Slack and Lewis 2015, p. 2). To accomplish this task, operations has to oversee and improve a large variety of different processes that span the entire value chain; including the design, procurement, production, delivery, and recovery of the products and services sold. It is beyond dispute that managing all these processes is a complex and challenging endeavor, and that any form of mismanagement may result in adverse outcomes that have a detrimental impact on the organization's profits. Yet, even though the coordination of all the different functions that operations is concerned with is by itself already an intricate mission, there is one feature of many real-life scenarios that substantially complicates operations management: the presence of uncertainties.

When uncertainty—be it on the supply, demand, or product side—enters the equation, then organizations cannot simply plan for how to proceed with their operations in the future, but instead they have to foresee all the different possible futures and create contingency plans for each of them (Van Mieghem and Allon 2015, p. 115). Clearly, this is an utmost daring—if not impossible—task and therefore, operations management has traditionally been concerned with simplifying this task by eliminating the root causes of the issue; that is, operations management also focusses on reducing uncertainties. This view that operations should reduce, or at least control, the organization's supply, demand and product uncertainties is well established, and accordingly, managers in practice routinely deal with operational risk management (Van Mieghem and Allon 2015, p. 369).

However, combatting uncertainties is not always an organization's most prudent strategy; in particular when the organization deals with innovative products. As the famous US physicist Brain Randolph Greene once noted: "*Exploring the unknown requires tolerating uncertainty*" (Greene 2006). In other words, uncertainty and risk are the mainsprings of innovation, and when engaging with innovative products, operations management—although going against its very nature—has to embrace uncertainty. Even more so since access to innovation is a key source of competitive advantage in many industries. As a result, operations has to find entirely different ways to manage innovative products as compared to standard products, where risk reduction is a primary concern. And uncovering these paths is at the heart of this thesis. More specifically, the following chapters investigate how operations should manage the procurement (Chapters II and III) and design (Chapter IV) of innovative products.

Chapters II and III are concerned with the design and management of so-called innovation or procurement contests, which have established themselves as a predominant purchasing mechanism for innovative products in practice (Cabral et al. 2006). The key benefit of using procurement contests as compared to bilateral contracting for research and development (R&D) is that a contest is an informationally very parsimonious mechanism that also spurs substantial innovation efforts in a firm's supplier base and allows for an ex post selection of the best innovation (Terwiesch and Xu 2008). Yet, to reap all these benefits contest holders need to correctly setup their contests and also manage the contest wisely as it unfolds.

How to manage an ongoing innovation contest is the key focus of Chapter II. In particular, Chapter II analyzes how feedback—the most practical form of in-contest interventions—from the contest holder to the contest participants can help in improving contest outcomes, and it also establishes optimal information structures for a contest holder's feedback policy. Specifically, the analysis identifies when, and when not, to give feedback as well as which type of feedback to give: public (which all participants can observe) or private (which only the focal participant can observe). The results uncover a nontrivial relationship between contest characteristics and optimal feedback choices. Additionally, Chapter II also examines whether the contest holder should mandate interim feedback or instead allow participants to seek feedback at their own discretion, and discusses how changing the granularity of feedback information affects its value to participants.

In contrast to Chapter II, Chapter III concentrates on a contest holder's design choices before the start of the contest. For rather simple innovations the question of successful contest design has received considerable attention in the academic literature (see, e.g., Taylor 1995, Moldovanu and Sela 2001, Terwiesch and Xu 2008), and as a result, scholars have gathered a sound understanding of how to conceptualize such contests. Unfortunately, many of these findings do not immediately transfer to contests that aim to source technologically complex products that consist of multiple interacting components. For such complex innovations, a central question for the buying firm is whether to procure the full product from a single supplier, or whether to buy the individual components from different suppliers. The analysis presented in Chapter III shows that the answer to this question depends on the magnitude of innovation that is required to develop the different components as well as the characteristics of the supplier base. Based on these findings, Chapter III provides managerial advice regarding the optimal contest format, and it also highlights which suppliers should be invited to participate in such procurement contests, and which not.

Chapter IV departs from the preceding chapters by considering how a firm should manage its internal new product development efforts, instead of procuring innovations only from external parties. In doing so, the analysis presented in this chapter examines how a firm should operationalize its design testing efforts. It is well known that design testing is an integral part of nearly all new product development initiatives because it enables firms to identify the best designs for their new products. Test results are usually collected by (teams of) experts, who must be incentivized not only to exert effort in testing the designs but also to report their findings truthfully. Motivated by this widespread challenge, Chapter IV addresses the following questions: How should a firm set up its design-testing process so that (i) the experts are adequately incentivized and, more importantly, (ii) the best design alternative is the one most likely to be selected? The presented analysis identifies the firm's optimal testing strategy and the optimal incentive structures; it also reveals how, exactly, delegation distorts a firm's testing process.

Taken together the findings presented in this thesis should be viewed as being two sides of the same coin. To retain their competitive edge and to improve their market position, firms have to constantly tap into new and innovative products. Such innovations can come either from outside parties—in which case procurement contests are an effective way of gaining access to these innovations—or they may be developed internally. Both ways of sourcing innovation are frequently used in reality, but unfortunately, both of them also introduce very challenging problems to a firm's operations management. It is the goal of this thesis to provide managerial solutions for some of the most pressing issues that managers face when trying to gain access to innovation—be it from external sources or from an internal R&D department.

[This page intentionally left blank.]

Chapter II

Feedback in Innovation Contests

with Jürgen Mihm[1]

2.1 Introduction

Firms have increasingly found it necessary to source their innovation from beyond their own boundaries (Chesbrough 2003). They often do not have, in house, the expertise needed to solve all the challenges that arise as a result of their ever more complex research and development (R&D) activities. Yet success often eludes innovation initiatives that involve outside parties; much depends on the suitability of the firm's sourcing mechanism. One mechanism that has garnered widespread interest is the *innovation contest*. In organizing such a contest, the firm specifies its goals at the outset (and often the metric by which it measures goal achievement) and promises an award to the solver who best fulfills those goals; at the end of the contest, the award is granted to the solver(s) with the best solution(s). The contest mechanism offers two key benefits: (i) it offers considerable flexibility in that the firm can choose a different set of participants for each contest; and (ii) it equips the firm with powerful incentives, since contestants compete fiercely to win the contest holder's approval and thus the award.

In light of these potential benefits, contests have been widely studied in the context of innovation and also in many other settings (Lazear and Rosen 1981, Moldovanu and Sela 2001, Siegel 2009, Ales et al. 2016). One consequence of this research interest is that a theory of contests has emerged. This theory focuses on how a contest holder can use different aspects of contest design to optimize the intensity of competition among contestants, thereby maximizing the effort exerted by contestants and, by extension,

[1]The research presented in this chapter is based on the paper "Sourcing Innovation: On Feedback in Contests", coauthored with Jürgen Mihm, which has been accepted for publication in *Management Science*.

the contest's effectiveness at providing incentives. The theory of contests has offered solutions for such diverse problems as optimal award structures (Ales et al. 2017), the optimal number of participants (Taylor 1995, Terwiesch and Xu 2008), and the optimal way of nesting contests within contests (Moldovanu and Sela 2006).

However, current theory has some gaps with respect to certain critical aspects. We highlight these gaps by considering Kaggle, an Internet platform that provides firms with the infrastructure for hosting contests on data-driven problems. When setting up such a contest, the firm must define its rules of engagement: the relevant metric (usually, out-of-sample accuracy of the predictions) and the reward(s) offered. After the contest announcement, data scientists compete against each other in 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. So in those respects, Kaggle's approach follows the general template of a contest. In one respect, however, it adds a fundamentally new feature: During the competition, data scientists enter preliminary versions of their code and receive feedback on how well it performs (usually in terms of how accurate its predictions are). Furthermore, Kaggle not only provides this performance feedback to the team itself but also maintains a public "leader board" so that each team (or individual participant) can observe its own performance relative to all competing submissions.

Thus Kaggle can be viewed as exemplifying a central question, faced by many contest organizers in practice, that has received but cursory attention in the academic literature: the question of optimal feedback (for notable exceptions, see Aoyagi 2010, Ederer 2010, Marinovic 2015). Performance feedback is a means by which the firm can systematically affect the amount of information held by each contestant—in particular, information about own and rivals' competitiveness—and thereby influence contestant behavior during the rest of the contest. Put differently, the firm can augment or diminish incentives by redistributing information and in this way can manipulate the contest's competitiveness. The question that then arises is: How, exactly, should a contest holder influence the information structure during a contest so that contestants are optimally incentivized?

Any comprehensive investigation of this issue must provide answers to the following three questions, which together constitute a feedback policy's *information structure*. (i) Which solvers should receive the feedback information? (ii) Who should decide which solvers receive feedback? (iii) What should be the information content of the performance feedback?

The importance of the first question rests on the fact that a contest holder can freely choose the recipients of feedback. More specifically, the firm may retain all information about the contest's competitiveness (no feedback), it may inform solvers about their respective individual performance but not about the performance of others (private feedback), or it may provide information about the performance of all contestants (public feedback). These information structures naturally induce different levels of competition and hence provide contestants with different incentives. However, it is not clear which policy is most appropriate for which situation. Real-world contest holders have experimented extensively with different forms. The default mode for Kaggle is to allow all contestants to observe each contestant's performance feedback. In contrast, the European Union (EU)—which regulates all major infrastructure, architectural design, and civil engineering contests organized within its jurisdiction—introduced in 2004 the "competitive dialogue procedure" (EU 2004) for the specific purpose of establishing a *private* feedback channel between contest holder and contestants. In 2010, 9% of the EU's entire public procurement budget was spent via this contest mechanism. The use of private feedback has proven so effective that, in 2016, the World Bank introduced a similar mechanism in its procurement regulations (World Bank 2016).[2]

With regard to the second question, it can be either the firm or a contestant who initiates feedback and hence a redistribution of information. In particular, the contest holder might mandate feedback or might simply provide a feedback option. In the latter case, contestants may strategically withhold their performance information to influence the contest's information structure. Should the contest holder allow for such strategic behavior? Again, companies have devised different approaches. Kaggle, for instance, often (though not always) makes feedback voluntary, whereas performance feedback is mandatory in any contest subject to the EU's competitive dialogue procedure.

The third question focuses on the accuracy of performance feedback. Clearly, any information about contestants' relative competitiveness will affect their incentives and thus their behavior. But should the firm divulge all of its available information or only some of it? Kaggle issues exact rankings of contestants' performance (i.e., their re-

[2]Of course, for public institutions such as the EU or the World Bank, the choice of feedback policy will likely depend also on transparency and compliance rules and thus involve more than pure efficiency considerations.

spective prediction accuracy). Yet the annual European Social Innovation Competition, which solicits ideas for building an inclusive economy, tends to provide less fine-grained (and thus merely "indicative") feedback.

In practice, the informational impact of different feedback policies is key to designing a successful contest; hence it is imperative for the contest-staging firm to answer each of those three questions. Yet the existing academic literature leaves them largely unanswered by implicitly restricting attention to the role of feedback that is *public, mandatory,* and *fully informative* (Aoyagi 2010, Ederer 2010) or at best to a specific form of public, mandatory, and noisy feedback (Marinovic 2015). Thus that literature covers too few of feedback's dimensions and options within dimensions to have much relevance for most practical settings. Furthermore, it concentrates exclusively on firms interested in promoting the average performance of their solvers. In innovation settings, however, firms are more likely to be interested in the best performance. We contribute to the literature on contests by offering a more complete and practically relevant description of how feedback can be used in contests—whether to improve average performance or to obtain the best performance. In so doing, we consolidate the most relevant feedback policies observed in practice within a broad framework and thereby deepen our theoretical understanding of when and how to use them.

The answers we find to our three guiding questions are as follows. First, and most importantly, contest organizers (and researchers) cannot neglect private feedback. Whereas public feedback always dominates in average-performance settings, a contingency arises for contests that seek to elicit the best performance: private (resp. public) feedback is optimal for contests with high (resp. low) uncertainty. This finding is in stark contrast to the existing literature's view, based on comparing only the cases of public and no feedback, that the feedback's role is the same for routine projects as for highly innovative projects. Second, public feedback may be underused when it is voluntary. Contestants always seek performance feedback under a private-feedback policy but never do so under a public-feedback policy, and inducing contestants via monetary incentives to ask for public feedback yields suboptimal results. Third, concerning the effect of information granularity on the value of feedback, we find no evidence that strategically hiding information—either by reducing the information content of feedback (e.g., providing rank-only feedback rather than detailed performance feedback) or by promulgating noise—can be used to improve contest outcomes.

2.2 Related Literature

The question of how best to motivate innovation and creativity is a central topic of academic inquiry (see, e.g., Erat and Krishnan 2012, Ederer and Manso 2013, Bockstedt et al. 2015, Erat and Gneezy 2016). Contests as a mechanism for eliciting innovation opportunities have become a focal point of attention, figuring prominently in both the economics and the operations management literatures. In the classification of Taylor (1995), this broad literature examines two different types of contests: (i) *innovation races*, in which contestants try to achieve a pre-defined and verifiable performance target (see e.g., Bimpikis et al. 2016, Halac et al. 2016); and (ii) *innovation contests* for solving open-ended problems, in which the firm cannot specify performance targets ex ante and rather tries to induce the best solution.

Our work falls into the second category because the assumption of a pre-defined performance target would be antithetical to our main goal: exploring how feedback can incentivize contestants to achieve *optimal output* on a given schedule. The literature on contests (in the narrow sense) was initiated by seminal research of Lazear and Rosen (1981), Green and Stokey (1983), and Nalebuff and Stiglitz (1983). Over the last decades, these contests have become an accepted paradigm in the study of settings that include lobbying, litigation, military conflict, sports, education, and of course R&D management (for an overview of applications, see Konrad 2009). The extant literature has addressed many contest design issues. Prominent among these is whether or not the contest should be open for everybody to enter; a larger number of entrants yields a larger number of trials (Terwiesch and Xu 2008), but restricting access increases the effort exerted by individual solvers (Taylor 1995, Fullerton and McAfee 1999, Moldovanu and Sela 2001).[3] Bid caps have been studied as a means of limiting access to a contest (Gavious et al. 2002), and so have more advanced mechanisms such as holding an auction for the right to participate (Fullerton and McAfee 1999). Another prominent issue is the optimal award structure (Che and Gale 2003, Siegel 2009, 2010), which depends on such contingencies as the solvers' respective cost functions (Moldovanu and Sela 2001), performance uncertainty (Ales et al. 2017), and whether the firm seeks the best solution or only to improve the average solution (Moldovanu and Sela 2006). Another major issue is the contest's temporal structure. Should the contest designer hold a single,

[3]This generalization is countered by Ales et al. (2016) and Körpeoğlu and Cho (2017), who give examples of contests for which individual solution efforts are *increasing* in the number of competitors.

overarching contest or rather a series of smaller, "cascading" contests?—see Moldovanu and Sela (2006) for a discussion. Finally, the literature has also analyzed more dynamic contest formats such as elimination and round-robin contests (Yücesan 2013) as well as repeated contests between the same contestants (Konrad and Kovenock 2009).

All of these models presume that the contest holder is relatively passive during the course of the contest. However, recently, scholarly attention has been shifting toward the actions that a contest holder could take as the contest unfolds (see, e.g., Gürtler et al. 2013). The most prominent of these actions is providing (or not) interim performance feedback.

The literature on feedback in contests is sparse.[4] Although generally acknowledged to be the first in this area, the paper by Yildirim (2005) does not address feedback per se and focuses instead 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; rather, these authors focus on optimally aggregating scores by combining intermediate and final reviews when the review process itself is noisy. However, there are four papers that do address feedback during contests in a more narrow sense. Goltsman and Mukherjee (2011) explore a setting in which solvers compete for a single prize by fulfilling two tasks at which each solver can either fail or succeed. Closer to our work, Aoyagi (2010), Ederer (2010), and Marinovic (2015) 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 has yielded only preliminary answers to some aspects of the three foundational questions that shape any feedback policy's information structure. First, all extant research restricts its attention to *public* feedback and neglects the class of *private* feedback (which is ubiquitous in practice); hence broader comparisons of different feedback policies have not been made. We solve the challenging case of private feedback and find nontrivial contingencies accounting for when private, public, or no feedback is preferable. Our results confirm the importance of private feedback for highly innovative settings and hence challenge extant research. As an aside, our analysis of private feedback contributes to the mathematical theory of contests by devising—to the best of our knowledge—the first closed-form solution of a stochastic contest with asymmetric private information. Second, previous research has considered only *mandatory* feedback. In other words, it implicitly assumes

[4]In the following we concentrate on theoretical work, but it is worth mentioning also the stream of empirical studies (see e.g., Gross 2017, Wooten and Ulrich 2017).

that solvers must provide the contest holder with intermediate solutions on which they receive feedback—an assumption often violated in practice. We examine all types of feedback with respect to mandatory versus voluntary feedback and establish the circumstances under which a firm should (or should not) make feedback mandatory. We also investigate the role that intermediate prizes designed to induce voluntary feedback play in this regard. Third, the existing literature simply presumes that the contest holder divulges all available information to contestants; the only exception is Marinovic (2015), who considers a specific form of noisy feedback. Yet feedback may in fact convey less fine-grained information, so we explore the effects of reducing the amount of feedback information conveyed. Finally, the contest literature on feedback has attended solely to the *average* performance of solvers. We answer each of the three central questions not only for a contest holder aiming to improve average performance but also for one looking for the best possible performance—a goal more typical of innovation settings. We show that the optimality of feedback policies hinges on this distinction.

2.3 Model Setup

Let us describe in more detail the characteristics of a typical innovation contest in terms of both the firm and the solvers so as to establish our base model (voluntary feedback and reduced information feedback are treated in Sections 2.6 and 2.7, respectively). The firm understands its own preference structure well enough that, when presented with a solution, the firm can express how much it values that solution. However, the firm cannot know the effort level expended by a solver 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 realizes that expected performance increases with effort. Yet solvers still experience uncertainty about how, exactly, effort links to performance. In addition, solvers are uncertain about the firm's preference structure and so, even after devising a solution, they cannot truly evaluate their performance. This latter uncertainty reflects that, for any true innovation, the firm cannot *fully* specify ex ante what criteria it values or how they should be weighted. These modeling requirements are typical for any innovation and R&D setting, and they place the foundation of our model squarely in the contest literature with stochastic ef-

Figure 2.1.: Structure of the Innovation Contest with Feedback.

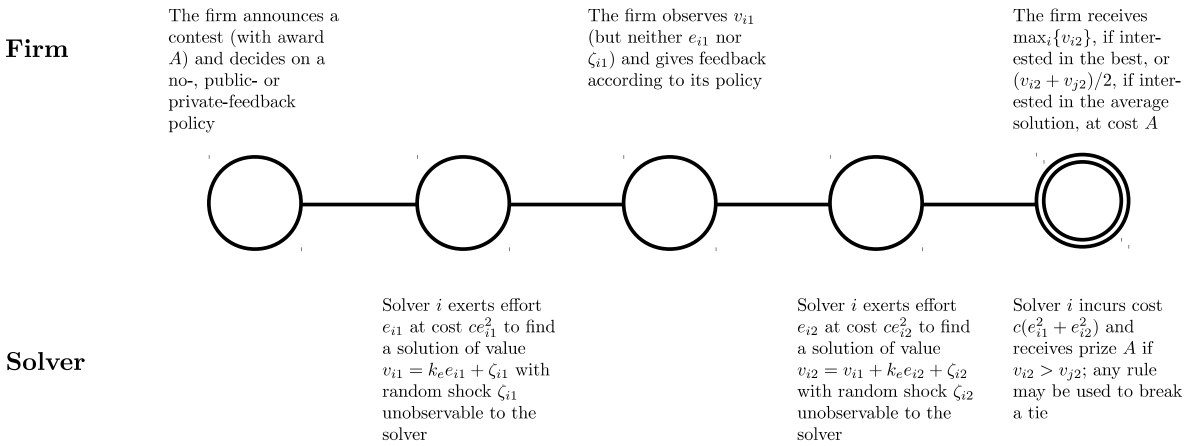

fort consequences which postulates a stochastic link between solvers' actions and contest outcomes (Taylor 1995, Fullerton and McAfee 1999, Ales et al. 2016).

Finally, 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*. We classify such feedback as public, private, or no feedback. The firm employs whichever feedback policy optimizes the contest's intended outcome—the highest average performance or best possible performance.

Formal Description of the Base Model.

In order to create a parsimonious model that nonetheless captures the essence of the scenario just outlined, we consider a firm that hosts a dynamic innovation contest over two rounds, $t \in \{1, 2\}$, with two risk-neutral solvers, i and j.[5] The primitives of the contest are common knowledge; its structure is depicted in Figure 2.1.

The process begins when 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 minimize technical complexity), we treat $A > 0$ as a parameter. Our decision variable for the firm at this stage is whether and, if so, how to give feedback. The firm may choose to give no feedback at all, to offer public feedback (i.e., both solvers

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

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 $e_{i1} \geq 0$[6] at private cost ce_{i1}^2, where $c > 0$. He finds an initial solution of value $v_{i1} = k_e e_{i1} + \zeta_{i1}$; here $k_e > 0$ is the sensitivity of effort and ζ_{i1} is a random shock that follows a uniform distribution, $\zeta_{i1} \sim \text{Uniform}(-a/2, a/2)$ with $a > 0$.

After the first round, each solver hands in his solution and the firm perfectly observes v_{i1}. However, solver i's effort is unobservable to the firm (and also to solver j); hence the firm cannot determine whether a high solution value stems from high effort, a large random shock, or both. In contrast, solver i knows how much effort he has invested; but since he cannot observe the realization of ζ_{i1}, he is uncertain about the true performance of his solution. To address that uncertainty, the firm provides interim performance feedback in accordance with its own policies. As is customary in the fledgling research field of feedback in contests, we assume that feedback is pre-committed, truthful and accurate (Aoyagi 2010, Ederer 2010)—although the "accurate feedback" assumption is relaxed in Section 2.7. It is clear that, in the absence of pre-committed truthfulness (i.e., if feedback does not convey a somewhat informative signal in a Bayesian sense), feedback is utterly meaningless. It is easy to prove that the firm would have a strong incentive to provide only feedback that maximizes future efforts *irrespective* of actual performance; naturally, each solver would anticipate this manipulation and discard the received information as uninformative.

Upon observing the firm's feedback, solver i updates his belief about the realization of first-round performances $v_1 = (v_{i1}, v_{j1})$ in accordance with 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 between the random shocks in round t with associated probability density function $g_{\Delta\zeta_t}$.

[6]An effort of 0 should be interpreted as the normalized minimal effort necessary to participate in the contest.

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

Model Implications.

A firm will naturally seek to employ the feedback policy that maximizes its expected profits. The relevant profit function is $\Pi_{\text{best}} = \mathbb{E}[\max\{v_{i2}, v_{j2}\}] - A$ if the firm is interested in the performance of the *best* solution only, or $\Pi_{\text{avg}} = \mathbb{E}[v_{i2} + v_{j2}]/2 - A$ 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's sole interest is in winning the contest. The utility that solver i receives from winning 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 $u_i = A \cdot \mathbb{P}(v_{i2} > v_{j2}) - \sum_t ce_{it}^2$ (we assume his outside option to be 0), and the effort he invests in the contest is determined by maximizing his expected utility.[7]

We are concerned with Perfect Bayesian Equilibria (PBE) of the contest. To avoid unnecessary technical complications during the analysis, we assume that $\kappa \equiv (a^2c)/(Ak_e^2) > 1$. For technical reasons, similar assumptions on the contest's inherent performance uncertainty are made in practically 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 higher κ, improvement effort is more expensive and the solution performance becomes more stochastic.

2.4 Solvers' Solution Efforts

In this section we analyze, for our base model, 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 whether the firm's

[7]Note that an outside option of zero ensures that a solver always participates in the contest because zero effort already guarantees him a nonnegative expected utility and in equilibrium his utility cannot be worse.

aim is to improve average performance or rather to attain the best performance. Each solver simply tries to win the contest.

We start by re-establishing familiar results in the context of our model, characterizing solvers' equilibrium efforts in the absence of feedback as a benchmark (Section 2.4.1); next we describe how providing public feedback affects the solution efforts of solvers (Section 2.4.2). In this section's main contribution, we then determine equilibrium levels of solution effort under a private-feedback policy (Section 2.4.3). Throughout the text, initial managerial implications are discussed in passing; however, our systematic comparison of feedback policies is deferred until Section 2.5. All mathematical derivations are presented in Appendix A.

2.4.1. No Feedback

In the benchmark case of no feedback, the firm does not provide any interim performance information to the solvers. As a result, each solver's two-stage effort choice problem reduces to a simultaneous, single-stage utility maximization problem.

Proposition 2.1. *The unique PBE under a no-feedback policy is symmetric, with*

$$e_1^{\text{no}} = e_2^{\text{no}} = \frac{Ak_e}{3ac}. \tag{2.1}$$

Proposition 2.1 parallels previous results of Taylor (1995), Fullerton and McAfee (1999), and Ales et al. (2016). Since neither solver receives any interim performance information and since the costs of effort are convex, it follows that solution efforts are identical across rounds. Moreover, because solvers are symmetric at the start of the contest, they always choose the same effort in equilibrium; hence they do not try to leapfrog each other. So under a no-feedback policy, it is the contest's inherent performance uncertainty that ultimately determines the contest winner.

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

2.4.2. Public Feedback

Next we study the implications of public feedback. In this case, after submitting his initial solution, each solver learns his own as well as his competitor's first-round performance. That is, public feedback perfectly reveals the solvers' first-round performance difference before the start of the second round, at which point solvers are therefore no longer symmetric.

Proposition 2.2. *The unique PBE under a public-feedback policy is symmetric, with*

$$e_1^{\text{pub}} = \mathbb{E}_{\Delta\zeta_1}[e_2^{\text{pub}}(\Delta\zeta_1)] = \frac{Ak_e}{3ac}, \tag{2.2}$$

$$e_2^{\text{pub}}(\Delta\zeta_1) = \frac{Ak_e}{2a^2c}\,(a - |\Delta\zeta_1|). \tag{2.3}$$

Mirroring Aoyagi (2010) and Ederer (2010), Proposition 2.2 has two main implications. First, it shows that each solver cares only about his relative performance and completely disregards the absolute performance information embedded in public feedback. Specifically: if the solvers' first-round performance difference is small (i.e., the contest is competitive), then second-round efforts are substantial and the solvers fight hard to win the contest; but if the first-round performance difference is sizable, then solvers reduce their solution efforts because the contest is de facto decided. Second, despite being asymmetric in the second round, both solvers expend the same amount of effort. In other words, the first-round leader pursues a simple blocking strategy: he tries to keep the follower at a 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. The follower just relies on a large positive second-round shock to reverse his fortune.

2.4.3. Private Feedback

We have just shown that, under a public-feedback policy, solvers set their second-round solution efforts as a function of their *relative* first-round performance. Yet that solver strategy is not viable under private feedback, since each solver receives information only about his own performance. Thus, only *absolute* performance information 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 which allows for the solvers' assessments of their chances to win to not be "coherent". 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 50%, although in reality those chances are merely 50%. And in contrast with the public-feedback scenario, solvers are never entirely certain whether they are ahead or behind their competitor. It is this asymmetric belief structure that drives asymmetric equilibrium solution efforts.

Proposition 2.3. *The unique PBE under a private-feedback policy is symmetric, with*

$$e_1^{\text{pri}} = \mathbb{E}_{\zeta_{i1}}[e_2^{\text{pri}}(\zeta_{i1})], \tag{2.4}$$

$$e_2^{\text{pri}}(\zeta_{i1}) = \begin{cases} -\frac{a+\zeta_{i1}}{k_e} + \frac{2a\kappa}{k_e}\ln\left(2\sqrt{\frac{\gamma_2}{3\gamma_1}}\sin\left(\frac{1}{3}\sin^{-1}\left(\sqrt{3\frac{\gamma_1}{\gamma_2}}\frac{a(1+12\kappa)+6\zeta_{i1}}{4\gamma_2}\right)\right)\right) & \text{if } \zeta_{i1} \in [-\frac{a}{2}, \gamma_3 x^2 - a\kappa], \\ -\frac{\zeta_{i1}}{k_e} + \frac{a\kappa}{k_e}\ln\left(\frac{\zeta + a\kappa}{\gamma_3}\right) & \text{if } \zeta_{i1} \in [\gamma_3 x^2 - a\kappa, \gamma_3 y^2 - a\kappa], \\ -\frac{\zeta_{i1}}{k_e} + \frac{2a\kappa}{k_e}\ln\left(\frac{\sqrt[3]{z(\zeta_{i1})}}{6} - 2\frac{\gamma_2}{\gamma 1}\frac{1}{\sqrt[3]{z(\zeta_{i1})}}\right) & \text{if } \zeta_{i1} \in [\gamma_3 y^2 - a\kappa, \frac{a}{2}]. \end{cases} \tag{2.5}$$

Here $z(\zeta_{i1}) = 12[-9(a(1/6 - 2\kappa) - \zeta_{i1}) + (12\gamma_2^3/\gamma_1 + 81(a(1/6 - 2\kappa) - \zeta_{i1})^2)^{1/2}]/\gamma_1$. *The constants are defined as* $\gamma_1 = p(ny - x)/(3nx^2 o)$, $\gamma_2 = py(n^3 x + y)/(nxo)$, *and* $\gamma_3 = p(n^2x^2 + y^2)/(2x^2 o)$, *where* $m = (1-6\kappa)/(1+6\kappa)$, $n = e^{1/(2\kappa)}$, $o = 3y^2 - n^2x^2 + 4n^3xy$, *and* $p = a(1+6\kappa)$ *and where* $x \in [e^{-1/(4\kappa)}, e^{-(1-1/\kappa)/(4\kappa)}]$ *and* $y \in [e^{1/(4\kappa)}, e^{(1+1/\kappa)/(4\kappa)}]$ *are the unique solutions to the following system of equations:*

$$mn^2x^4 - 4mn^3x^3y - 3(m+n^2)x^2y^2 - 4n^{-1}xy^3 + y^4 = 0, \tag{2.6}$$

$$\frac{1-6\kappa^2}{\kappa(1+6\kappa)} + m\ln(y) - \ln(x) + \frac{n^2x^4 + 8n^3x^3y + 9(1+n^2)x^2y^2 + 8n^{-1}xy^3 + y^4}{6x^2(3y^2 - n^2x^2 + 4n^3xy)} = 0. \tag{2.7}$$

This proposition presents—to the best of our knowledge—the first solution of a contest with asymmetric private information but it is rather unwieldy; we offer a more tractable approximation in Corollary 2.1. Our numerical analyses indicate that the

Figure 2.2.: Equilibrium Second-Round Effort under Private Feedback.

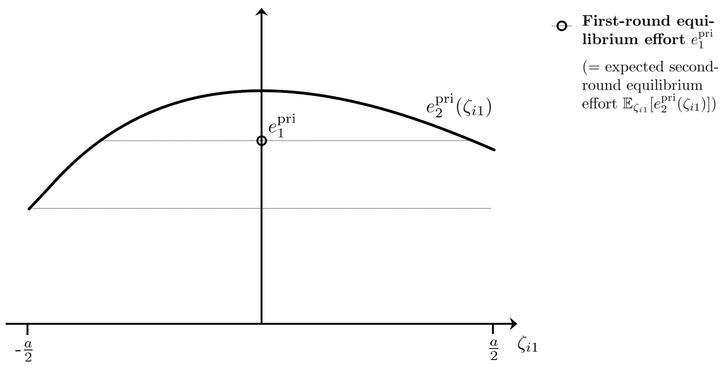

Notes: The functions are based on the following set of parameters: $A = 1$, $a = 1$, $k_e = 1$, $c = 1.01$.

corollary yields an exceptionally good approximation even for low κ, which makes it a good starting point for reflecting on Proposition 2.3.

Corollary 2.1. *Define* $\widetilde{\gamma}_3 = a(1+6\kappa)e^{(\kappa-1)/(2\kappa^2)}/(2(1+2e^{1/\kappa}))$*, and let*

$$\widetilde{e}_2(\zeta_{i1}) = -\frac{\zeta_{i1}}{k_e} + \frac{a\kappa}{k_e}\ln(\zeta_{i1}+a\kappa) - \frac{a\kappa}{k_e}\ln(\widetilde{\gamma}_3). \tag{2.8}$$

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

Figure 2.2 plots the equilibrium effort functions e_1^{pri} and $e_2^{\text{pri}}(\zeta_{i1})$ for different first-round shocks. The graph makes salient that Proposition 2.3 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 second round contribute equally to a solver's overall performance. Second, a solver's second-round effort $e_2^{\text{pri}}(\zeta_{i1})$ is not monotonically increasing in ζ_{i1}. In fact, $e_2^{\text{pri}}(\zeta_{i1})$ has an inverted U-shape; it increases with ζ_{i1} for $\zeta_{i1} \leq 0$ but decreases with ζ_{i1} for $\zeta_{i1} > 0$. Thus solvers with a moderate first-round performance (i.e., $\zeta_{i1} = 0$) exert substantial efforts 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 competitive whereas exceptionally good- or ill-performing solvers perceive the contest as more or less decided. Most importantly, however, unlike the public-feedback scenario, under private feedback the bad solvers reduce their efforts to a greater extent than do the good solvers; formally, $e_2^{\text{pri}}(-\zeta_{i1}) < e_2^{\text{pri}}(\zeta_{i1})$ for all $\zeta_{i1} > 0$ (observe the asymmetry in Figure 2.2). This finding stems from the absence of relative performance information. A solver with a high first-round shock can never be certain that he is ahead, so he invests more effort to maintain his chances of winning in case the competitor is equally strong—even though that is unlikely. Hence, private feedback induces well-performing solvers to invest relatively more effort; it makes them relatively more risk averse. This asymmetric response to feedback is the central feature that distinguishes private from public feedback.

But does this mean that less fortunate solvers can leapfrog better solvers by increasing their second-round efforts? The answer is No. To see this, note that solver i's final performance v_{i2}^{pri} is increasing in ζ_{i1}. That is: 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 strategic behavior of solvers. In equilibrium, no solver ever allows a less fortunate solver (i.e., one with a lower first-round shock) to overtake him in the second round through effort alone. So 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.

2.5 The Optimal Feedback Policy

Having characterized the solvers' equilibrium solution efforts under the different feedback policies, we are now ready to answer our main research question: Which feedback policy is the best for each of the two stipulated objectives? We first discuss the optimal feedback policy for maximizing average performance (Section 2.5.1); we then shift our focus to maximizing the performance of the best solution (Section 2.5.2).

2.5.1. 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 $\Pi_{\text{avg}} = \mathbb{E}[v_{i2} + v_{j2}]/2 - A = \mathbb{E}\big[\sum_{i,t} e_{it}\big]/2 - A = \mathbb{E}\big[\sum_t e_{it}\big] - A$. That is, maximizing average performance is equivalent to maximizing the sum of a solver's (ex ante) expected first- and second-round equilibrium efforts. Proposition 2.4 compares the expected first- and second-round effort choices of a solver as well as the firm's expected profits for the cases of no feedback, public feedback, and private feedback.

Proposition 2.4. *The following statements hold:*
(i) $e_1^{\text{pri}} < e_1^{\text{pub}} = e_1^{\text{no}}$;
(ii) $\mathbb{E}_{\zeta_{i1}}[e_2^{\text{pri}}(\zeta_{i1})] < \mathbb{E}_{\Delta\zeta_1}[e_2^{\text{pub}}(\Delta\zeta_1)] = e_2^{\text{no}}$;
(iii) $\Pi_{\text{avg}}^{\text{pri}} < \Pi_{\text{avg}}^{\text{pub}} = \Pi_{\text{avg}}^{\text{no}}$.

The first noteworthy result of Proposition 2.4 is that, in each round, the ex ante expected effort of each solver is identical under a no-feedback and a public-feedback policy. This result can be explained by *public* feedback having two opposed effects on a solver's second-round effort choice. On the one hand, if the revealed first-round performance difference is low ($|\Delta\zeta_1| < a/3$), then each solver understands that the contest is highly competitive and is motivated thereby to expend more effort than under a no-feedback policy. On the other hand, if the performance difference is large ($|\Delta\zeta_1| > a/3$), then solvers are discouraged from investing effort because they believe that the contest is practically decided. In equilibrium, these countervailing effects of motivation and de-motivation offset each other; thus, $\mathbb{E}_{\Delta\zeta_1}[e_2^{\text{pub}}(\Delta\zeta_1)] = e_2^{\text{no}}$. Clearly, when deciding on his first-round solution effort, each solver anticipates this balance between motivation and de-motivation effects and therefore chooses to exert 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 policies. 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 if his first-round performance was middling.[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 (and will) form only a belief about the performance difference. Yet given the inherent randomness of performance, each solver knows that his competitor is unlikely to have achieved the

[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.

same performance. For this reason, solvers respond only moderately to the motivation effect of private feedback.

Second, private feedback has a strong de-motivating effect on relatively low-performing solvers. As Figure 2.2 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 these asymmetric effort choices. As a result, we observe a phenomenon that does not arise under a public-feedback regime—namely, 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 the focal solver can win becomes smaller and so he begins to shirk. In short: private feedback reduces the contest's competitiveness, which in turn leads solvers to reduce their effort.

This phenomenon also has a strong effect on a solver's effort in the *first* round. Since effort in the second round is reduced, solvers refrain from wasting 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 not observed in 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. It is therefore optimal for the firm to choose either a no-feedback or a public-feedback policy. And whereas the firm is indifferent between these two policies, solvers strictly prefer a no-feedback policy.

2.5.2. Finding the Best Solution

In practice, most innovation contests are designed to elicit one exceptional solution 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 $\Pi_{\text{best}} = \mathbb{E}[\max\{v_{i2}, v_{j2}\}] - A$. Proposition 2.5 establishes that, for certain types of innovation contests, private feedback is the optimal policy.

Proposition 2.5. (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 reflects the identity of expected efforts under these two feedback policies.

The key result of Proposition 2.5 is that public (resp., private) feedback is optimal if $\kappa < \underline{\kappa}$ (resp., if $\kappa > \overline{\kappa}$).[9] To better understand this result, recall that $\kappa = (a/k_e)^2/(A/c)$. The numerator 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, 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 use κ to denote *performance uncertainty.* Proposition 2.5 implies that, for innovation contests in which effort is 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.

Yet how can this result be explained—especially since, according to Proposition 2.4(iii), private feedback induces lower average performance than public feedback (or no feedback)? We can answer this question by considering Figure 2.3, which compares solver i's expected second-round effort conditional on his first-round shock ζ_{i1} under a private-feedback (solid line) and a public-feedback (dashed line) policy. The figure's left (resp. right) panel shows the functions for low (resp. high) performance uncertainty κ. Two key observations can be made here. First, comparing the solid and the dashed lines plotted in each panel reveals 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. Comparing the two panels reveals also that, for low performance uncertainty κ, the reduction in average effort under private feedback is much greater than for high κ. 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 finding captures the need of top performers to protect their good position

[9]The complexity of the equilibrium emerging under a private-feedback policy makes it difficult to find the dominant strategy for $\underline{\kappa} \leq \kappa \leq \overline{\kappa}$. However, numerical simulations indicate that $\underline{\kappa} = \overline{\kappa}$; hence there exists a unique threshold for κ above which a private-feedback policy maximizes the firm's expected profits.

Figure 2.3.: Ex ante Expected Second-Round Efforts under Private and Public Feedback.

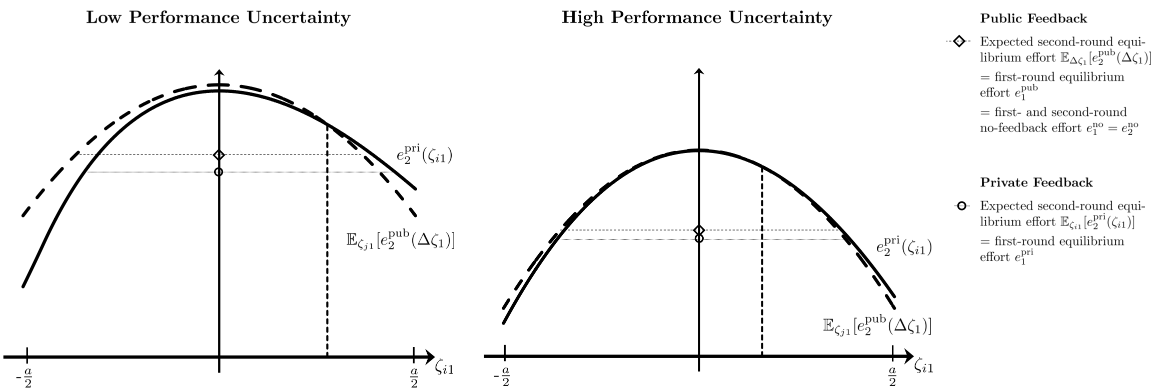

Notes: The graphs compare solver i's (expected) equilibrium second-round effort conditional on ζ_{i1} under private feedback (solid line; $e_2^{\text{pri}}(\zeta_{i1})$, as stated in Proposition 2.3) and under public feedback (dashed line; $\mathbb{E}_{\zeta_{j1}}[e_2^{\text{pub}}(\Delta\zeta_1)]$, with $e_2^{\text{pub}}(\Delta\zeta_1)$ as in Proposition 2.2); we take the expectation over ζ_{j1} in the public-feedback case in order to make it directly comparable to the private-feedback case. In the left panel, performance uncertainty is low ($\kappa = 1.01$); in the right panel, performance uncertainty is high ($\kappa = 4$). The vertical dotted line marks the unique intersection point of the two curves. Parameters are: $A = 1$, $a = 1$, and $k_e = 1$ (both panels); $c = 1.01$ (left panel); and $c = 4$ (right panel).

more determinedly under private than under public feedback owing to the lack of relative performance information. Moreover, Figure 2.3 shows that the fraction of solvers for whom private feedback increases their effort is small under low performance uncertainty but is large 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 encourages greater effort from the best solvers. Thus the optimal feedback policy is the one that best balances the average effort provision against the likelihood that a top-performing solver participates in the contest. Consider the left panel of Figure 2.3. For low κ, the decline in average effort under private feedback is relatively pronounced while 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 shows that the decline 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 farther to the left). In this case, then, private feedback is the optimal feedback policy.

This finding—that the optimal feedback policy is tightly linked to the relative importance of effort and uncertainty—has not been addressed in the extant literature and it has two immediate managerial implications. First, when setting up an innovation contest, it is crucial that the firm identifies the extent to which a solver's performance depends on stochasticity. For instance, there is seldom much uncertainty associated with contests that seek to foster incremental innovation. In such contests, the hosting firm should provide public feedback. However, private feedback is the preferred choice for ideation contests that aim to develop novel concepts, new ideas, or breakthrough research (all of which are associated with high levels of uncertainty). Even more, our results concern communications between solvers: whenever the effort–performance link becomes tenuous, idea exchange between solvers becomes detrimental to firm goals and so the firm should minimize any form of communication between competing solvers. Second, if performance uncertainty is substantial then a firm should not focus on improving the average second-round effort; instead, it should choose a feedback mechanism that "pampers" potential first-round top performers—even if in realization such top performers may not be present in the contest.

From the social welfare standpoint, Proposition 2.5 (in conjunction with Proposition 2.4) provides another important result for contests that are inherently uncertain: the private-feedback policy is not only optimal for the firm but can also be *socially efficient*. More precisely: apart from maximizing the firm's expected profits, a private-feedback policy also allows solvers to reduce their expected efforts. As a consequence, both the firm and the solvers may well prefer private feedback to either public or no feedback in settings characterized by high performance uncertainty.

2.6 Voluntary Feedback

So far we have assumed that the firm, by dictating the type of feedback, also obliges solvers to submit intermediate solutions. Those intermediate solutions are the firm's vehicle for providing solvers with interim performance feedback. In reality, however, instead of enforcing intermediate submissions, a contest holder may also simply offer an option for feedback. In such cases, feedback is voluntary and it is each solver's deliberate choice whether or not to seek feedback. Clearly, a solver will only do so if he sees a significant benefit in submitting his intermediate solution. We now explore how different potential benefits affect a solver's decision.

For concreteness, we extend our base model (as depicted in Figure 2.1) by allowing each solver to decide—after his first-round efforts—whether or not to submit an interim solution. The firm then follows its announced feedback policy and provides the solvers with accurate feedback on any submitted solutions. That is, if feedback is private then solver i receives feedback on v_{i1} only if he has submitted his solution beforehand; under a public-feedback policy, solver i can receive feedback on v_{j1} also—but only if solver j has submitted his interim solution. The following proposition characterizes the solvers' equilibrium behavior.

Proposition 2.6. (i) *Under a public-feedback policy, no solver voluntarily discloses his first-round solution in equilibrium.*

(ii) *Under a private-feedback policy, each solver submits his first-round solution in equilibrium.*

The key insight here is that a solver's behavior as regards submitting interim solutions depends critically on the contest holder's feedback policy. It is intuitive that under a private-feedback policy, each solver always seeks feedback as there are no negative externalities from requesting feedback. Without disclosing any information to his competitors, each solver receives more refined information on his performance, allowing him to optimally adjust his second-round efforts. In contrast, under a public-feedback policy, solvers refrain from submitting their interim solutions because the threat of an escalation of effort provision outweighs any potential benefits. More precisely, public feedback introduces a relative benchmark that induces solvers to invest exceptionally high effort levels if the contest is competitive. However, such bold effort choices are utterly futile because each solver invests the same amount of effort in equilibrium, and therefore the chances of winning are unaffected. Instead, it is only the costs of effort that skyrocket.

The same logic continues to hold when accounting for another potential benefit of feedback: the reduction of uncertainty in the second round. In fact, as Proposition 2.2

indicates, the problem of effort escalation becomes even stronger because with a lower performance uncertainty (i.e., a lower a) equilibrium efforts are becoming even higher.

We conclude that whereas pure information incentives are strong enough for a solver to reveal his first-round performance under private feedback, they are insufficient under a public-feedback policy. Thus to further strengthen the solvers' incentives the firm may need to resort to monetary incentives by granting a so-called "milestone" award to the solver with the best first-round solution. It is evident that such financial prospectives will—if large enough—induce each solver to submit his intermediate solution. However, as the next proposition shows, this is never in the firm's interest.

Proposition 2.7. *Under a public feedback policy, it is never optimal for the firm to grant a milestone award.*

The higher the milestone award the more the firm shifts the solvers' incentives away from winning the overall contest towards winning the first round. As a result, the introduction of a milestone award drastically dilutes a solver's second-round incentives, and thus his equilibrium efforts. Per se, however, the firm is not interested in these intermediate submissions, but it only cares about the final solution qualities. Thus, the higher the milestone award the more misaligned are the firm's and the solvers' objectives. This is why the firm strictly prefers to not grant any milestone awards.

In conclusion, we have shown that in the absence of monetary rewards, submitting interim solutions and populating a publicly available leader board are not decisions that a rational solver would make. It is thus an important alley for future empirical research to investigate why, in practice, many solvers are nonetheless eager to share their solution quality with competitors.

2.7 Partial Information Feedback

Until now we have assumed that if the firm provides feedback then this feedback is perfect; in other words, the firm always reveals fully accurate information on the solvers' performance. However, the firm may either prefer or (for practical reasons) be required to provide only partial feedback. Two canonical cases are practically relevant. First, the firm may provide information that is less detailed—for example, by publishing the solvers' rankings but no specific information about their performance. Second, the firm's performance feedback may be disturbed by some level of noise.

2.7.1. Rank-Only Feedback

We first consider the case in which the firm reduces the granularity of feedback by providing only information on the solvers' ranking in the contest. Before delving into the details, it is important to recognize that such rank-only feedback cannot be conceptualized as private information; revealing a rank is inherently public feedback. It is not material for a solver which other solver holds which rank, but rather that one of the other solvers holds each of the ranks above or below him.

As compared with a public-feedback policy, rank-only feedback may change solver behavior in two ways. When first-round performances are extremely close, providing rank-only feedback reduces competition because the "blurriness" of the rank information leaves none of the solvers fully aware of how close the competition actually is. When first-round performances instead vary widely, rank information results in solvers underestimating their relative performance difference and so induces them to exert more second-round effort than actually required by the situation. Our next proposition compares the relative strength of these effects.

Proposition 2.8. *The unique PBE under a rank-only feedback policy is identical to the unique PBE under a no-feedback policy. In particular,* $e_1^{\mathrm{r}} = e_1^{\mathrm{no}}$ *and* $e_2^{\mathrm{r}} = e_2^{\mathrm{no}}$.

Proposition 2.8 holds a surprise. In comparison with a no-feedback policy, providing rank-only feedback has *no* effect on the equilibrium behavior of solvers; that is, solvers completely ignore their respective rankings. This outcome is in stark contrast to the fate of accurate performance feedback, which is always used by solvers to adjust their second-round behavior. The implications for practice are striking. If a contest holder wants its feedback to influence the second-round efforts of solvers, then this feedback must include information about each solver's actual performance; rank-only feedback will not have the desired effect.

2.7.2. Feedback with Noise

Next we analyze the implications of noisy feedback, which are important for a firm that cannot (or prefers not to) guarantee perfectly accurate feedback. Consider again the case of Kaggle. Some contest holders provide interim performance feedback by using only a sample of the full data set on which to test the solvers' algorithms; this approach helps prevent overfitting of the data during the contest. However, the final ranking is based not on that sample but rather on the full data set. So from the solvers' perspective, the interim performance feedback is not entirely accurate: it is disturbed by noise.

Two questions arise. First, does the introduction of noise change the balance between no feedback, public feedback, and private feedback? Second, how does noisy feedback compare to entirely accurate feedback? Is it possible for the firm to *benefit* from introducing noise into the feedback mechanism? Here we set out to start answering these questions.

Treating all the possible facets of noise is beyond the scope of this paper. Instead, we investigate a simple example of noisy feedback that nonetheless captures the essence of masking a solver's true performance. More specifically, we assume that the firm gives perfectly accurate feedback with publicly known probability q, and with probability $1-q$ it transmits absolutely uninformative feedback (i.e., white noise).

Proposition 2.9. (i) *Under a public-feedback policy, the firm's expected profits* $\Pi_{\text{avg}}^{\text{pub}}$ *and* $\Pi_{\text{best}}^{\text{pub}}$ *are invariant with respect to* q*.*

(ii) *For any fixed* $q > 0$*, there exists a* $\overline{\kappa} < \infty$ *such that* $\Pi_{\text{best}}^{\text{pri}} > \Pi_{\text{best}}^{\text{pub}}$ *for all* $\kappa > \overline{\kappa}$*.*

Part (i) of this proposition establishes that, under a public-feedback policy, noise does not affect the firm's expected profits in terms of either average performance or best performance. The implication is that, under a public-feedback policy, contest holders cannot use noise strategically to improve contest outcomes. In contrast, noise does affect the firm's profits under a private-feedback policy; however, exact analytical expressions are difficult to derive. Yet our numerical studies reveal that profits are monotonic in q: for small κ, a no-feedback policy (i.e., $q = 0$) maximizes the firm's profits; for large κ, accurate private feedback (i.e., $q = 1$) is optimal. It seems once again that noise cannot be deployed strategically to improve contest outcomes.

Combining Proposition 2.9(i) for public feedback and part (ii) for private feedback indicates that our results about the preferability of different feedback types are robust to the introduction of noise. That is, noise does not impact the ranking of the different feedback policies and hence the selection of a feedback policy should not be affected by the accuracy of the feedback.

2.8 Robustness and Extensions

We explored the sensitivity of our conceptual results to key modeling choices. Our results are exceptionally robust to changes in those assumptions, as we summarize in this section.

Solver asymmetry. Across solvers, the distribution of random shocks may be asymmetric. Two possible sources of such asymmetry are differences in mean and differences in variance; the former (resp., latter) signifies inherent differences in solvers' capabilities (resp., innovation processes). A closed-form analysis of an extended model establishes that our results are robust to both sources of asymmetry; moreover, it seems that our assumption of solver symmetry is actually conservative with respect to the true value of private feedback.

Alternative feedback policy. Another form of partial information feedback is to inform solvers only of whether (or not) they have surpassed a certain performance threshold. Our formal investigation of such a threshold feedback policy shows that it can never improve on the performance of a fully informative policy—in accordance with our results in Section 2.7.

Cost of effort. For some contests, solvers may be more concerned with their total expended effort than with their respective individual efforts in each round. In such cases, a more appropriate cost-of-effort function would be $c(e_{i1} + e_{i2})^2$. We can demonstrate analytically that allowing for effort interaction effects between the two rounds does not alter our results in any meaningful way.

Performance uncertainty. Depending on the contest's particular context and the innovation process of solvers, random shocks can follow a multitude of distribution functions. After conducting a large number of numerical experiments with normal and beta distributions, we can report that our results are robust.

2.9 Conclusions

Contests are a frequently used mechanism for providing incentives when companies and government agencies source innovation. Prize competitions organized via the Internet (as hosted by Kaggle) are contests, and so are sourcing efforts organized via the European Union's "competitive dialogue procedure" and many of the approaches taken by private companies to source custom-engineered innovations. Feedback has been extensively used in practice to improve both the efficiency and the efficacy of contests. However, a rigorous understanding of when and how to provide which kind of feedback—and of when to refrain from giving feedback—is rather limited. The primary goal of this paper is to begin building a more comprehensive understanding of feedback in contests.

Our main contribution consists of charting a practically relevant landscape—one that determines how feedback can be used in contests—by addressing three questions that together define a feedback policy's information structure: Who receives the feedback information? Who decides about which contestants receive feedback? What should be the information content of the performance feedback that is given? Answering these questions allows us to analyze many forms of feedback that actually are used in contests and to prescribe beneficial policies for a wealth of settings. In doing so we build new insights and challenge existing ones.

It is remarkable that—almost irrespective of whether feedback is voluntary and of whether feedback includes all or only some of the available information—firms need only focus on two straightforward factors when choosing whom to provide with any feedback: the *contest objective* (average versus best performance) and the contest's inherent *performance uncertainty*. If the firm is concerned about the solvers' average performance, then either no feedback or public feedback is preferable to private feedback. The same preference obtains if the firm is interested in the best performance, provided that performance uncertainty is low. However, private feedback is the optimal choice if the firm seeks the best possible performance and performance uncertainty is high. Hence, private feedback is most suitable for innovation settings.

Our findings have immediate managerial implications. Contest holders that aim to raise the overall effort level among all solvers should refrain from giving private feedback; thus, if performance information is released then it should be made public. Incremental innovation contests likewise do not benefit from private feedback; 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 rely solely on private feedback.

As for who should decide on whether feedback is provided, one must bear in mind that voluntary feedback can function only if the solvers have an incentive to seek it. Generally speaking, that incentive may be of two different forms: an informational advantage or a monetary reward. As regards informational advantages, solvers always (resp. never) ask for feedback under a private-feedback (resp. public-feedback) policy. That is, the informational advantage outweighs the informational cost with private but not with public feedback. Monetary rewards intended to induce feedback-seeking behavior do work, but they are never optimal from the firm's perspective.

Our results on voluntary public feedback have consequences for practice. If the contest holder anticipates benefits from having solvers share information about their performance, then it should find ways to make such feedback mandatory and not leave feedback up to the solver. Under private feedback, however, the opposite is true. Since in this case the solvers will always ask for feedback, contest holders should offer feedback only if they truly want to provide it. Intermediate prizes are never advisable—that is, from an incentive perspective.

Finally, if the feedback recipients have been identified and if the choice of voluntary or mandatory feedback has been made, then there remains the question of how granular the feedback should be. There is no evidence—in the cases examined here—that reducing feedback granularity and/or accuracy benefits the contest holder. Even more interestingly, we show that feedback lacking specific performance-related information (e.g., rank-only feedback) will likely be disregarded by solvers and thus fail to achieve its intended effect.

One aspect that must be considered when interpreting our results is that we did not explicitly incorporate the cost of providing feedback. We made this choice for two reasons. First, in nearly all practical cases the cost differences among no feedback, public feedback, and private feedback are small compared with the benefits of providing tailored incentives to solvers (in our Kaggle and EU examples, the cost of giving feedback is negligible when compared with the potential benefits of an optimal feedback policy). Second, including such costs would be trivial, both technically and conceptually, because adding a cost term to the firm's profits has no effect on the equilibrium analysis. Decision makers can simply subtract the cost differences between feedback policies from their respective benefit differences (as derived in this paper) to determine the overall trade-off.

Our model has limitations that should be mentioned. In order to maintain tractability and develop a parsimonious model, we made some assumptions about the purpose of contests; in particular, we focus on the *incentive* effects of feedback in contests. However, feedback may be used also to guide solvers in terms of where to direct their search efforts. Understanding this directional aspect of contest feedback requires an approach that differs fundamentally from ours and is a promising avenue for future research. Another limitation is our focus on a contest's *rational* motivation effects, which are amenable to game-theoretic analysis. Yet real-world contests may also effect their outcomes by way of psychological inducements. Within the game-theoretically rational framework we consider only two competing solvers; this assumption ensures tractability but also introduces symmetries that may be absent in contests with three or more solvers. Hence, there is no guarantee that our results hold for more than two players; thus our results should be interpreted with some caution for large contests. Finally, we assume an additive performance relation between effort and uncertainty. This approach can be viewed as a first-order Taylor approximation of a more general relationship, and adding higher-order terms could capture additional effects—for example, an increase in uncertainty with greater solution efforts.

Previous research on contests has not broadly explored the repercussions of many practically relevant feedback policies. The aim and principal contribution of this paper is to fill several critical gaps in the literature and to build a deeper understanding of feedback in contests. It is only by considering the many variants and aspects of feedback that managers can reasonably hope to make the contest mechanism—a method often relied upon in practice for sourcing innovation—more efficient and effective.

Chapter III

How to Procure Complex Innovation

with Zhi Chen and Jürgen Mihm[1]

3.1 Introduction

Access to innovative products is a key source of competitive advantage in many industries (Pisano 2015). Traditionally, firms guaranteed themselves such access by pursuing internal research and development (R&D) projects and by making substantial investments into new product development. More recently, however, firms have extensively tapped into an alternative source of innovation: their suppliers (Cabral et al. 2006). One industry that stands out in demonstrating how effective such a use of suppliers as source for innovation can be is the automotive industry. Consider—as a prime example—the recent advances in the area of automotive lighting systems.

Since its inception in 1898, when the Electric Vehicle Company of Hartford introduced the first electric headlamps for its Columbia Electric Car, electric lighting systems have become the unquestioned standard for automotive lighting systems. Over the course of the last century, the initial tungsten filament technology was replaced by two major innovations—the introduction of Halogen lighting in 1962 and high-intensity discharge lamps (also known as Xenon lamps) in 1991. These new technologies led to great improvements in driver and traffic safety, but did not spur much further innovation. It is only with the more recent introduction of LED headlamps in 2003 by Hella—a leading automotive supplier—that automotive lighting systems have seen a drastic ac-

[1]The research presented in this chapter is based on a paper entitled "Sourcing Complex Innovation: Integrated System or Individual Components?", coauthored with Zhi Chen and Jürgen Mihm.

celeration in innovation, and most of these innovations were introduced not by the car manufacturers themselves but by their suppliers.

But why are today's automotive suppliers developing new technologies at an ever more rapid pace? One major reason is a shift in the car manufacturers' procurement process for complex innovations such as an automotive lighting system. Instead of relying on direct contracting with one of their suppliers, car manufacturers nowadays frequently organize so-called procurement or innovation contests to simultaneously gain access to the innovation capabilities of their full supplier base. In hosting such a procurement contest, car manufacturers only announce the minimum specifications for the innovation to be developed, but detail development rests in the responsibility of each individual supplier participating in the contest. The promised award—typically a supply contract for the winning design—incentivizes the suppliers to engage in development activities and makes them present their solutions to the car manufacturer. The car manufacturer—after evaluating all submitted designs—can then *ex post* select the best design alternative, and award the supply contract accordingly. The advantages of using a procurement contest are evident: Not only is a procurement contest an informationally very parsimonious purchasing mechanism, but it also induces high innovation efforts in the car manufacturer's supplier base and it allows for an ex post selection of the best product alternative (Rogerson 1989). Yet to fully reap all these benefits it is imperative for a car manufacturer to carefully design the setup of his procurement contests. This is the key challenge that we study in this paper.

Most notably, procurement contests should be tailored towards the properties of the desired product and the structure of the supplier base. With regard to the product characteristics, a central feature of *complex* innovations is that they are made of multiple components that only together form the full product (or in engineering terminology, the integrated system). Consider again our example of an automotive lighting system. Broadly speaking, such lighting systems consist of two components: a light module, which is part of the car's exterior design and is responsible for the light emission; and an electronic control module, which steers the different functionalities of the light module. In the face of such a complex—or multi-component—innovation, a central question for any car manufacturer is whether to buy the full product from a single supplier, or whether to source the different components from different suppliers. In the former case, the car manufacturer hosts a *system contest* and requires suppliers to submit full products, whereas in the latter case, the car manufacturer hosts a separate *component contest* for each individual component.

Answering the question of when to prefer which contest format (system vs. component contest) is at the very heart of our research. We conjecture that many different dimensions—both on the product as well as the supplier side—will have an impact on a car manufacturer's optimal choice. One of the key influence factors is certainly the magnitude of innovation that is required to develop each of the components, and hence the integrated system. Specifically, are both components rather incremental innovations or is one of them—or even both—radically new? And what is the technological relationship between the components? That is, are the components technological substitutes or complements? Besides these product-level features, other important aspects relate to the size of the supplier base, the degree of heterogeneity in the supplier base as well as the level of performance correlation exhibited by each supplier. In other words, how many suppliers can participate in the procurement contest, are there structural performance differences among the suppliers, and is a supplier who is good at developing one component (e.g., the light module) also likely to be good at developing the other component (e.g., the control module), or not?

In this paper, we develop a game-theoretic model that takes all these factors into account, and we investigate how the individual factors impact the optimality of either contest format (system or component contest). In doing so, we also shed light on the optimal properties of a contest holder's supplier base, and we study which kind of suppliers the contest holder should invite to participate in the different contest formats. Our detailed contributions are as follows.

First, we show that a firm should use a system contest if all product components are merely incremental innovations and the firm's supplier base is relatively small. In all other cases—that is, if at least one component demands radical innovation or the firm has access to a large supplier base—the firm should opt for hosting a separate procurement contest for each individual component. For incremental innovation it is the firm's key ambition to provide significant incentives for development effort, and since a system contest offers the highest stakes (i.e., the biggest supply contract) it is the firm's preferred option in less innovative environments. However, the benefits of high effort incentives may be overshadowed if controlling and managing uncertainty moves to the firm's center of attention—a situation that arises naturally in the presence of radical innovation or when the supplier base is large. As an interesting side result we find that these general guidelines are independent of the technological relationship between components as well as the level of heterogeneity and performance correlation in the supplier base.

Second, we finely delineate how the firm should organize its supplier base, and we also identify which suppliers should be invited to participate in the different contest formats. A couple of managerial rules emerge: First, the firm should work towards a relatively homogeneous supplier base as any form of performance handicaps deters effort incentives, both for relatively strong as well as relatively weak suppliers. Second, when hosting component contests the firm benefits from inviting the same set of suppliers to the individual contests, instead of inviting different sets of suppliers, and these benefits are largest when components are technological complements. Third, in a system contest, the firm's optimal choice of suppliers depends on the innovativeness of the desired product. For incremental innovations the firm should seek for suppliers that exhibit only minor levels of performance correlation across components, whereas the firm should strive for higher performance correlation when radical innovation is involved.

Finally, we also investigate how systematic differences in performance correlation across the different contest formats affect our structural results. We find that all our results are fully robust to such differences, with one notable exception. If the magnitude of performance correlation in a system contest is substantially larger than in a component contest, then a system contest is optimal even when the full product requires radical innovation. The reason for this finding is that with large differences in performance correlation between system and component contests, a system contest allows the firm to more effectively exploit the large uncertainties associated with radical innovation, which otherwise is the key strength of a component contest.

3.2 Related Literature

The design of effective procurement mechanisms has been a longstanding concern in the academic literature (Vickrey 1961, Rob 1986, Laffont and Tirole 1993, Elmaghraby 2000, Beil 2010). It is beyond dispute among scholars and practitioners alike that any successful procurement policy must be tailored to the properties of the desired product. As outlined in Cabral et al. (2006), one product feature that is of paramount importance when selecting the right procurement policy is the degree of innovativeness. In practice, products can be broadly divided into two categories: standard products that are well-specified and exhibit only minor performance uncertainties; and innovative products that have yet to be developed and whose final performance is thus merely a speculation.

Most of the existing procurement literature has concentrated on devising optimal procurement mechanisms for standard products (see, e.g., Che 1993, Cachon and Zhang 2006, Chen 2007, Wan and Beil 2009, Hu et al. 2013, and references therein). Since the specifications of these products—and thus their performance—are known before the procurement process is initiated, buyers aim to find the supplier who can deliver the product most efficiently, and extant research has investigated many such dimensions of efficiency. More precisely, most of the early literature has primarily concentrated on how to lower procurement costs (Rob 1986, Dasgupta and Spulber 1990); however, subsequent work has accounted for such diverse factors as integrating product quality (Che 1993, Beil and Wein 2003) and delivery lead-time (Cachon and Zhang 2006) into the procurement process, reducing transportation costs (Chen et al. 2005, Kostamis et al. 2009), engaging in supplier qualification (Wan and Beil 2009, Wan et al. 2012), improving supplier reliability (Yang et al. 2009, Chaturvedi and Martínez-de-Albéniz 2011, Yang et al. 2012), anticipating future changes in competitive structures (Li and Debo 2009), building a stable supplier base in the long run (Chaturvedi et al. 2014), and motivating suppliers to invest in cost-reduction efforts (Arozamena and Cantillon 2004, Li 2013, Li and Wan 2016).

All these papers—albeit investigating very different dimensions of procurement efficiency—agree in their proclamation of auctions as the preferred procurement mechanism. In fact, since its inception by Vickrey (1961) auction theory has become the focal mechanism in procurement theory, and over the course of time, classical auction models have been enriched to include the peculiarities of many different operational settings (see, e.g., Beil and Wein 2003, Kostamis et al. 2009, Wan and Beil 2009). However, when procuring an innovative product neither the buyer nor the suppliers can envision the final product's true performance at the outset of the procurement process. As a result it is impossible for a supplier to credibly quote his product's future performance, let alone that the buyer can identify ex ante (i.e., before product development) which supplier will deliver the best product. As a result the use of traditional auction formats is prohibitive for innovative goods, and hence buyers have to resort to alternative procurement mechanisms. Both in theory and practice it is the contest mechanism that has proven most effective for sourcing innovative products (Lichtenberg 1988, Rogerson

1989, Elmaghraby 2000, Scotchmer 2004, Cabral et al. 2006). The key advantage of hosting a procurement contest is its ease of practical implementation. In a procurement (or innovation) contest, the buyer only sets a prize and a time deadline at the outset of the contest, and after the deadline awards the prize to the contestant that submitted the best product (Taylor 1995). In particular, there is no need that the performance of the submitted products be verifiable, since the buyer has no incentives to manipulate the outcome of the contest given that he has to pay the award to one of the contestants. As a result, contests are not only easily implementable, they are also informatively parsimonious (Cabral et al. 2006).

Inspired by these benefits of contests, a wealth of research has evolved around the question of how to effectively design and manage innovation contests. Building on the pioneering work of Lazear and Rosen (1981), Green and Stokey (1983) and Nalebuff and Stiglitz (1983), the contest literature provides rich insights into such fundamental contest design issues as the design of optimal award schemes (Che and Gale 2003, Ales et al. 2017), the restriction of access to the contest (Taylor 1995, Moldovanu and Sela 2001, Terwiesch and Xu 2008), or how to manage the contest as it unfolds (Aoyagi 2010, Mihm and Schlapp 2017). Taken together, the above mentioned papers have led to a coherent theory of how to source a single innovative product via contests. In this paper, however, we are interested in how a buyer should design a procurement contest when simultaneously sourcing multiple heterogeneous, but dependent products. To the best of our knowledge, this question has not been addressed in prior work.[2]

In particular, we make the following contributions to the theory of procurement and contests. First, we demonstrate how different formats of procurement contests can be used to source complex innovative products that consist of multiple components, and we show how the choice of contest format is governed by the level of innovation required and the size of the supplier base. Also, by considering the technological interactions between the individual components we are able to answer the above questions for a large variety of different products commonly found in practice.

[2] It is worth mentioning that the literature on multi-object auctions is concerned with a similar question, but only for standard products (see, e.g., Armstrong 2000, Palfrey 1983, Elmaghraby and Keskinocak 2004, Hausch 1986, Avery and Hendershott 2000). As an immediate implication, the tradeoffs governing the buyer's choice of procurement format are substantially different. Most notably, for innovative products the buyer seeks to balance development effort incentives with performance uncertainty considerations—two aspects that are absent in procurement auctions for standard goods.

As our second key contribution we shed light on the ideal characteristics of a buyer's supplier base when hosting procurement contests. Besides the mere size of the supplier base, we also consider more subtle factors such as performance correlation across different component development projects and the heterogeneity of the supplier base, and we analyze how the optimal supplier base depends on the characteristics of the complex innovation to be procured.

3.3 The Model

Consider a firm (from hereon "the buyer") organizing a procurement contest among $n \geq 2$ suppliers with the goal of buying an innovative product that consists of two components $j \in \{1, 2\}$. The buyer can choose between two different contest formats to source the product: a system contest, and a component contest. In a *system contest* the buyer asks the suppliers to submit solutions for the full product (i.e., the integrated system) and promises an award $A > 0$ to the supplier that delivers the product with the best performance. On the contrary, in a *component contest*, the buyer holds a separate contest for each of the two components, and offers an award $A_1 = pA$ (resp. $A_2 = (1-p)A$) to the supplier that delivers the best performing component $j = 1$ (resp. $j = 2$), with $p \in [0, 1]$.

The goal of the buyer is to maximize the performance S of the product to be procured. Since the full product consists of two components, its performance depends on two key attributes: (i) the individual performance of each component, and (ii) the technological relationship $t : \mathbb{R}^2 \to \mathbb{R}$ between the two components. Specifically, in line with extant research, we consider that—depending on the practical context—components can be either technological substitutes or complements (see, e.g., Bhaskaran and Krishnan 2009, Roels 2014, Gurvich and Van Mieghem 2015, and references therein). In the former case the buyer can trade off the two components' individual performances, and the performance of the integrated system is simply the sum of the individual components' performances; i.e., $t(x, y) = x + y$. To the contrary, in the latter case both components are essential for the performance of the integrated system, and total performance is thus the minimum of the two components' performances; i.e., $t(x, y) = \min\{x, y\}$.

To answer our research questions, it is imperative to express the performance of the full product S as a function of the chosen contest format and the technological

Table 3.1.: Performance of the Integrated System.

	System contest	Component contest
Tech. substitutes	$S^{sys}_{sub} = \max_i\{v_{i1} + v_{i2}\}$	$S^{cpo}_{sub} = \max_i\{v_{i1}\} + \max_i\{v_{i2}\}$
Tech. complements	$S^{sys}_{cml} = \max_i\{\min\{v_{i1}, v_{i2}\}\}$	$S^{cpo}_{cml} = \min\{\max_i\{v_{i1}\}, \max_i\{v_{i2}\}\}$

relationship between the two individual components. We achieve this by introducing the variable v_{ij} which denotes the performance of component $j \in \{1, 2\}$ as developed by supplier $i \in \{1, \ldots, n\}$. In a system contest, each supplier i submits an integrated system with performance $S_i = t(v_{i1}, v_{i2})$, and the buyer eventually procures the product with the highest performance; that is, $S^{sys} = \max_i\{S_i\}$. In a component contest, however, the buyer procures the two best components $\hat{v}_1 = \max_i\{v_{i1}\}$ and $\hat{v}_2 = \max_i\{v_{i2}\}$ individually, and the performance of the full product is then $S^{cpo} = t(\hat{v}_1, \hat{v}_2)$. Table 3.1 summarizes the performance of the integrated system as a function of the contest format and the technological relationship between the components. We now elaborate in more detail on our model setup.

3.3.1. Sequence of Events

The procurement process begins when the buyer publicly announces the contest format (i.e., system or component contest), the total award A, and in the case of a component contest the prize split p. To focus on the buyer's choice of contest format (and to simplify the mathematical exposition), we treat the total award A as a parameter. As such, the buyer's primary decisions are which contest format to choose, and, if applicable, how to divide the total award between the two different component contests (i.e., the prize split p).

After the buyer's public contest announcement, each supplier $i \in \{1, \ldots, n\}$ decides whether to participate in any of the offered contests. When participating, supplier i invests an unobservable solution effort $e_{ij} \geq 0$ to develop component $j \in \{1, 2\}$ at private costs $c(e_{ij})$, where c is a twice continuously differentiable, increasing, and strictly convex function with $c(0) = 0$ and $c'(0) = 0$. The performance of component j as developed by supplier i is given by $v_{ij} = r(e_{ij}) + \zeta_{ij}$, where $r(\cdot)$ captures the deterministic effort-performance relationship and ζ_{ij} represents a stochastic performance component. We assume that r is a twice continuously differentiable, strictly increasing, and concave function with $r(0) = 0$ and $\lim_{x\to\infty} c'(x)/r'(x) = \infty$. The realization of ζ_{ij} is unobserv-

Figure 3.1.: Sequence of Events.

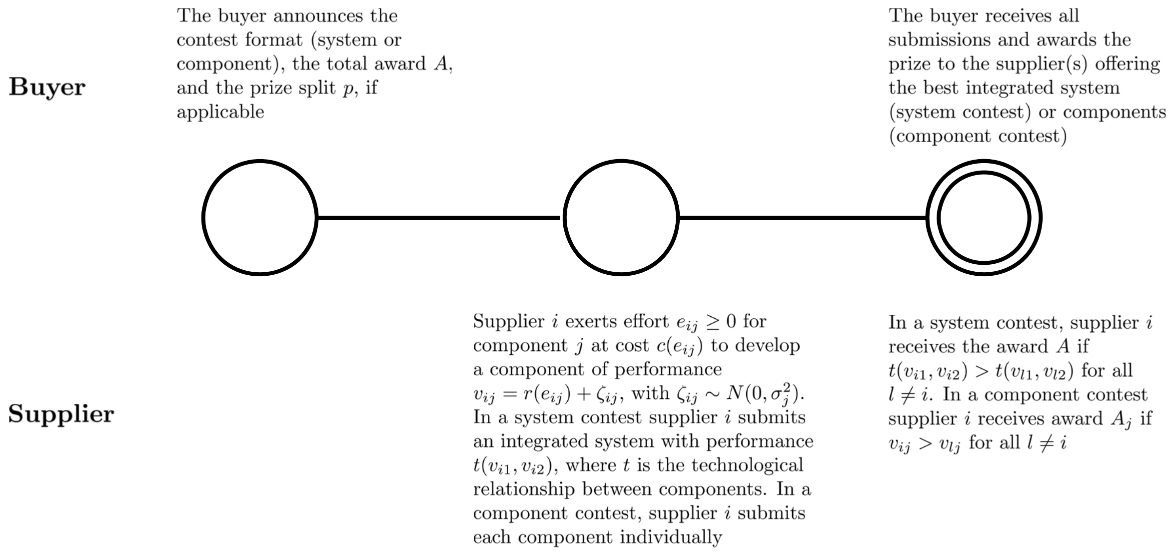

able to all suppliers and the buyer. For each supplier i, $\zeta_i = (\zeta_{i1}, \zeta_{i2})$ follows a bivariate Normal distribution with marginal distributions $\zeta_{i1} \sim N(0, \sigma^2)$ and $\zeta_{i2} \sim N(0, k^2\sigma^2)$ and correlation $\rho \geq 0$, where $\sigma > 0$ and $k \geq 1$. We assume that ζ_i is independent across suppliers.

Finally, after receiving the suppliers' submissions, the buyer evaluates their performances and awards the pre-announced prize to the supplier with the best performance. In particular, in a system contest, supplier i wins prize A if $t(v_{i1}, v_{i2}) > t(v_{l1}, v_{l2})$ for all $l \neq i$. Analogously, in component contest j, supplier i wins prize A_j if $v_{ij} > v_{lj}$ for all $l \neq i$. In all cases, ties can be broken by invoking an arbitrary rule. All primitives of the model are common knowledge and Figure 3.1 summarizes the sequence of events.

3.3.2. Model Implications

The buyer's goal is to maximize the expected performance of the best integrated system that can be built based on the suppliers' submissions; i.e., $\Pi = \mathbb{E}[S]$, where S is as displayed in Table 3.1. In contrast to the buyer, the suppliers are not concerned with the absolute performance of their submissions, but instead they are only interested in their relative performance as compared to their competitors. More precisely, each supplier's primary interest is to win any of the contests that he participates in. In a system contest, supplier i gains a utility of $U_i = A - c(e_{i1}) - c(e_{i2})$ if he wins the contest, and $U_i = -c(e_{i1}) - c(e_{i2})$ if he loses the contest. Similarly, in a component contest, supplier i's utility from winning component contest j is $U_{ij} = A_j - c(e_{ij})$, and losing leads to a utility of $U_{ij} = -c(e_{ij})$. We normalize the utility of each supplier's outside option to zero, which implies that all suppliers find it worthwhile to participate in any of the offered contests. Additionally, in Section 3.6 we extend our base model to account for heterogeneities in the suppliers' award valuations. Specifically, we later consider scenarios in which each supplier i receives a different utility $\alpha_i A$, $\alpha_i \in (0, 1]$, from winning an award of size A.

We are interested in symmetric pure-strategy perfect Bayesian Nash equilibria of the different contest formats. For such equilibria to exist, we need to invoke two additional technical assumptions.

Assumption 3.1. *For components that are technological substitutes (resp. complements), we assume that $\sigma > \sigma_{sub}$ (resp. $\sigma > \sigma_{cml}$), where σ_{sub} (resp. σ_{cml}) is the smallest σ such that for all $\sigma > \sigma_{sub}$ (resp. $\sigma > \sigma_{cml}$) supplier i's expected utility function has a unique maximum for any contest format given that the other supplier's play $e_{-i,j} = x_j I$ for all $j \in \{1, 2\}$, where $x_j \geq 0$ and I is a vector of ones.*

In a nutshell, Assumption 3.1 requires that the performance uncertainty involved in any of the contest formats must be sufficiently large—a condition that is typically true for innovative products. In addition, it is a common assumption used throughout the entire contest literature (see, e.g., Nalebuff and Stiglitz 1983, Aoyagi 2010, Ales et al. 2017, Mihm and Schlapp 2017). To better understand the intuition behind this assumption suppose to the contrary that performance uncertainty would be negligible. Then any infinitesimally small additional amount of effort by a supplier would lead to almost sure winning of the contest, and as a result, no pure-strategy equilibrium could exist. Moreover, it is straightforward to show that the thresholds σ_{sub} and σ_{cml} always exist because each supplier's utility function becomes strictly concave for large enough σ.[3]

Assumption 3.2. $3r''(x)/r'(x) < (c'''(x)r'(x) - c'(x)r'''(x))/(c''(x)r'(x) - c'(x)r''(x))$ *for all* $x \geq 0$.

[3]In fact, when deriving more explicit sufficient conditions for the existence of pure-strategy equilibria, papers in the contest literature typically focus on verifying concavity of a supplier's utility function (Aoyagi 2010, Ales et al. 2017).

Although Assumption 3.2 looks unwieldy, it has an important managerial interpretation. Specifically, it ensures that the buyer is actually interested in procuring both components. Without this condition in place, the buyer could find it optimal to buy only one component, and—at the end of the procurement process—be left with an incomplete product. We discard such a situation from further consideration as it does not mirror the practical situations that we are interested in. To gain further intuition for Assumption 3.2 note that a sufficient condition for it to hold is $c'''(x)/c'(x) \geq r'''(x)/r'(x)$ for all $x \geq 0$. That is, each supplier's marginal cost of effort—compared to the marginal return on effort—must increase sufficiently quickly. In other words, Assumption 3.2 precludes situations in which a supplier's cost-return-ratio improves overly strong as he invests more effort. Furthermore, besides its managerial relevance, Assumption 3.2 is also technically only mildly restrictive as many standard functional relationships for r and c that are used in the contest literature satisfy this condition (such as, e.g., polynomial or logarithmic r, and polynomial or exponential c).

3.4 The Optimal Contest Format

In this section we characterize the buyer's optimal choice of contest format, assuming that each supplier's performance shocks ζ_{i1} and ζ_{i2} are independent across components (i.e., $\rho = 0$). The study of correlated performance shocks is deferred to Section 3.5. Focussing, for now, on independent performance shocks allows us to clearly elicit how the two main contextual parameters of our model setup affect the buyer's choice of contest format: (i) the magnitude of innovation involved in developing each component as measured by σ; and (ii) the relative degree of innovativeness between the two components as measured by k. We begin with establishing the optimal contest format for a product whose components are technological substitutes (Section 3.4.1), and then move on to characterize the optimal contest format for a product with complementary components (Section 3.4.2). Throughout, to facilitate readability and to concentrate on the managerial implications of our analysis, we do not present any mathematical derivations in the main body of the manuscript. The technically interested reader is referred to Appendices B.2 and B.3, where we provide closed form derivations of all equilibrium results.

3.4.1. Technological Substitutes

To answer the question of which contest format is superior for the buyer, it is imperative to understand the main benefits of the different contest formats and how these benefits are affected by the absolute magnitude of innovation required (i.e., σ) and the relative degree of innovativeness between components (i.e., k). Recall that the performance of component j as developed by supplier i, v_{ij}, depends on both the supplier's development effort e_{ij} as well as a technological shock ζ_{ij}. The optimal contest format is thus the one that best allows the buyer to incentivize significant development efforts while at the same time offering him a chance to hedge against harmful technological shocks. A system contest is particularly strong in the former attribute as the monetary stakes in this contest format are highest and hence effort incentives are largest. In contrast, a component contest allows the buyer to choose (and combine) *ex post* the two best components and thus provides an excellent hedge against high technological uncertainties as found in many innovation projects. How these different benefits influence the suppliers' development efforts in the different contest formats—and how these efforts change in σ and k—is the core of Lemma 3.1.

Lemma 3.1. *Suppose components are technological substitutes.*

(i) *For $k = 1$, the unique symmetric perfect Bayesian Nash equilibrium for both system and component contest satisfies the following properties:* (a) *The buyer's optimal award split in a component contest is $p^* = 1/2$;* (b) $e_1^{sys} = e_2^{sys} > e_1^{cpo} = e_2^{cpo} > 0$.

(ii) *For $k \to \infty$, the unique symmetric perfect Bayesian Nash equilibrium for both system and component contest satisfies the following properties:* (a) *The buyer's optimal award split in a component contest is $p^* = 1$;* (b) $e_1^{sys} = e_2^{sys} = 0$ *and* $e_1^{cpo} > e_2^{cpo} = 0$.

Lemma 3.1(i) is concerned with situations in which both components require an equaling level of innovation (i.e., $k = 1$); that is both components bear the same level of risk during component development. In these cases, both components have ex ante an equally strong effect on the performance of the integrated system, and this is why both the buyer (in terms of the award split) as well as the suppliers (in terms of effort) spread their bets equally across the two components.

To understand the intuition behind this result in more detail, consider initially the buyer's and the suppliers' decisions in a component contest. The buyer seeks to maximize the performance of the submitted components, and this is akin to maximizing the suppliers' development efforts across the two component contests. Since each supplier invests the more effort in a contest the higher the promised award is, the buyer has to carefully balance the distribution of awards between the two different component contests. But why is it now optimal to split the total award equally (i.e., $p^* = 1/2$)? Recall that each supplier's marginal return on effort is decreasing, whereas his marginal costs are increasing. It is thus most efficient for the buyer if the suppliers invest equal efforts in each of the two component contests. And the only way for the buyer to induce such a balanced distribution of efforts is to offer the same award in each contest. As a result, equilibrium efforts are also identical across component contests: $e_1^{cpo} = e_2^{cpo}$.

A similar reasoning applies to system contests as well. Yet this time it is not the buyer that is most interested in an efficient distribution of efforts across components (as the buyer only cares about the integrated system's final performance), but it is the suppliers that try to maximize the efficiency of their effort provision. As before, for each supplier the return on effort is diminishing whereas marginal costs are increasing, and thus each supplier splits his efforts evenly between components: $e_1^{sys} = e_2^{sys}$.

Finally, the result in Lemma 3.1(i) that has the strongest implications on the buyer's choice of contest format is the relation between equilibrium efforts in a system as compared to a component contest: When components require a similar magnitude of innovation (i.e., $k = 1$), then suppliers unequivocally invest higher efforts in a system contest. This monotonicity is noteworthy as a change from system to component contest introduces two diametrically opposite effects. On the one hand, the award that the suppliers compete for is twice as high in a system contest as in each of the individual component contests (A vs. $A/2$), and higher financial stakes clearly promote more effort provision. On the other hand, each supplier encounters more uncertainty in a system contest than in a component contest, and higher levels of risk deter the suppliers' to invest in development effort. Yet, Lemma 3.1(i) reveals that the first effect, which favors an increase in effort, always dominates the latter effect for components with a similar degree of innovativeness. Interestingly, Lemma 3.1(ii) reverses this finding when one component requires a significantly higher level of innovation than the other component (i.e., $k \to \infty$).

Specifically, for components with a sizeable difference in innovativeness the main benefit of a system contest—inducing higher development efforts—is entirely absent. This is happening because with a large k, the performance of the integrated system is primarily determined by the performance of the more innovative (and thus riskier) component, and for such a component it is the technological uncertainty that ultimately determines performance. In other words, development efforts play a subordinate role in determining the final performance of the integrated system and therefore suppliers refrain from incurring overly high effort costs.

To the contrary, in a component contest effort still matters. In particular, for the less innovative component the suppliers'—and also the buyer's—return on effort is substantially larger than for the more innovative component, and therefore suppliers fully dedicate their development efforts to the less innovative component. This explains why for $k \to \infty$, equilibrium development efforts are higher in a component than a system contest. We are now well equipped to discuss the buyer's optimal contest format choice.

Proposition 3.1. *Suppose components are technological substitutes. Then, the following profit relations hold in equilibrium:*

(i) *There exists a threshold* $\underline{\sigma}_{sub} > 0$ *such that if* $\sigma \in (\sigma_{sub}, \underline{\sigma}_{sub})$*, there exists a threshold* $\underline{k}_{sub} > 1$ *such that* $\Pi^{sys}_{sub} > \Pi^{cpo}_{sub}$ *if* $k \in [1, \underline{k}_{sub})$*. Moreover,* $\underline{\sigma}_{sub} \to 0$ *as* $n \to \infty$*.*

(ii) *For any fixed* $\sigma > \sigma_{sub}$ *there exists a threshold* $\overline{k}_{sub} < \infty$ *such that* $\Pi^{cpo}_{sub} > \Pi^{sys}_{sub}$ *if* $k \in (\overline{k}_{sub}, \infty]$*.*

(iii) *For any fixed* $k \geq 1$ *there exists a threshold* $\overline{\sigma}_{sub} < \infty$ *such that* $\Pi^{cpo}_{sub} > \Pi^{sys}_{sub}$ *if* $\sigma \in (\overline{\sigma}_{sub}, \infty]$*.*

Proposition 3.1 establishes—and the left panel of Figure 3.2 visualizes—the buyer's optimal contest format choice for a product with substitute components, highlighting the role that the magnitude of innovation involved plays when choosing the optimal contest format. A couple of managerial insights abound.

First and foremost, Proposition 3.1 shows that the buyer should use a system contest to procure the integrated system only if two conditions are simultaneously satisfied: (1) Both components should be incremental innovations only (i.e., $\sigma < \underline{\sigma}_{sub}$ and $k < \underline{k}_{sub}$), and (2) the supplier base should be relatively small (i.e., small n). If either of these two conditions is violated—that is, if at least one component requires radical innovation (i.e., high σ or high k) or the buyer's supplier base is large—the buyer should opt for a component contest and procure the components individually. The intuition for this result is instructive. For incrementally innovative components and in the presence of a small supplier base, it is a supplier's development efforts that have the biggest impact on the final product's performance. As such, the buyer chooses the contest format that offers the highest effort incentives, which is a system contest. However, Lemma 3.1 has shown that these effort incentives deteriorate quickly when at least one of the components requires more radical innovation. Under radical innovation—that is, when the effort performance link becomes highly tenuous—the possibility to ex post combine the best components—the key feature offered by a component contest—becomes more and more valuable, and therefore the buyer implements a component contest in these cases.

Second, Proposition 3.1 also reveals how the size of the supplier base impacts the buyer's contest format choice. Two forces are at work here. From a supplier's perspective, the presence of a larger number of competitors diminishes each supplier's chances of winning the contest, thereby inducing each supplier to invest less effort into component development. From the buyer's perspective, the probability that a single supplier develops the best two components drastically shrinks as n increases, and simultaneously, chances rise that for each of the components one of the suppliers will develop an exceptionally good component. Taken together, these implications of a larger supplier base lead the buyer to implement a component contest when n is sufficiently large, irrespective of the level of innovation involved in component development. As a side note, this finding also implies that the buyer should not always choose the contest format that promises the highest supplier efforts.

From a practical perspective, Proposition 3.1 gives clear advice. When the buyer has access to a large supplier base, which is oftentimes the case in crowdsourcing environments, then a separate contest should be dedicated to each required component. However, for very complex technological products like automotive lighting systems, buyers can typically tap into only a handful of suppliers, and in these situations, it is the level of innovation that ultimately determines whether a system or component contest should be used.

3.4.2. Technological Complements

In the previous section, we have established which contest format is optimal for a buyer wishing to procure a product whose components are technological substitutes. At this point, it is then still an open question whether the particular form of the technological relationship between the individual components affects our results—and thus our man-

Figure 3.2.: The Optimal Contest Format.

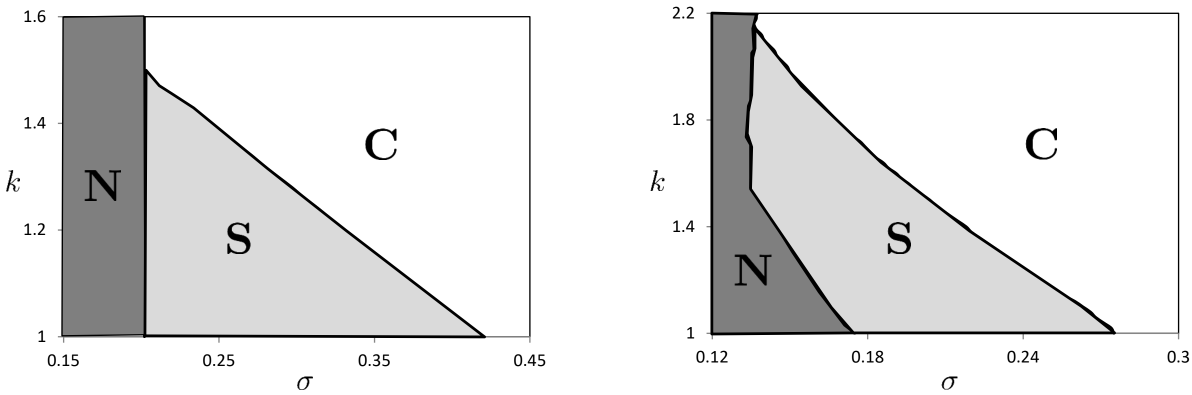

Notes: The buyer's optimal choice of contest format for technological substitutes (left panel) and technological complements (right panel). The buyer prefers a component contest in the white area (region C), and a system contest in the light gray area (region S). No pure strategy equilibrium exists in the dark gray area (region N). The parameters are: $r(x) = x$, $c(x) = x^2$, $n = 2$, $A = 1$.

agerial recommendations—or not. Answering this question is the primary purpose of this section.

As a starting point for our discussion, suppose that the buyer hosts a component contest—i.e., a separate contest for each component—and consider a supplier participating in at least one of these contests. In each contest, regardless of whether the supplier participates in both or only one contest, the supplier entirely focuses on developing the best component for this contest, without worrying about how this component may contribute to the overall performance of the integrated system. Put differently, the use of a component contest induces the suppliers to consider each component in isolation, and as a result, it is irrelevant for their effort choices whether the components are technological substitutes or complements. Clearly, this is not true for a system contest. The following Lemma compares the suppliers' equilibrium efforts between a system and a component contest for different levels of absolute and relative innovation.

Lemma 3.2. *Suppose components are technological complements.*

(i) *For $k = 1$, there exists a symmetric perfect Bayesian Nash equilibrium for both system and component contest that satisfies the following properties:* (a) *The buyer's optimal award split in a component contest is $p^* = 1/2$;* (b) $e_1^{sys} = e_2^{sys} > e_1^{cpo} = e_2^{cpo} > 0$.

(ii) *For $k \to \infty$, the unique symmetric perfect Bayesian Nash equilibrium for both system and component contest satisfies the following properties:* (a) *The buyer's optimal award split in a component contest is $p^* = 1$;* (b) $e_1^{cpo} > e_1^{sys} > e_2^{cpo} = e_2^{sys} = 0$.

Lemma 3.2 reveals that many of our observations for substitute components immediately transfer to the case of technological complements. To be concrete, for components that demand an equaling level of innovation (i.e., $k = 1$) it remains true that the buyer and the suppliers split their investments equally across both components, and that the suppliers always invest more development effort in a system contest. At the same time, however, Lemma 3.2(ii) indicates that when one of the components requires a significantly larger amount of innovation than the other component, then the technological relationship between components has an impact on supplier efforts. More specifically, even when the relative degree of innovation becomes very large (i.e., $k \to \infty$), each supplier continues to invest effort in the less innovative component (i.e., component 1) during a system contest—a finding that is at odds with our findings for substitutable components (cf. Lemma 3.1(ii)). Why is this happening? Recall that for a product with technologically complementary components the performance of the integrated system is the minimum of the two components' individual performances. In other words, the integrated system is only as good as its weakest component. Given that k is large, component 2 is associated with a high technological uncertainty, thus making it very likely that the development of this component results in an extreme success or a severe failure. In the latter case effort expended into the development of component 1 is futile, but in the former case the performance of component 1 is the limiting factor for the integrated system's performance. Consequently, each supplier sustains an effort investment into component 1 in order to improve the performance of the component that is likely to be the weakest link and to increase his chances of winning the contest.

Despite this difference between Lemmas 3.1(ii) and 3.2(ii), a key result from Section 3.4.1 remains valid irrespective of whether components are technological substitutes or complements: when the relative degree of innovation between components is sizeable, then the overall effort that a supplier invests into the development of both components is maximized under a component contest. Based on this observation, it is no longer surprising that structurally our results regarding the buyer's optimal contest format choice carry over from the case of technological substitutes to the case of technological complements. Proposition 3.2 and the right panel of Figure 3.2 make this argument explicit.

Proposition 3.2. *Suppose components are technological complements. Then, the following profit relations hold in equilibrium:*

(i) *There exists a threshold* $\underline{\sigma}_{cml} > 0$ *such that if* $\sigma \in (\sigma_{cml}, \underline{\sigma}_{cml})$*, there exists a threshold* $\underline{k}_{cml} > 1$ *such that* $\Pi^{sys}_{cml} > \Pi^{cpo}_{cml}$ *if* $k \in [1, \underline{k}_{cml})$*. Moreover,* $\underline{\sigma}_{cml} \to 0$ *as* $n \to \infty$*.*

(ii) *For any fixed* $\sigma > \sigma_{cml}$ *there exists a threshold* $\overline{k}_{cml} < \infty$ *such that* $\Pi^{cpo}_{cml} > \Pi^{sys}_{cml}$ *if* $k \in (\overline{k}_{cml}, \infty]$*.*

(iii) *For any fixed* $k \geq 1$ *there exists a threshold* $\overline{\sigma}_{cml} < \infty$ *such that* $\Pi^{cpo}_{cml} > \Pi^{sys}_{cml}$ *if* $\sigma \in (\overline{\sigma}_{cml}, \infty]$*.*

As in the case of technological substitutes, the buyer uses a system contest when both components require only incremental innovation and the available supplier base is limited in size. In contrast, for products that require more radical innovation or in situations with many potential suppliers the buyer seeks to implement a component contest. In the latter situation, the buyer sacrifices the suppliers' solution efforts in favor of having the option to ex post select the best components out of a large set of alternatives. This finding is reminiscent of traditional parallel path results as described in, e.g., Dahan and Mendelson (2001) and Boudreau et al. (2011), but transfers them to contests involving multiple interacting products.

From a managerial perspective, the most important finding of Proposition 3.2 is that the buyer's optimal choice of contest format does not depend on the technological relationship between components. Specifically, independent of whether an integrated system consists of technological substitutes or complements, the answer to the question of which contest format to select depends solely on the components' absolute and relative level of innovation and the size of the available supplier base.

3.5 Performance Correlation

To clearly elicit the key tradeoffs immanent to a buyer's contest format choice, we have so far ignored the effects of performance correlation across components on both, the suppliers' solution efforts and the buyer's choice of contest format. However, given that performance correlation between different components developed by the same supplier is an integral aspect of many product development initiatives in reality, we now set out to study the implications of performance correlation as represented by the parameter ρ.

For a supplier, performance correlation across development projects for different components may arise naturally for manifold reasons. For instance, both development projects might be overseen by the same project manager who imposes his "beauty ideals" on both components; or the supplier might simply use similar technological approaches to tackle the different development projects. And whereas each supplier can proactively influence his own level of performance correlation, the buyer has no lever to directly control the amount of correlation exhibited by each individual supplier. In other words, from the buyer's perspective, performance correlation is a given trait inherent to each supplier's development processes. Yet this does not imply that the buyer cannot influence the amount of performance correlation present in his procurement contest. To the contrary, he can do so by, e.g, selecting which suppliers are invited to participate in which contest. For instance, in a component contest, the buyer could decide to allow each individual supplier to participate in only exactly one component contest, thereby effectively breaking correlation between the two component contests. Alternatively, he could oblige each supplier to participate in both component contests, thus introducing a higher level of performance correlation across the two contests. Similarly, in a system contest, the buyer could choose to invite only suppliers with a very focussed or a more diverse technology base, thereby directly influencing the latent level of performance correlation in his procurement mechanism.

Many options to influence the level of performance correlation in the procurement contest abound, but just what amount of performance correlation would be optimal for the buyer? And how is the buyer's choice of contest format affected when system and component contests exhibit different levels of performance correlation? Answering these questions is the purpose of this section.

3.5.1. Performance Correlation in Component Contest

When hosting a component contest, does the buyer prefer more or less performance correlation between the two different components? The following Proposition highlights that the answer to this question critically depends on the technological relationship between the components; that is, whether components are technological substitutes or complements.

Proposition 3.3. *In equilibrium, the buyer's expected profit from hosting a component contest satisfies the following sensitivities:*

(i) *Suppose components are technological substitutes. Then,* Π_{sub}^{cpo} *is invariant with respect to* ρ.

(ii) *Suppose components are technological complements. Then,* Π_{cml}^{cpo} *increases in* ρ.

The first noteworthy result of Proposition 3.3 is that in a component contest involving an integrated system with substitutable components performance correlation has no effect on the buyer's profits. To see why, recall that in a component contest, each supplier considers each of the two contests in isolation, and therefore his development efforts are not affected by performance correlation. Also, the substitutable nature of the components precludes the buyer from exploiting the benefits of correlation on the overall level of technological uncertainty. Hence, since neither the suppliers' efforts nor the overall level of technological uncertainty are affected by performance correlation in the case of technological substitutes, the buyer's profits are invariant in ρ. This is not true for technological complements (Proposition 3.3(ii)). As for technological substitutes supplier development efforts remain unaffected by performance correlation, but this time the buyer can exploit the effect of correlation on the overall level of technological uncertainty in the procurement mechanism. Specifically, the higher the performance correlation between the individual components, the more similar are their performances. And since the integrated system is only as good as its weakest component, the buyer benefits from higher levels of performance correlation.

Proposition 3.3 also has immediate implications for practice. A buyer always (weakly) benefits from performance correlation in a component contest—and this finding is independent of whether components are incremental or radical innovations, the technological relationship between components, and the size of the supplier base. But most importantly, a buyer can take simple, yet effective measures to drive up performance correlation in his procurement contest. For instance, the buyer should always oblige suppliers to participate in both contests, instead of limiting access to only one contest. Similarly, the buyer should preferably invite only such suppliers that are known to develop components with well-balanced performances, and forgo inviting suppliers that are overly specialized on only one component.

3.5.2. Performance Correlation in System Contest

In the preceding section, we have established that for a component contest the effect of performance correlation on the buyer's profits crucially depends on the technological relationship between components. Rather surprisingly, it is not the technological relationship that determines the sensitivity of the buyer's profits in a system contest; instead it is the magnitude of innovation required to develop the components as well as the size of the supplier base that plays a vital role.

Proposition 3.4. (i) *Suppose components are technological substitutes. Then,* Π_{sub}^{sys} *increases in* ρ *if and only if* $n\sigma^2(1+k^2+2k\rho) \geq 2A\eta'(A\mu^{(n)}/(n\sigma\sqrt{1+k^2+2k\rho}))$.

(ii) *Suppose components are technological complements. Then,* Π_{cml}^{sys} *decreases in* ρ *if* k *and* σ *are sufficiently small, and increases if* k *or* σ *are sufficiently large.*

The main finding of Proposition 3.4 is as follows. For products that are made of two incrementally innovative components (i.e., low σ and k) and for which the supplier base is narrow (i.e., small n) the buyer prefers a low degree of performance correlation, whereas in any other situation—that is, if at least one component is a radical innovation or the supplier base is large—the buyer wishes to have substantial levels of performance correlation in his procurement contest. To better understand this finding, recapitulate that from the buyer's perspective, the success of the procurement contest hinges on two factors: the suppliers' development efforts and the performance uncertainty. More specifically, in contests involving low levels of innovation and a small supplier base it is particularly the supplier effort that drives the performance of the final product. In contrast, in contests involving radical innovation or a large supplier base the final product performance is predominantly determined by the inherent technological uncertainty. Clearly, in the former case the buyer prefers conditions that lead suppliers to engage in high development efforts, whereas in the latter case the buyer prefers high levels of uncertainty—and the higher the performance correlation is, the higher is the overall technological uncertainty and the lower are the suppliers' effort incentives.

Proposition 3.4 provides clear advice for managers overseeing the procurement of a complex innovation that is sourced through a system contest. Whether performance correlation is beneficial for the procurement mechanism, or not, depends on the magnitude of innovation required and the size of the supplier base, but not on the technological relationship between components. And the buyer can steer the degree of performance correlation in his procurement contest by inviting only those suppliers that offer the right level of performance correlation; that is, suppliers that employ the right technologies and that have an appropriate organizational structure.

3.5.3. Heterogeneity in Performance Correlation

We conclude our study of performance correlation by investigating how structural differences in correlation between the different contest formats may influence the buyer's optimal contest format choice. To be more concrete, we would expect that in a system contest—where each supplier develops the full product instead of only an isolated component—performance correlation is at least as high as in a component contest (for a given set of suppliers). To include this possibility into our model setup, we let ρ^{sys} and ρ^{cpo} be the performance correlation in a system and component contest, respectively, and we now examine how the buyer's choice of contest format changes with $\Delta\rho = \rho^{sys} - \rho^{cpo} \geq 0$.

Proposition 3.5. (i) *Suppose components are technological substitutes. Then the same preference ordering between contest formats as given in Proposition 3.1 applies for any* $\Delta\rho \geq 0$.

(ii) *Suppose components are technological complements. Then: (a) For small* k *and small* σ, $\Pi^{sys}_{cml} > \Pi^{cpo}_{cml}$ *for any* $\Delta\rho \geq 0$. *(b) There exists a threshold* $\sigma_\rho > 0$ *such that for all* $\sigma \in (\sigma_\rho, \infty]$ *there exists thresholds* $\Delta\underline{\rho}$ *and* $\Delta\overline{\rho}$, *with* $0 < \Delta\underline{\rho} < \Delta\overline{\rho} < 1$, *such that* $\Pi^{cpo}_{cml} > \Pi^{sys}_{cml}$ *for any* $\Delta\rho < \Delta\underline{\rho}$ *and* $\Pi^{sys}_{cml} > \Pi^{cpo}_{cml}$ *for any* $\Delta\rho > \Delta\overline{\rho}$.

The results presented in Proposition 3.5 emphasize that our initial results regarding the buyer's optimal contest format choice—as presented in Propositions 3.1 and 3.2—remain valid even when the different contest formats exhibit varying levels of performance correlation. In particular, whereas a system contest is the buyer's optimal choice whenever both components are incremental innovations, the component contest becomes the buyer's preferred contest format for more innovative components. However, part (b) of Proposition 3.5(ii) shows that there is one notable exception to this general rule: When the heterogeneity in performance correlation between system and component contest becomes sizeable (i.e., $\Delta\rho > \Delta\overline{\rho}$), then the buyer prefers a system contest even when both components are radical innovations (i.e., large σ). But why is this happening? Recall that for radical innovations, supplier efforts have only a minor influence on product performance; instead it is the technological uncertainty that predominately determines the performance of the integrated system. And with increasing $\Delta\rho$, relative more correlation—and thus also more uncertainty—is introduced into the system contest, whereas the component contest exhibits relatively less correlation—and thus less uncertainty. Eventually, there exists a threshold $\Delta\overline{\rho}$ beyond which the system contest format displays a higher degree of technological uncertainty than a component contest, and this is why the system contest becomes the optimal choice for the buyer in the face of radically innovative components.

3.6 A Heterogeneous Supplier Base

Until now we have assumed a homogeneous supplier base in the sense that each supplier $i \in \{1, \dots, n\}$ is endowed with an identical effort-performance tradeoff (i.e., identical r and c across suppliers) and that all suppliers equally value the award A promised by the buyer. Clearly, such a symmetry assumption is a simplification of reality, and to study the impact of a heterogeneous supplier base on the buyer's contest format choice, we now relax the latter assumption of equal award valuations. In reality, each supplier's perceived value of the award A depends on the exact terms of the final procurement contract, and different suppliers may be more or less effective in adhering to these terms, which naturally creates asymmetries in their valuation of the contract.

To simplify the exposition we restrict our attention to a setting with two suppliers, and we assume that supplier 1 is more effective in fulfilling the supply contract than supplier 2. More precisely, supplier 1 receives a utility of A from winning an award of size A, whereas supplier 2 only gains αA, with $\alpha \in (0, 1]$. As a result of these heterogeneous award valuations, the two suppliers invest different amounts of effort during the procurement contest: Supplier 1, who has a higher valuation for winning the contest, always invests more development effort than supplier 2, but both suppliers reduce their efforts as α decreases. In other words, supplier asymmetry has a negative impact on effort provision. But how does this negative effect influence the buyer's choice of contest format?

As Figure 3.3 visualizes, heterogeneity in the supplier base only has a minor impact on the buyer's preferences across contest formats. In particular, independent of the magnitude of supplier asymmetry, the buyer implements a system contest if both

Figure 3.3.: The Optimal Contest Format with a Heterogeneous Supplier Base.

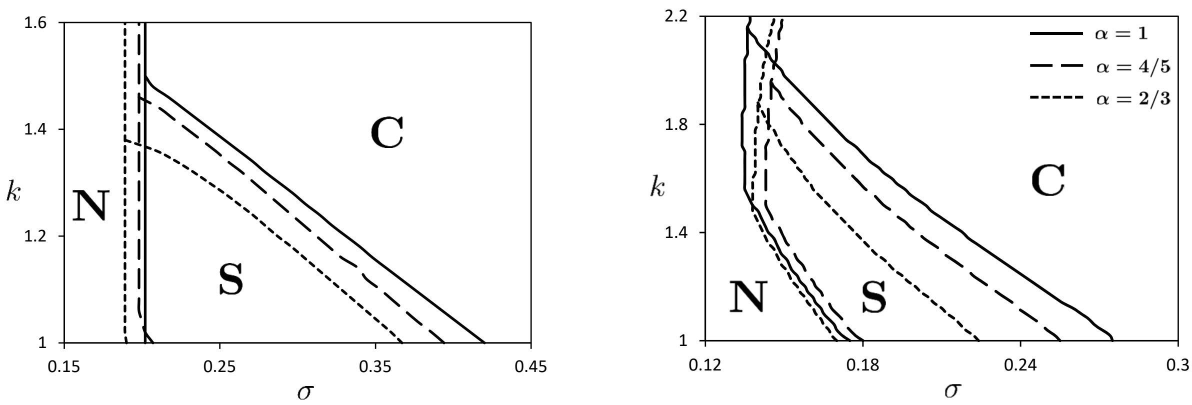

Notes: The buyer's optimal choice of contest format for technological substitutes (left panel) and technological complements (right panel) for different degrees of heterogeneity in award valuations α. The buyer prefers a component contest in region C, and a system contest in region S. No pure strategy equilibrium exists in region N. The parameters are: $r(x) = x$, $c(x) = x^2$, $n = 2$, $A = 1$.

components are incremental innovations, and chooses a component contest otherwise. We conclude that our key structural results—as presented in Propositions 3.1, 3.2 and 3.5—seem to be fully robust to the introduction of heterogeneity into the supplier base. As an aside, our results also indicate that a buyer should always strive for a symmetric and well-balanced supplier base.

3.7 Conclusions

In practice, procurement contests have become a popular tool for buyers to gain access to innovative products developed by their supplier base. Yet, tapping into the suppliers' innovation efforts is a challenging endeavor. In particular, to fully exploit the benefits of hosting a procurement contest, buyers need to be able to tailor their contests to the characteristics of the products to be acquired and to their supplier base. This is particularly true so if the desired innovation is a technologically complex product such as an automotive lighting system. For such complex products, the most prominent question for the buyer is whether to procure the entire system from a single supplier, or whether to source the product's individual components from different suppliers. This question is at the very heart of our research, and clearly, the answer to it also has strong implications on how the buyer should conceptualize his procurement contest.

First, we find that a buyer should use a system contest whenever all product components are merely incremental innovations and the firm's supplier base is relatively small. In contrast, in the presence of radical innovation or a large supplier base, the buyer should choose to implement a component contest. Put differently, when deciding on the optimal contest format managers only have to consider two dimensions: (i) the magnitude of innovation required for each component; and (ii) the size of the supplier base. Relying only on these two factors significantly reduces the complexity of the managerial decision problem in practice. Both the required level of innovation and the number of potential suppliers are well-known to the firm, and at the same time managers need not worry about more subtle—and less tangible—dimensions such as the exact technological relationship between components or the degree of heterogeneity and correlation within the supplier base.

Second, our results give clear advice on how a buyer should conceptualize his supplier base in the long term, and which suppliers he should invite to the different procurement contests. As a general managerial guideline, the buyer should always strive for a relatively homogeneous supplier base; strong differences in the suppliers' capabilities have a severely negative effect on each supplier's incentives to exert development efforts. In contrast, whether the buyer prefers higher or lower levels of performance correlation in his supplier base depends on the contest format as well as the innovativeness of the product. It is interesting to observe that there are again only two important dimensions that a firm has to consider, and both of them are again relatively easy to observe. Our results therefore encourage managers to actively engage in the management of their supplier base when it comes to complex innovations, because this management problem might be less complex than perceived at first glance.

Our model has limitations that should be mentioned. To develop a parsimonious model and maintain mathematical tractability we have made some simplifying assumptions regarding the characteristics of the considered products. In particular, we restrict the number of components in the integrated system to be two, thereby disregarding more complex or more hierarchical product architectures. We believe that incorporating hierarchical product architectures may lead to important and interesting additional insights, and we see this as a promising avenue for future research. Also, we have considered only the cases of perfectly substitutable and perfectly complementary components, leaving aside the large class of products that exhibit intermediary forms of technological relationship. Lastly, we have assumed that the buyer treats all suppliers equally. In reality, however, the buyer may strategically handicap certain suppliers to improve contest outcomes and to influence the evolution of his supplier base.

Chapter IV

Managing Delegated R&D Testing

with Gerrit Schumacher[1]

4.1 Introduction

Scholars have long acknowledged the crucial role that design testing plays in the success of any research and development (R&D) or new product development (NPD) initiative (Simon 1969, Allen 1977, Clark and Fujimoto 1989, Wheelwright and Clark 1992, Thomke 1998). This view is confirmed by the significant amount of time and resources that firms invest in activities related to testing their new products. Thus the automotive industry spent more than $100 billion (US) on R&D activities in 2015, with a large portion of this budget dedicated to design testing (Jaruzelski and Hirsh 2016); Airbus spent more than seven years evaluating different design options for its next-generation A380 aircraft before deciding on the final design (*The Economist* 2007); and the high-tech sector is expected to invest 40% of its information technology budget in new testing processes—such as "virtual" testing and robotics—by 2019 (Buenen and Muthukrishnan 2016). These numbers indicate that firms continuously seek to improve their testing processes, from both a technological and a managerial perspective, as a way of reducing costs and resource consumption yet without compromising the quality of their testing efforts.

Prior academic work has identified many aspects that bear on the efficacy of a firm's design-testing process—including such diverse factors as test efficiency, testing costs, lead times, and learning effects—and thus on the firm's optimal testing strategy (see, e.g., Weitzman 1979, Dahan and Mendelson 2001, Loch et al. 2001, Erat and Kavadias

[1]The research presented in this chapter is based on a paper entitled "Delegated Testing of Design Alternatives: The Role of Incentives and Testing Strategy", coauthored with Gerrit Schumacher.

2008). So far, however, one critical aspect has received only scant attention in the literature: *design* testing is often delegated to self-interested experts who may pursue their own respective agendas. It is an organizational reality in many industries that information asymmetry between senior management and these testing experts distorts the outcomes of a testing process (Sommer and Loch 2009, Mihm 2010, Schlapp et al. 2015). The goals of this paper are to determine (i) precisely how information asymmetry affects a firm's testing process and (ii) how the firm can mitigate the negative effects of associated agency issues by devising appropriate incentive structures and an adequate testing strategy.

One can more clearly understand the effect of information asymmetry on a firm's testing process by considering a wind turbine manufacturer that seeks to set up a new "wind farm"—a grid-connected installation of multiple wind turbines—in a predetermined location.[2] Wind farms are built to convert the wind's kinetic energy into electricity (Krohn et al. 2009). A wind farm is a viable (i.e., an economically rational) contender in the production of electricity only if it satisfies three basic requirements: "(1) produce energy, (2) survive, and (3) be cost effective" (Manwell et al. 2009, p. 505).

In fact, the "produce energy" requirement has become a moot point. Wind turbine manufacturers can now produce a variety of wind turbine designs that have proven their technological effectiveness through standardized testing procedures and widespread application in practice. Moreover, technical developments are pushing modern wind turbines closer to the theoretical efficiency limits dictated by Betz's law[3] (Burton et al. 2011, p. 63); hence current wind turbine designs are an excellent choice also for future wind farms. Yet one crucial question remains: Which turbine design is the best choice for a given wind farm location?

The answer to this question is closely tied to evaluating Manwell et al.'s requirements (2) and (3). A wind farm can be cost effective and long-lived only if the wind turbines used are technically reliable and do not result in strong negative externalities on the environment. In other words: wind farm builders are looking for the wind turbine designs that best match site-specific wind conditions, climatic factors, and regulatory constraints yet have minimal effects on animal wildlife, emit little noise, and do not generate severe

[2]The siting of wind farms usually proceeds in close collaboration with regulatory bodies. Therefore, wind turbine manufacturers can influence but not ultimately control decisions about where new wind farms will be located.

[3]According to Betz's law, no turbine can capture more than 59.3% of the wind's kinetic energy.

electromagnetic interference (Manwell et al. 2009, pp. 321ff). Unfortunately, many of these factors are known only imperfectly ex ante, and there are no standardized testing procedures for evaluating them (in contrast to the purely mechanical testing of the energy produced by a particular wind turbine design). Instead, wind farm builders must rely on teams of experts that acquire and deliver information regarding the suitability of different wind turbine designs for a given location. Since there is no standard testing procedure, this information acquisition process—and the interpretation of the acquired information—is a process that creates *tacit* knowledge and that relies strongly on the experts' prior experience and knowledge, the quality of information sources, the synthesis of implicit information, and often also on gut feeling.

As a result, it is almost impossible for a wind farm builder to assess the quality of the information on which expert recommendations are based, let alone to verify that the experts have actually shared their knowledge (and all acquired information) with senior management. To overcome this information asymmetry, wind farm builders must adequately incentivize those experts to investigate design suitability in a thorough manner and to share the test outcomes with senior management in a truthful manner. These incentives must, of course, be aligned with the firm's overall testing strategy (e.g., parallel vs. sequential testing). Hence the questions that arise are: What is the firm's optimal testing strategy, and what are the corresponding optimal incentive structures? Our paper's main contribution is development of a game-theoretic model that delivers answers to these questions.

In particular, this study makes three main contributions to the literature. First we show that—almost regardless of the firm's testing strategy—the optimal compensation schemes that adequately incentivize the experts have a surprisingly simple two-payment structure: a success bonus; and a consolation award if an expert's design is not chosen for development. The optimal balance between these two payments depends on the informational quality of an expert's test outcomes. For rather simple designs that can be tested with high precision, the firm should place a strong emphasis on the success bonus. Yet designs that are more complex and subject to a less precise testing process demand more tolerance for failure, so in these cases the firm should offer higher consolation awards. In short, a one-size-fits-all approach to incentives is not advisable.

Second, we find that the firm's optimal testing strategy depends primarily on two parameters: the testing costs and the test efficiency. For low testing costs, the firm prefers a parallel testing strategy whereby all design alternatives are tested simultaneously; for higher testing costs, a sequential testing approach is the firm's preferred choice. We also address the question of how best to conceptualize a sequential strategy under delegation; that is, in which order should the designs be tested, and how many experts should be hired for the testing process? Our analysis reveals that the greater the heterogeneity in test efficiency across the design alternatives, the more experts the firm should hire. The reason is that employing many experts makes it more difficult for any one expert to extract high information rents—an issue that is most salient when the quality of testing is unbalanced across designs. With regard to the optimal order in which design alternatives should be tested, we challenge results in the literature that argue for the optimality of "reservation price rules" (Weitzman 1979, Adam 2001, Erat and Kavadias 2008). We establish that, when there is information asymmetry, it might be better for the firm to test less promising designs *first* in order to reduce agency costs.

Our third main contribution is to show that information asymmetry always results in a suboptimal testing process; of perhaps even more importance is our finding that the negative effect of information asymmetry is greater on parallel than on sequential testing strategies. These results indicate that, under delegation, a parallel testing approach is less suitable than promised by extant research (Dahan and Mendelson 2001, Loch et al. 2001). This finding likely also explains why so few parallel testing efforts are observed in practice despite sharply reduced testing costs in recent years: firms simply want to avoid the high agency costs associated with parallel design testing.

4.2 Related Literature

The challenges associated with managing the design process of a novel product have been a long-standing and central concern in the NPD literature (Simon 1969, Allen 1977, Clark and Fujimoto 1989, Wheelwright and Clark 1992, Thomke 1998, Loch et al. 2001, Pich et al. 2002, Erat and Kavadias 2008, Sommer et al. 2009). In his foundational work, Simon (1969, pp. 128f) describes the product design process "as involving, first, the generation of alternatives and, then, the testing of these alternatives". This view has served as the foundation of much of the subsequent academic literature, and as such it has triggered numerous extensions (see, e.g., Clark and Fujimoto 1991, Wheelwright and Clark 1992, Thomke 1998, 2003). Following the seminal classification of Simon (1969), the extant literature can be divided into two broad categories. The first group of studies focuses on the *search* dimension of product design by investigating successful strategies for finding design alternatives. In contrast, research in the second group emphasizes the *testing* dimension of product development by analyzing how best to evaluate the performance of a given design alternative.

The literature on optimal *search* dates back to the pioneering work of March and Simon (1958) and Simon (1969), who were among the first to describe organizational problem solving as a search process. This notion of viewing the innovation process as a search over a complex design landscape inspired the subsequent proposal of different conceptual models to describe the underlying search spaces. The two most influential models of search spaces are the exploration–exploitation trade-off described in March (1991) and Manso (2011) and the NK landscape of Kauffman and Weinberger (1989) and Kauffman (1993). Building on these conceptualizations of a search space, scholars have extensively investigated how the efficiency of the search process—and thus the firm's innovation performance—changes with the complexity of the problem (Ethiraj and Levinthal 2004, Mihm et al. 2003, Billinger et al. 2014), organizational hierarchy (Rivkin and Siggelkow 2003, Siggelkow and Rivkin 2005, Mihm et al. 2010), team structure (Kavadias and Sommer 2009, Girotra et al. 2010), unforeseeable uncertainties (Sommer and Loch 2009, Sommer et al. 2009), the particular search strategy employed (Sommer and Loch 2004, Kornish and Ulrich 2011), and competition (Oraiopoulos and Kavadias 2014). More recently, Erat and Krishnan (2012), Lewis (2012), and Ulbricht (2016) have analyzed how delegation affects both the breadth and overall performance of a search process. All these cited papers focus on how to discover a set of potentially promising design alternatives, which is the quintessential first step in an innovation process. However, we are concerned with the second step in that process: determining the most reliable way to select the best alternative from among the candidates. As a consequence of that different focus, the formal model we propose differs considerably from those in the search literature.

Much closer to our work is the literature on design *testing* as initiated by Weitzman (1979). In his terminology, any design alternative can be considered a "black box", and uncertainty about its value can be resolved only by costly testing activities. This generic model of a testing process has become a building block for almost all research on design testing, and it has proven itself flexible enough to accommodate two very different kinds of testing processes: feasibility testing and selection testing.

The primary purpose of *feasibility testing* is to discover whether (or not) a given design is technologically feasible. Answering this question requires that the design in question be repeatedly tested until there is sufficiently strong evidence either for or against its feasibility. Prime application areas for this testing procedure include the pharmaceutical industry, where a new molecule is tested and retested during clinical trials to evaluate whether it reliably produces the desired effects. Research on feasibility testing seeks to answer such questions as when in the development process to test the design as well as how many tests to pursue and at what level of fidelity to a real-world counterpart (see, e.g., Thomke and Bell 2001, Terwiesch and Loch 2004). In the economics literature, the question of how to motivate an expert's participation in this dynamic information acquisition process has recently gained traction (Gromb and Martimort 2007, Gerardi and Maestri 2012, Hörner and Samuelson 2013), leading to a theory of optimal incentives for feasibility testing. There are two ways in which our paper is connected to that stream of literature. First, from previous work on feasibility testing we borrow the insight that nearly all testing processes are imperfect and so, even with the most thorough of testing efforts, there will still be uncertainty about a design's true value. Second, we answer the question of how the firm should manage delegated testing activities when it is concerned with selection testing—the second main challenge in design testing. Thus our work complements the extant literature and broadens our knowledge about devising optimal incentives and strategies for delegated design testing.

In contrast to feasibility testing, which focuses on the technological feasibility of a single design, *selection testing* is concerned with choosing the best alternative out of a set of different candidate designs. For instance, as explained in the Introduction, wind farm builders are confronted with such a selection issue when choosing a particular wind turbine design for a new wind farm. The existing literature describes two diametrically opposed testing strategies to tackle this issue: sequential and parallel testing. Weitzman (1979) advocates the use of a sequential testing approach, in which the different designs are tested in sequence and the testing process can be stopped after each design test. In his seminal contribution, he establishes the now classical reservation price rule (a.k.a. "Pandora's rule") for determining both the order in which to test the alternatives and also when to stop testing. Adam (2001) and Erat and Kavadias (2008) study how these results are affected by the firm's ability to learn between design tests. In contrast to those papers, Dahan and Mendelson (2001) promote the use of a parallel testing approach in which all design alternatives are tested simultaneously. Loch et al. (2001)

build on these results by directly comparing the performance of sequential and parallel testing strategies. They find that sequential testing is preferable when design tests are expensive and test efficiency is low whereas parallel testing is preferable when testing processes are slow and the firm needs quick results.

It is remarkable that past work on selection testing has not considered the role played by information asymmetry in the design of an optimal testing process—that is, given the ubiquity of such asymmetries in practice. Our principal contribution is to investigate how, exactly, delegation affects a firm's optimal testing strategy; we derive simple yet optimal incentive structures that counter the effects of information asymmetries. The analysis yields several new insights regarding the management of delegated testing processes. First, our derivation of the firm's optimal testing strategy reveals that information asymmetry is much more detrimental to a parallel than to a sequential testing strategy. This finding implies that, when testing is delegated, parallel strategies are probably less effective than advertised (e.g., Dahan and Mendelson 2001, Loch et al. 2001). Second, the existing literature is silent about how to conceptualize a sequential testing approach; that is, should the firm hire multiple experts to test the different design alternatives, or should it rather assign all testing activities to a single expert? We show that the multi-expert approach is preferable when test efficiency varies considerably across the design alternatives whereas the single-expert approach is preferable when test efficiency is relatively homogeneous. As a corollary we also find that, under delegation, the classical reservation price rules (promulgated in Weitzman 1979, Adam 2001, Erat and Kavadias 2008) no longer generate the optimal order for testing alternative designs. Finally, we derive the optimal incentive structures for delegated design testing, which turn out to be extremely simple irrespective of the chosen testing strategy.

4.3 Model Setup

Consider a firm engaged in NPD; it faces the challenge of selecting one out of $N \geq 1$ possible design alternatives for the new product being developed. The value that the firm receives from choosing a certain design depends on two factors: the design's technological feasibility, which is uncertain at the outset and can end up being either good or bad; and the design's inherent economic potential. The firm's goal is to choose and develop the design alternative that offers the highest value upon implementation.

At the core of our model is a testing phase in which the firm can acquire costly but imperfect information regarding each design's technological feasibility and, eventually, its value. Yet the firm cannot access this information directly. Instead, it must delegate the desired testing activities to experts who then collect information about the designs' technological feasibility through experiments, simulations, and/or prototype building. Each design test gives the corresponding expert more refined information about the evaluated design's feasibility, which enables that expert to provide a more informed recommendation vis-à-vis the firm's decision on whether or not to develop the focal design. Once the firm has collected enough recommendations, it develops the most promising design alternative; if no alternative is sufficiently convincing, then the development process may be abandoned.

In the real world, the delegated nature of this testing process gives rise to information asymmetry between the firm and the experts. Thus it is difficult if not impossible for the firm to verify the informational quality of an expert's recommendation. Two factors are responsible for this adverse situation. First, testing activities are costly for the expert, who may therefore choose to be less diligent with respect to some tests than others. However, firms cannot ascertain the diligence of experts because the interpretation of testing outcomes relies critically on an expert's gathering and synthesis of information and—most notably, owing to the tacit nature of such knowledge—prior experience. Second, experts may be reluctant to share their testing outcomes truthfully with the firm. This form of information asymmetry captures the reality that experts can use their recommendations strategically to influence the firm's design choice. So in practice, delegated design testing involves two different forms of information asymmetry: *moral hazard* during the testing phase and *adverse selection* during the recommendation phase. It follows that the firm must offer an appropriately designed compensation scheme if it hopes to incentivize experts to test the focal design(s) thoroughly and then to communicate the testing outcomes truthfully. This compensation scheme, in turn, should be carefully coordinated with the firm's second major decision: the choice of testing strategy. In particular, the firm must decide about whether the design alternatives should be tested in parallel or in sequence, how many experts to employ, how many design alternatives to test, and (if tested in sequence) the best order in which to test the different designs. In the rest of this section we provide more detail on our model setup and assumptions.

4.3.1. Delegated Design Testing

The value V_i that the firm receives from developing design $i \in \mathcal{N} = \{1, \ldots, N\}$ depends on the design's technological feasibility $\Theta_i \in \{G, B\}$, its inherent economic potential $v_i > 0$, and the development costs $K \geq 0$. We assume more specifically that, once developed, a technologically feasible design ($\Theta_i = G$) generates a value of $v_i - K$ whereas an unfeasible design ($\Theta_i = B$) results in a loss of $-K$. Prior to development, however, each design's technological feasibility is uncertain; that is, neither the firm nor the experts know the design's true feasibility ex ante. To simplify the presentation, we assume also that each state is ex ante equally likely.[4] Experts engage in costly testing activities in their efforts to (partially) resolve this uncertainty. We represent the testing activities for design i by an expert's testing effort e_i, which can be either high ($e_i = h$) or low ($e_i = l$) and is not observable by the firm. An expert who engages in high-effort testing incurs a private cost $c > 0$, whereas the costs of low-effort testing are normalized to zero. Of course, the chosen testing effort affects the quality of collected information. Formally, we model testing outcomes for each design i as an imperfect signal $s_i \in \{g, b\}$, which is received only by the expert testing design i and that indicates whether the design is technologically feasible ($s_i = g$) or not ($s_i = b$). We denote the precision (or quality) of this signal $q(e_i)$ because it depends on the expert's testing effort: an expert who exerts high effort receives a signal of quality $q(e_i = h) = q_i \in (1/2, 1]$; in contrast, low effort leads to an uninformative signal $q(e_i = l) = 1/2$. We assume that signals are stochastically independent across designs.

After receiving the signal s_i, the expert updates—in accordance with Bayesian rationality—her belief about design i's technological feasibility. Using this refined information, the expert gives the firm an unverifiable recommendation $r_i \in \{g, b\}$, which states whether design i is considered to be technologically feasible ($r_i = g$) or unfeasible ($r_i = b$).

4.3.2. The Firm's Decisions

The ultimate goal of a firm is to maximize expected profits by selecting the most valuable design for its product while minimizing the costs of a delegated testing process. For this problem to be relevant, we assume that $q_i v_i \geq K \geq (1 - q_i) v_i$ for all $i \in \mathcal{N}$. Otherwise,

[4] The sole purpose of this assumption is to reduce the complexity of our mathematical exposition. From a structural standpoint, all our results continue to hold for arbitrary prior probabilities.

the firm's decision would be a trivial one: if $q_i v_i < K$ then the firm could safely exclude design i from consideration because of the economic irrationality of developing such a design; at the other extreme, if $(1 - q_i)v_i > K$ then design i is so promising that the firm would develop it even without prior testing.

Designing an optimal testing process requires that the firm determine the testing strategy and also a scheme for compensating the experts. With regard to the former decision, we assume that the firm can select among three different testing strategies: parallel testing, multi-expert sequential testing, and single-expert sequential testing. Under a *parallel* testing strategy, the firm first decides on the number $|\mathcal{I}_P|$ and identity $\mathcal{I}_P \subseteq \mathcal{N}$ of design alternatives to test; it then assigns a separate expert to each design $i \in \mathcal{I}_P$, and all testing processes are carried out simultaneously. After reviewing the experts' recommendations, the firm decides which design (if any) to develop. Under a *multi-expert sequential* testing strategy, the firm again decides on the number (here, $|\mathcal{I}_M|$) and identity ($\mathcal{I}_M \subseteq \mathcal{N}$) of design alternatives to test—but it also determines the order in which the different designs will be tested. The firm then assigns a different expert to each design $i \in \mathcal{I}_M$ and the designs are tested, one after the other, in the order specified. After each design test, the firm can choose to stop the testing process and develop the latest design alternative; we assume that the firm always does so after receiving a good/ feasible recommendation for the current design alternative. An intuitive consequence of this assumption is that, once a firm in this position continues with the testing process, it can no longer implement any previously tested design.[5] Throughout the modeled testing process, all payments are discounted at a constant rate of $\delta \in (0, 1]$ after each design test. Under a *single-expert sequential* testing strategy, the firm makes the same decisions as in the multi-expert case (viz., deciding on $|\mathcal{I}_S|$, on $\mathcal{I}_S \subseteq \mathcal{N}$, and on the order of tested designs); the only difference here is that just one expert is assigned to perform all the tests. As before, the firm can stop the testing process after each design test and develop the latest design alternative (which always occurs if the firm is given a good recommendation) yet does not have the option of developing any formerly tested design.

[5] As discussed in Sections 4.4.2 and 4.6, this assumption greatly reduces mathematical complexity yet has almost no bearing on the generalizability of our results. Furthermore, if the firm were unwilling to develop favorably recommended designs, then experts would have no motive to expend effort testing design alternatives.

As regards the firm's schemes for compensating the experts, we allow the firm to offer asymmetric and nonlinear contracts that include any combination of action- and evidence-based payments.[6] An *action-based* payment is contingent on a specific action taken by the firm—for instance, developing design i.[7] In contrast, an *evidence-based* payment depends on a design's true technological feasibility Θ_i. In practice, any pay-for-performance contract must be action-based and/or evidence-based; the reason is that only such criteria are verifiable and thus enforceable by courts. In our model, then, an expert who is testing design i is eligible for the following payment types: (i) a "success bonus" u_{ig} if the firm successfully develops design i (i.e., $\Theta_i = G$); (ii) an "allowance" u_{ib} if the firm's development of design i fails (i.e., $\Theta_i = B$); (iii) a "consolation award" u_{ia} if the expert tested design i but the firm did not choose it for development; and (iv) a "termination bonus" u_{it} if none of the available design alternatives is chosen for development. We assume throughout that the firm must make all wage payments immediately when due; that is, it cannot hold back any wages. It is intuitive that the evidence-based payments u_{ig} and u_{ib} incentivize an expert to test the design thoroughly and also to recommend the design's development in the event of a favorable signal. In contrast, u_{ia} and u_{it} reflect the firm's tolerance for failure. That is to say, the firm appreciates an expert's testing efforts even in the case of negative testing outcomes; hence this payment incentivizes experts to refrain from recommending a bad design for development.

Given such compensation schemes, the utility π_i received by an expert for testing design i is the (discounted) sum of all his wage payments net of his effort costs. We follow the principal–agent literature in assuming that all experts are risk neutral and protected by limited liability; in other words, the compensation of each expert must be nonnegative at all times. The firm is risk neutral, too, and its profit Π consists of the realized value of the developed design (net of any development costs) *minus* the compensation paid to experts.

4.4 Incentives for Delegated Testing

In this section we characterize the optimal compensation schemes, for the different testing strategies, given that the set of design alternatives $\mathcal{I} \subseteq \mathcal{N}$ (with $|\mathcal{I}| = n \geq 1$)

[6]Note that, without loss of optimality, we do not need to consider fixed wages because in optimum, such a fixed wage must be zero as experts are shielded by limited liability.

[7]Payments cannot depend on the *expert*'s action because such actions cannot be verified.

is exogenously fixed.[8] In line with the revelation principle, we limit our attention to contracts that incentivize all experts to evaluate design(s) thoroughly and to reveal test outcomes truthfully. We begin by identifying the optimal compensation scheme for parallel testing (Section 4.4.1), after which we characterize the optimal contract for multi-expert sequential testing and derive the optimal testing order (Section 4.4.2). Finally, we examine the optimal incentive scheme and testing order for the single-expert sequential testing strategy (Section 4.4.3).

4.4.1. Parallel Testing

Under a parallel testing strategy, the firm assigns a different expert to each design $i \in \mathcal{I} = \{1, \ldots, n\}$, and all experts then perform their testing activities simultaneously. This setup allows the firm to receive refined information on each design's value and to use this information to select, ex post, the most promising alternative for development. The firm's overarching goal is to maximize expected profits, which amount to the expected market value of the developed design *net* of development costs and compensation paid to experts. Formally, the firm solves the following optimization problem P (whose mathematical derivation, along with all other formal proofs, has been relegated to Appendix C):

$$P: \quad \max_{u,y} \ \Pi := \sum_{j=1}^{n} \frac{1}{2^j} \left(\sum_{i=1}^{n} [y_i^{(j)}(q_i v_i - q_i u_{ig} - (1-q_i)u_{ib} - (2^j - 1)u_{ia} - 2^{j-n}u_{it})] - K \right) \tag{4.1}$$

$$\text{s.t.} \ \ y_i^{(j)}(q_i u_{ig} + (1-q_i)u_{ib}) \geq y_i^{(j)}(u_{ia} + 2^{j-n}u_{it}) \qquad \forall i, j \in \mathcal{I} \tag{4.2}$$

$$y_i^{(j)}((1-q_i)u_{ig} + q_i u_{ib}) \leq y_i^{(j)}(u_{ia} + 2^{j-n}u_{it}) \qquad \forall i, j \in \mathcal{I} \tag{4.3}$$

$$y_i^{(j)}(u_{ig} - u_{ib}) \geq y_i^{(j)} 2^{j+1} c/(2q_i - 1) \qquad \forall i, j \in \mathcal{I} \tag{4.4}$$

$$\begin{aligned} &y_i^{(j)}[q_i(v_i - u_{ig}) - (1-q_i)u_{ib} - \sum_{k \neq i} u_{ka}] \\ &\geq y_i^{(j)} \sum_{l=1}^{n} y_l^{(j+1)}[q_l(v_l - u_{lg}) - (1-q_l)u_{lb} - \sum_{k \neq l} u_{ka}] \qquad \forall i \in \mathcal{I}, j \in \mathcal{I}\backslash\{n\} \end{aligned} \tag{4.5}$$

[8]We endogenize the firm's choice of design alternatives in Section 4.5.1.

$$\sum_{i=1}^{n} y_i^{(j)} = 1, \ \ \sum_{j=1}^{n} y_i^{(j)} = 1, \ \ y_i^{(j)} \in \{0,1\} \qquad \forall i,j \in \mathcal{I} \tag{4.6}$$

$$u_{ig}, u_{ib}, u_{ia}, u_{ia} + u_{it} \geq 0 \qquad \forall i \in \mathcal{I} \tag{4.7}$$

Although complex at first sight, this optimization problem P has an intuitive structure. As a starting point, note that $y_i^{(j)}$ is an indicator variable that reflects whether design i is the firm's jth most preferred alternative (constraint (4.6) ensures that this mapping is indeed one-to-one). That is, if $y_i^{(j)} = 1$ then the firm chooses design i for development only if it receives an unfavorable recommendation for all designs with a lower ranking $j' < j$. Clearly, a design's attractiveness is not exogenously given and instead depends endogenously on the offered compensation scheme. This fact is reflected by (4.5), which guarantees that the firm makes an ex post optimal selection decision. Conditions (4.2)–(4.4) represent each expert's incentive compatibility constraints, which depend on the relative attractiveness of the design she has been assigned to evaluate. Thus (4.2) and (4.3) ensure that each expert truthfully reveals, respectively, a "good" and a "bad" signal. These constraints eliminate the adverse selection problem during the recommendation phase. Condition (4.4) similarly negates the moral hazard problem during the design-testing phase because it ensures that each expert prefers high-effort to low-effort testing. Finally, experts are protected by limited liability; hence (4.7) ensures that all wage payments are nonnegative. The following proposition characterizes—under mild conditions on the properties of designs in $\mathcal{I}$—the optimal incentive structures for parallel testing and the resulting firm profits.[9]

Proposition 4.1. *Suppose the designs in $\mathcal{I}$ can be ordered such that $q_i v_i \geq q_{i+1} v_{i+1} + 2^{i+1} c[q_i/(2q_i - 1) - 2q_{i+1}/(2q_{i+1} - 1)]^+$ for all $i \in \mathcal{I} \backslash \{n\}$, where $[x]^+ = \max\{0, x\}$. Then the following statements hold.*

(i) *Under a parallel testing strategy, the optimal contract that induces truth telling and high-effort testing for each design satisfies $u_{ig} = 2^{i+1} c/(2q_i - 1)$, $u_{ib} = 0$, $u_{ia} = 0$, and $u_{it} = 2^{n+1}(1 - q_i)c/(2q_i - 1)$ for all $i \in \mathcal{I}$. Moreover, $u_{it}/u_{ig} = 2^n(1 - q_i)/2^i < 2^{n-i-1} \leq 2^{n-2}$ for all $i \in \mathcal{I}$.*

(ii) *Ex ante, the firm's expected profit is $\Pi_P = \sum_{i=1}^{n} ((q_i v_i - K)/2^i - 2c/(2q_i - 1))$.*

[9] The condition in Proposition 4.1 holds unless $\mathcal{I}$ contains a design for which the test is of exceptionally low efficiency.

Perhaps the most remarkable aspect of this proposition is the simplicity of the optimal contract's structure. For each design (and thus for each expert) $i \in \mathcal{I}$, the firm need offer only two payments: a success bonus u_{ig} if design i is successfully developed, and a termination bonus u_{it} if none of the design alternatives is chosen for development (i.e., if test outcomes indicate that all designs are technologically unfeasible). But what respective roles do these payments play in the firm's incentive system? As (4.4) reveals, the primary purpose of u_{ig} is to motivate expert i to engage in high-effort testing. Put differently, u_{ig} is a purely individual incentive that resolves each expert's moral hazard concern. In contrast, u_{it} is a common (or shared) incentive that collectively compensates the experts if all designs are considered to be technologically unfeasible. It therefore decreases the likelihood of an expert giving a positive recommendation despite receiving a negative test outcome—and thereby induces truth telling; see constraint (4.3).

It is intuitive that, when effort becomes less rewarding (i.e., effort costs c increase) and recommendations become less reliable (the signal quality q_i decreases), expert i becomes more reluctant to invest high effort and to report test outcomes truthfully. Under these circumstances, the firm must provide stronger incentives; this explains why u_{ig} and u_{it} are increasing in c and decreasing in q_i. Conversely, in the extreme case of perfect testing ($q_i = 1$), expert i has no incentive to misrepresent the test outcomes because such a false recommendation would be easily detected by the firm and so would not benefit him. It follows that if testing is perfect then the firm can forgo payment of any shared incentives ($u_{it} = 0$).

Whereas neither u_{ig} nor u_{it} depends directly on a design's inherent economic potential v_i, these terms *are* affected by the total number n of designs to be tested and by a design's relative value, which we also index via i. In particular, u_{ig} increases with i because, with a higher index i, it becomes more likely for tests to indicate that a design with a smaller index is technologically feasible—which would render futile the expert's testing efforts. Similarly, with a higher n it becomes less likely that all n designs are technologically unfeasible; hence the termination bonus u_{it} is correspondingly less likely to be paid out. To compensate the experts for this reduced payment probability, the firm must offer a higher u_{it}.

Proposition 4.1(i) also sheds light on the severity of adverse selection—as compared with moral hazard—when the firm employs a parallel testing strategy. For designs that are extremely promising ex ante (i.e., those with a small index i), the ratio u_{it}/u_{ig} is high; the implication is that, for these designs, the firm's central concern is to incentivize truth telling. For ex ante less promising designs (those with a large index i) the ratio is low; in that event, the firm becomes relatively more concerned with incentivizing experts to exert high effort. This shifting priority of the optimal compensation scheme has an appealing explanation. From an expert's perspective, admitting that her own testing outcome is bad increases the odds of receiving *no* payments. This conclusion follows from the extreme unlikelihood of all tested designs being pronounced technologically unfeasible. The firm's concern on this point is especially strong for designs that show the most promise. At the same time, an expert assigned to a promising design has a significant intrinsic motivation to exert high effort because the potential rewards from receiving a good test outcome are high. In contrast, an expert assigned to a less promising design fears that her efforts are futile because, in all probability, a more promising design will receive a good recommendation and so the results of her testing could be irrelevant to the firm. Therefore, relatively higher individual incentives must be offered to the experts who are assigned to test less promising designs.

Finally, Proposition 4.1(ii) reveals that the firm's expected profit is decreasing in c and increasing in q_i. The reason is that, with a higher c and a lower q_i, expert i is less willing to exert high effort and to disclose the test outcomes truthfully. Hence the incentive misalignment between the firm and expert i widens, which enables the expert to extract higher information rents.

4.4.2. Multi-Expert Sequential Testing

Under a multi-expert sequential testing strategy, the firm assigns a different expert to test each design $i \in \mathcal{I} = \{1, \ldots, n\}$ and announces the testing sequence; then the experts carry out their design tests one after the other. A sequential testing approach allows the firm to stop the testing process when it receives a positive recommendation (i.e., so it can start developing that design) or to test the next viable alternative when it receives a negative recommendation.

For a given $\mathcal{I}$ and any testing order, the firm must solve the following optimization problem to derive the optimal compensation schemes. For expositional simplicity, we relabel the design alternatives such that a design's index i is identical to its position in the testing order; thus design $i = 1$ is tested first, $i = 2$ second, and so on. The optimization problem M is expressed formally as follows:

$$M: \quad \max_{u} \; \Pi(u) := \sum_{i=1}^{n} \frac{\delta^{i-1}}{2^i}(q_i v_i - K) - \sum_{i=1}^{n} \frac{\delta^{i-1}}{2^i}(q_i u_{ig} + (1-q_i)u_{ib} + u_{ia} + 2^{i-n}\delta^{n-i}u_{it}) \tag{4.8}$$

$$\text{s.t. } q_i u_{ig} + (1-q_i)u_{ib} \geq u_{ia} + 2^{i-n}\delta^{n-i}u_{it} \qquad \forall i \in \mathcal{I} \tag{4.9}$$

$$(1-q_i)u_{ig} + q_i u_{ib} \leq u_{ia} + 2^{i-n}\delta^{n-i}u_{it} \qquad \forall i \in \mathcal{I} \tag{4.10}$$

$$u_{ig} - u_{ib} \geq \frac{4c}{2q_i - 1} \qquad \forall i \in \mathcal{I} \tag{4.11}$$

$$u_{ig}, u_{ib}, u_{ia}, u_{it} \geq 0 \qquad \forall i \in \mathcal{I}\backslash\{n\}, \qquad u_{ng}, u_{nb}, u_{na} + u_{nt} \geq 0 \tag{4.12}$$

The structure of the optimization problem M is similar to the firm's optimization problem P under a parallel testing strategy. Specifically, constraints (4.9) and (4.10) ensure that all experts truthfully reveal their testing outcomes, (4.11) guarantees that each expert engages in high-effort testing, and (4.12) accounts for the experts' limited liability. In the next proposition we derive the firm's optimal compensation schemes, describe the optimal testing order, and state the resulting firm profits for a multi-expert sequential testing strategy.

Proposition 4.2. (i) *Under a multi-expert sequential testing strategy and for any given testing order, the optimal contract that induces truth telling and high-effort testing by all experts satisfies, for all* $i \in \mathcal{I}$: $u_{ig} = 4c/(2q_i - 1)$, $u_{ib} = 0$, $u_{ia} = 4(1-q_i)c/(2q_i - 1)$, *and* $u_{it} = 0$. *In addition,* $u_{ia}/u_{ig} = 1 - q_i < 1/2$.

(ii) *It is optimal to test the designs in* $\mathcal{I}$ *in decreasing order of* $R_i \equiv q_i v_i - 4c/(2q_i - 1)$.

(iii) *Ex ante, the firm's expected profit is* $\Pi_M = \sum_{i=1}^{n} \delta^{i-1}(q_i v_i - K - 4c/(2q_i - 1))/2^i$.

As in the optimal contract for parallel testing, there are only two payments in the optimal compensation scheme for multi-expert sequential testing. To resolve each expert's moral hazard, the firm must again reward expert i with a success bonus u_{ig} in the event design i is developed successfully. Yet in this case the firm does not rely on shared incentives to induce truth telling; that is, $u_{it} = 0$. Instead the firm provides an individual consolation award u_{ia} to reimburse expert i for his effort costs whenever the firm dismisses the design he tested owing to the subsequent unfavorable recommendation.

This focus on purely individual incentives has two immediate consequences. First, the optimal payments depend not on a design's position in the testing order but only on its informational quality q_i: the higher the informational quality, the more aligned are the interests of firm and experts and so the lower is the compensation offered. Second, as the ratio u_{ia}/u_{ig} indicates, if testing is sequential then moral hazard is a much greater concern than adverse selection, especially when q_i is high. This finding may be better understood if one notes that, from an expert's perspective, recommending a technologically unfeasible design for development leads to zero income ($u_{ib} = 0$). So once an expert has invested high effort in testing, there is hardly any point in trying to pass off a technologically unfeasible design as a good one. Even so, motivating the expert to engage in high-effort testing at the outset requires a high effort incentive (high u_{ig}).

The optimal testing order is given in part (ii) of Proposition 4.2, the essence of which is that the firm should test designs in decreasing order of their expected net contribution R_i. This result is in the spirit of Weitzman's (1979) reservation price rule but extends it to include the costs of delegation. That is, the testing order depends not only on the designs' expected values $q_i v_i$ but also on the experts' information rents. Since these information rents are decreasing in q_i and invariant with respect to v_i, it follows that the firm—as compared to the reservation price rules advocated by Weitzman (1979), Adam (2001), and Erat and Kavadias (2008)—more strongly prefers first to test designs of high informational quality. The structure of R_i indicates that the firm's optimal testing order is myopic: a design's expected net contribution depends only on its own properties and so is independent of other design alternatives.

Recall our argument that an expert's information rents are decreasing in the quality of her information and increasing in her effort costs. From these relations it clearly follows that the firm's expected profit should be increasing in q_i and decreasing in c. Proposition 4.2(iii) confirms this intuition and also underscores how the firm's profit is adversely affected when δ, the time value of money, is low.

Finally, we emphasize that the compensation scheme and testing order presented in Proposition 4.2 remain optimal even if the firm is allowed to develop formerly tested (yet rejected) design alternatives, that is, if the firm can test "with recall". To see this, note that if the firm receives a good recommendation for some $i \in \mathcal{I}$ then it can immediately realize an expected profit of $q_i v_i - K - 4q_i c/(2q_i - 1)$ by developing design i right away. If instead the firm chooses to test the next design alternative, then the optimal ordering in Proposition 4.2(ii) implies that the firm's expected continuation profit is strictly smaller.

We conclude that, as soon as the firm receives a good recommendation for a particular design, it is optimal to develop this design at once—rendering a recall option superfluous.

4.4.3. Single-Expert Sequential Testing

Under a single-expert sequential testing strategy, the firm assigns a single expert to test in sequence the design alternatives in $\mathcal{I}$. It is easy to see that—as compared with a multi-expert strategy—such reliance on the testing efforts of only a single expert will have a strong bearing on the required incentives. On the one hand, a single expert is much more inclined (than is one in a group of experts) to acknowledge an unfavorable test outcome because there is always the chance of finding a technologically feasible design later in the testing process. On the other hand, it is extremely difficult to continue incentivizing an expert to exert high testing efforts. Thus an expert's behavior during the testing process is strongly affected by his anticipation of future actions and payments.

For a given $\mathcal{I}$ and any testing order, the firm's incentive design problem is as follows. As in the preceding section, we relabel the design alternatives such that a design's index i is identical to its position in the testing order. The incentive design problem S is then

$$S: \quad \max_{u} \ \Pi(u) := \sum_{i=1}^{n} \frac{\delta^{i-1}}{2^i}(q_i v_i - K) - \sum_{i=1}^{n} \frac{\delta^{i-1}}{2^i}(q_i u_{ig} + (1-q_i)u_{ib} + u_{ia}) \tag{4.13}$$

$$\text{s.t.} \quad q_i u_{ig} + (1-q_i)u_{ib} \geq u_{ia} + \delta\hat{\pi}_i \qquad \forall i \in \mathcal{I} \tag{4.14}$$

$$(1-q_i)u_{ig} + q_i u_{ib} \leq u_{ia} + \delta\hat{\pi}_i \qquad \forall i \in \mathcal{I} \tag{4.15}$$

$$u_{ig} - u_{ib} \geq \frac{4c}{2q_i - 1} \qquad \forall i \in \mathcal{I} \tag{4.16}$$

$$\begin{aligned} &\hat{\pi}_{i-1} = (q_i u_{ig} + (1-q_i)u_{ib} + u_{ia} - 2c + \delta\hat{\pi}_i)/2 \qquad \forall i \in \mathcal{I} \\ &\hat{\pi}_n = 0 \end{aligned} \tag{4.17}$$

$$u_{ig}, u_{ib}, u_{ia} \geq 0 \qquad \forall i \in \mathcal{I} \tag{4.18}$$

Some peculiarities of the optimization problem S warrant further discussion. First, with a single expert it is unnecessary to have an additional termination bonus u_t that compensates her when none of the design alternatives is developed. In fact, such a payment can—without loss of optimality—be folded into the expert's consolation award for testing the last design alternative (u_{na}). This follows because both payments have the same requirements and are executed simultaneously. Second, $\hat{\pi}_{i-1}$ as defined in (4.17) is the expert's expected continuation utility immediately *before* testing design i. Since the expert's decision-making process accounts for her own future utility, it is only natural for $\hat{\pi}_i$ to become an integral part of her incentive constraints; see (4.14) and (4.15). More precisely: as compared with a multi-expert sequential testing strategy, a single sequentially testing expert is more (resp. less) likely to report an unfavorable (resp. favorable) signal truthfully. Our explanation is that the expert may enjoy additional information rents by artificially keeping the testing process alive (i.e., by concealing a signal of feasibility). As shown by Proposition 4.3, that possibility has important consequences for the optimal compensation scheme.

Proposition 4.3. (i) *Under a single-expert sequential testing strategy and for any given testing order, the optimal contract that induces truth telling and high-effort testing for all designs satisfies, for all $i \in \mathcal{I}$: $u_{ig} = 4c/(2q_i - 1) + [\delta\hat{\pi}_i/q_i - 4c/(2q_i - 1)]^+$, $u_{ib} = 0$, and $u_{ia} = [4(1 - q_i)c/(2q_i - 1) - \delta\hat{\pi}_i]^+$. Moreover, $u_{ia}/u_{ig} \leq 1 - q_i < 1/2$.*

(ii) *If the design alternatives in $\mathcal{I}$ can be ordered such that $q_i v_i \geq q_{i+1}v_{i+1}$, $q_i \geq q_{i+1}$, and $(1 - q_i)4c/(2q_i - 1) \leq \delta\hat{\pi}_i \leq 4q_i c/(2q_i - 1)$ for all $i \in \mathcal{I}\backslash\{n\}$, then it is optimal to test in increasing order of i.*

(iii) *Ex ante, the firm's expected profit is $\Pi_S = \sum_{i=1}^{n} \delta^{i-1}(q_i v_i - K - \max\{4q_i c/(2q_i - 1), \delta\hat{\pi}_i, 4c/(2q_i - 1) - \delta\hat{\pi}_i\})/2^i$.*

Although the optimal contract for a single-expert sequential testing strategy is structurally similar to that for a multi-expert strategy, there are some important differences. First of all, under single-expert sequential testing, the firm needs to place more emphasis on motivating high-effort testing and less on inducing experts to report truthfully. Correspondingly, the success bonus u_{ig} is higher under single-expert than multi-expert sequential testing while the consolation award u_{ia} is substantially lower. In fact, it may be optimal for the firm to offer no consolation award at all ($u_{ia} = 0$). That would be the case for sufficiently large values of $\hat{\pi}_i$, the expert's expected continuation utility. Here the single expert anticipates substantial future payments if the testing process continues; therefore, under sequential testing, that expert will always truthfully report an unfavorable signal. This dynamic has the effect of eliminating the adverse selection problem.

Second, and in contrast to the multi-expert strategy detailed previously, the optimal payments related to each design i are not myopic in the single-expert setting; instead those payments depend on the informational quality of all designs tested *after* design i. This result is a natural and direct consequence of the expert's strategic behavior, and as such it bears implications for the optimal testing order. Following the logic of Weitzman (1979) and Proposition 4.2(ii), one might well suppose that it is still optimal to test the designs in decreasing order of their expected net contribution. However, that supposition is not true in general. We can see from Proposition 4.3(ii) that such a testing order is optimal only if the expert anticipates a moderate level of continuation utilities (i.e., only if $(1-q_i)4c/(2q_i-1) \leq \delta\hat{\pi}_i \leq 4q_ic/(2q_i-1)$). However, if continuing with the testing process promises continuation utilities that are exceptionally high or low, then the firm should *not* test the designs in decreasing order of attractiveness. It might rather be optimal to test the least promising designs first—with the goal of reducing the expert's strategic rent extraction.

Finally, Proposition 4.3(iii) gives the firm's expected profit under a single-expert sequential testing strategy and yields a rather surprising result. Unlike the other testing strategies, under single-expert sequential testing the firm's expected profit need not increase with quality q_i. Because of the expert's strategic behavior, a higher q_i for one design might result in the expert extracting higher information rents from the other designs being tested; this would, of course, *reduce* the firm's overall profits.

4.5 Comparison of Testing Strategies

So far, we have characterized the optimal incentive structure for the three different testing strategies given that the set of design alternatives was exogenously fixed. As a next step, we relax this assumption and determine the optimal number and identity of designs to test for each of the three testing strategies (Section 4.5.1). We then build on these results by deriving the optimal testing strategy as a function of our main contextual parameters (Section 4.5.2). Finally, we offer some insights regarding the question of how delegation, which entails information asymmetry, alters the relative ranking of testing strategies—that is, from the ranking in an otherwise identical setting but *without* information asymmetry (Section 4.5.3).

4.5.1. Optimal Set of Design Alternatives

In Section 4.4 we derived optimal compensation schemes for the different testing strategies while assuming that the firm intended to test a fixed set $\mathcal{I} \subseteq \mathcal{N}$ of design alternatives. A more realistic scenario is one in which the set of designs to test is not given exogenously but instead is chosen by the firm. In this section, then, we characterize the optimal sets of designs to be tested for the different testing strategies. Our results reveal that these sets vary considerably across those strategies.

Proposition 4.4. (i) *Under a multi-expert sequential testing strategy, the optimal set of designs to be tested is* $\mathcal{I}_M = \{i \in \mathcal{N} \mid q_i v_i - K - 4c/(2q_i - 1) \geq 0\}$.

(ii) *Under a parallel testing strategy, the optimal set of designs to be tested satisfies* $\mathcal{I}_P \subseteq \mathcal{I}_M$.

(iii) *Let* $\mathcal{I}_S$ *be the optimal set of designs to be tested under a single-expert sequential testing strategy, and let* n *be the last design in the optimal testing order. If* $q_n v_n - K - 4c/(2q_n - 1) \geq 0$ *and if* $q_n v_n \leq q_i v_i$ *and* $q_n \leq q_i$ *for all* $i \in \mathcal{I}_S \backslash \{n\}$, *then* $\mathcal{I}_S \subseteq \mathcal{I}_M$.

Part (i) of this proposition offers a detailed characterization of the optimal identity ($i \in \mathcal{I}_M$) and number ($|\mathcal{I}_M|$) of designs to test under a multi-expert sequential testing strategy. In particular, the firm should test any design $i \in \mathcal{N}$ for which the expected value $q_i v_i$ exceeds the *sum* of (a) the expert's information rents $4c/(2q_i - 1)$ and (b) the development costs K. That is, the firm should test only those designs that promise ex ante a positive contribution margin.

A similar argument applies to the optimal set of designs to be tested under a parallel testing strategy $\mathcal{I}_P$. However, as indicated by Proposition 4.1(iii) and Proposition 4.2(iii), the experts' information rents under parallel testing are much higher than under sequential testing, which explains why the firm always tests fewer designs than under a multi-expert sequential testing strategy. There are two reasons for this difference. First, under parallel testing the firm does not have the option to stop the testing process prematurely. Second, experts testing a relatively undesirable design know that the firm will probably consider their recommendations to be irrelevant; it is therefore costly for the firm to motivate these experts to exert high testing efforts. Whereas the first dynamic has been well established by previous academic work (see, e.g., Loch et al. 2001), little attention has been paid to the second source of inefficiency.

Finally, Proposition 4.4(iii) derives some properties of the optimal set of designs to be tested under a single-expert sequential testing strategy. A few observations merit discussion here. We note that if the last design in the testing order is also the least promising alternative—yet still offers a positive contribution margin—then the firm always tests fewer designs than under a multi-expert sequential testing strategy. In other words, the expert's strategic behavior induces the firm to make reductions in the number of designs to test. Yet this generalization does admit some exceptions. In some instances, it might be profitable for the firm to include an ex ante unprofitable design in its test set $\mathcal{I}_S$ for the sole purpose of influencing the expert's continuation utility and thereby reducing his strategic behavior. In such cases, $q_n v_n - K - 4c/(2q_n - 1) < 0$ and so the firm may find it optimal to increase the number of designs to test: $\mathcal{I}_S \supset \mathcal{I}_M$.

In sum: the firm tests only ex ante profitable designs under both the parallel and multi-expert sequential testing strategy; under single-expert sequential testing, however, it may be optimal for the firm to test ex ante unprofitable designs in order to curtail rent extraction by experts.

4.5.2. Optimal Testing Strategy

Given the optimal contract structures and the optimal set of design alternatives for the different testing strategies, we can now turn to our main research question: What is the firm's optimal testing strategy under delegation? Propositions 4.1–4.3 have already indicated that the answer to this question depends mainly on two contextual parameters: the costs of effort (c) and the informational quality of test outcomes (q_i). It seems clear that these two parameters determine how much information rent the firm must sacrifice in order to align the experts' interests with the firm's agenda. The next proposition confirms this intuition.

Proposition 4.5. *Let* $\Pi_P^* = \Pi_P(\mathcal{I}_P)$, $\Pi_M^* = \Pi_M(\mathcal{I}_M)$, *and* $\Pi_S^* = \Pi_S(\mathcal{I}_S)$. *Then the following statements hold.*

(i) *If* $\delta < 1$, *then there exists a* $\underline{c} > 0$ *such that* $\Pi_P^* > \max\{\Pi_M^*, \Pi_S^*\}$ *for all* $c < \underline{c}$.

(ii) *Let* $\mathcal{I}_P$ *be the optimal set of designs to be tested under a parallel testing strategy. If those designs can be ordered such that* $q_i v_i \geq q_{i+1} v_{i+1} + 2^{i+1} c[q_i/(2q_i - 1) - 2q_{i+1}/(2q_{i+1} - 1)]^+$ *for all* $i \in \mathcal{I}_P \backslash \{n\}$, *then* $\max\{\Pi_M^*, \Pi_S^*\} > \Pi_P^*$ *provided that* $c > \overline{c} \equiv \sum_{i \in \mathcal{I}_P}((1 - \delta^{i-1})(q_i v_i - K)/2^i)/\sum_{i \in \mathcal{I}_P}((2(1 - (\delta/2)^{i-1})/(2q_i - 1))$.

(iii) *Let $\mathcal{I}_M$ be the optimal set of designs to be tested under a multi-expert sequential testing strategy, and let those designs be optimally ordered according to Proposition 4.2(ii). Then $\Pi_S^* \geq \Pi_M^*$ provided that $q_{i+1} \geq \underline{q}_i \equiv 1/2 + \delta(2q_i - 1)/(4q_i - \delta(2q_i - 1))$ for all $i \in \mathcal{I}_M \backslash \{n\}$. Moreover, $\underline{q}_i \leq \min\{q_i, 5/6\}$.*

In Figure 4.1, the left panel illustrates the key properties of the optimal testing strategy. First, a parallel testing strategy is undertaken only when the testing costs are sufficiently small ($c < \underline{c}$). Otherwise, the burden of paying all experts immediately for their testing efforts is greater than the value of information received; in that case, the firm decides to implement a sequential testing strategy. Thus design tests that are more expensive—and the resulting higher information rents—make sequential testing more economical. In this respect, parts (i) and (ii) of Proposition 4.5 extend previous findings of Loch et al. (2001) to testing processes that are prone to information asymmetry.

However, it remains an open question exactly how the firm should implement a sequential testing strategy. That is, should the firm hire multiple experts to test the different design alternatives, or should it rather assign all testing activities to a single expert? Proposition 4.5(iii) shows that this question's answer is closely tied to the test efficiency, q_i, of the different designs. Relying on a single expert is especially beneficial when the informational quality of the different design tests is relatively homogeneous (i.e., when $q_{i+1} \geq \underline{q}_i$). In contrast, if test efficiency is heterogeneous across designs then the firm is better-off assigning a different expert to each design alternative. The explanation for this finding is instructive. When the designs' test efficiencies are very different, then the firm prefers testing the designs of highest informational quality first and testing those of lowest quality last (cf. Proposition 4.3(ii)). However, the informational rents extracted from the firm by a single expert increase with any decline in the informational quality of a design test. Hence the expert tries to keep the testing process alive as long as possible—even if that requires reporting a negative assessment of what is actually a good design. There can be no question that exposure to such strategic behavior is suboptimal for the firm, which should therefore rely instead on multi-expert sequential testing.

The left panel of Figure 4.1 also shows the role that the discount factor δ plays in the firm's choice of an optimal testing strategy. As expected, sequential testing strategies are preferable when the time factor is less critical for the firm—that is, as δ increases. Less obvious, though, is δ's effect on the firm's preferred sequential testing strategy. We

Figure 4.1.: The Firm's Optimal Testing Strategy.

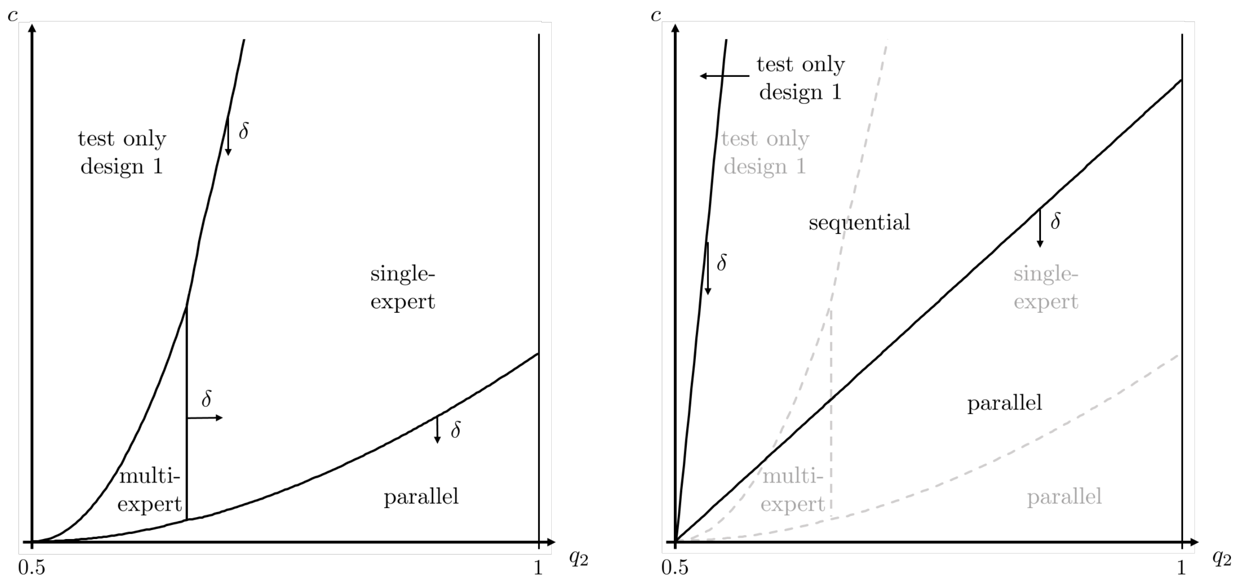

Notes: The graphs plot, for an example with $N = 2$ design alternatives, the firm's optimal testing strategy under delegation (left panel) and first-best (right panel) conditions. Under a *parallel* testing strategy it is optimal to test both designs simultaneously. Under any sequential testing strategy (i.e., *multi-expert*, *single-expert*, or *first-best*) it is optimal to test design $i = 1$ first and $i = 2$ second. The other parameter values are $v_1 = 1000$, $v_2 = 400$, $q_1 = 0.9$, $K = 200$, and $\delta = 0.9$.

find that a higher δ facilitates the single expert's strategic extraction of rent because prolonging the testing process is then less costly for her. It follows that the size of the region in which the firm prefers multi-expert to single-expert sequential testing increases with δ (cf. the sensitivity of $\underline{q}_i$ as given in Proposition 4.5(iii)).

Finally, Proposition 4.5 and Proposition 4.4 together reveal an interesting non-monotonicity in the optimal number of designs to test. For very low testing costs c, the firm pursues a parallel testing strategy and simultaneously tests a moderate number $|\mathcal{I}_P|$ of designs. As c increases, the firm moves to a sequential testing strategy and, in so doing, increases the number of designs to test (recall that $|\mathcal{I}_M| \geq |\mathcal{I}_P|$ by Proposition 4.4(ii)). Yet when c becomes too large, design testing becomes so costly that the firm is impelled to reverse course and reduce the number of design tests. These results contradict the conventional wisdom—which is true in the *absence* of information asymmetry—that lower testing costs unequivocally lead to more design tests.

4.5.3. Costs of Delegation

Our aim in this section is to discover precisely how information asymmetries distort the firm's design-testing process. We start by describing, as a basis for comparison, the firm's first-best testing strategy: one in which both experts and firm behave as a single entity. Then, in Proposition 4.6, we compare this first-best strategy with the optimal testing strategy under delegation.

If the incentives of experts and the firm are aligned, then the latter need not pay any action- or evidence-based bonuses to motivate the former to engage in high-effort testing and to reveal their testing outcomes truthfully. So absent incentive misalignment, the firm can simply reimburse the experts for their testing efforts by paying them their effort costs c for each design test conducted. Given the resulting lack of information asymmetry, the firm's first-best (fb) expected profit under a sequential (seq) testing strategy is given by $\Pi^{\text{fb}}_{\text{seq}} = \sum_{i \in \mathcal{I}^{\text{fb}}_{\text{seq}}} \delta^{i-1}(q_i v_i - K - 2c)/2^i$; here $\mathcal{I}^{\text{fb}}_{\text{seq}} = \{i \in \mathcal{N} \mid q_i v_i - K - 2c \geq 0\}$ is the optimal set of designs to be tested, and the firm tests the designs in decreasing order of $q_i v_i$. Analogously, the firm's first-best expected profit under a parallel (par) testing strategy is $\Pi^{\text{fb}}_{\text{par}} = \sum_{i \in \mathcal{I}^{\text{fb}}_{\text{par}}}((q_i v_i - K)/2^i - c)$; here the designs in $\mathcal{I}^{\text{fb}}_{\text{par}}$ are ordered in decreasing order of $q_i v_i$, and $\mathcal{I}^{\text{fb}}_{\text{par}} \subseteq \mathcal{I}^{\text{fb}}_{\text{seq}}$. Our final proposition leverages these insights to identify how information asymmetry affects the firm's optimal testing strategy.

Proposition 4.6. (i) *Under delegated sequential testing, the optimal set of designs to be tested is a subset of the first-best set*: $\mathcal{I}_S, \mathcal{I}_M \subseteq \mathcal{I}^{\text{fb}}_{\text{seq}}$.

(ii) *Under delegated parallel testing, the firm may choose a completely different set of designs to test than under first-best conditions; thus there are cases in which* $\mathcal{I}_P \cap \mathcal{I}^{\text{fb}}_{\text{par}} = \emptyset$.

(iii) *Suppose* $q_i = q$ *for all* $i \in \mathcal{N}$. *If* $\Pi^{\text{fb}}_{\text{seq}} \geq \Pi^{\text{fb}}_{\text{par}}$, *then* $\max\{\Pi^*_M, \Pi^*_S\} \geq \Pi^*_P$; *however, the converse is not true in general.*

The main finding of Proposition 4.6 is that information asymmetry has fundamentally different effects on parallel than on sequential testing. Consider first the implications of delegation on the optimal design of a sequential testing strategy. Part (i) of the proposition shows that—as expected—the presence of information asymmetry results in a suboptimal testing process. In particular, the firm is testing too few designs and therefore stops the testing process too early; that is, $\mathcal{I}_S, \mathcal{I}_M \subseteq \mathcal{I}^{\text{fb}}_{\text{seq}}$. This result reflects that an expert's information rent makes design testing unequivocally more expensive (for the firm) than under first-best conditions.

One might suppose a similar reasoning to apply also with regard to parallel testing. In this respect, however, part (ii) of Proposition 4.6 holds a surprise. Note that even though the firm tests too few designs under a sequential testing strategy, it does still test those designs that are also the most promising ones under first-best conditions. Yet this statement does not necessarily hold for a parallel testing strategy. In fact, Proposition 4.6(ii) reveals that the optimal sets of designs to be tested with and without information asymmetry may be disjoint; under delegation, then, the firm may test an *entirely* different set of design alternatives. How can we explain this split? Recall from our discussion after Proposition 4.1 that the information rents extracted from the firm by experts are *decreasing* in the quality of those experts' information. Hence the firm never tests designs that offer relatively poor information quality—that is, with almost complete disregard for their economic potential v_i. In contrast, under first-best conditions the firm's testing costs are constant and thus do not depend on a design's informational quality; in that case, it makes sense for the firm always to test those designs promising the highest expected value $q_i v_i$. Evidently, these different priorities under first-best and delegated testing can lead to disjoint optimal test sets. This phenomenon is most likely to occur when some design alternatives are of exceptionally high economic value v_i but low test efficiency q_i.

Finally, Proposition 4.6(iii) hints at an important managerial insight: delegation strongly favors sequential testing. This result is also clearly illustrated in the right panel of Figure 4.1, which plots the firm's optimal testing strategy vis-à-vis the first-best benchmark. Under symmetric test efficiencies ($q_i = q$ for all $i \in \mathcal{N}$), we can demonstrate formally that if the firm prefers sequential testing under first-best conditions then it does so under delegation as well. Although we are unfortunately not able to generalize this result analytically to heterogeneous test efficiencies, our numerical experiments confirm that the claim does indeed hold much more generally; see the right panel of Figure 4.1, which allows for such heterogeneity. Our finding has implications both for the academic literature and for practice. It questions at a fundamental level the claims of those (e.g., Dahan and Mendelson 2001, Loch et al. 2001) who have praised the effectiveness of parallel testing strategies. In the presence of information asymmetry, the benefits of such a parallel approach may be outweighed by high agency costs. This result finds further support in practice. In recent years, testing costs have declined significantly owing to technological advancements in the realms of robotics, virtualization, and computer-assisted test systems (among others). Following conventional wisdom—and previous academic insights—these developments should have led firms to focus more strongly on parallel testing strategies. However, there is no substantial evidence to date that such a general trend is underway (though the software industry is a notable exception). Our results offer a plausible and straightforward explanation for this observation: firms are reluctant to incur the high costs of delegation that come with parallel testing.

4.6 Conclusions

Design testing is an integral part of virtually any new product development initiative because it enables firms to identify the best possible designs for their new products. In reality, however, managing such testing processes is a daunting challenge. The reason is that in most cases the firm does not itself conduct the desired testing activities and so has no direct accesses to the precious information; the firm must instead rely on the recommendations of experts, who may be pursuing their own agendas. This delegated nature of the testing process gives rise to information asymmetry between the firm and the experts, which can result in a worrisome misalignment of objectives. The primary goal of this paper is to understand how the firm can set up an effective testing process that will reliably select the best design alternative—that is, notwithstanding the adverse consequences of delegation. More precisely, our main contribution is to provide insights on the questions of (i) which testing strategy the firm should choose and (ii) how the firm can optimally incentivize the experts it hires.

It is remarkable that, regardless of the chosen testing strategy, the optimal compensation scheme—one that motivates experts to test their designs with high effort and to reveal their test outcomes truthfully—always involves but two payments: a success bonus if an expert's design is developed and turns out to be technologically feasible, and a reward that reimburses an expert for his efforts in case the firm dismisses the design he is testing or terminates the testing process altogether. We show in addition that the balance between these two payments is fundamentally different for designs with different levels of test efficiency. Designs that can be tested with high precision require a strong focus on individual success bonuses, whereas designs that are evaluated with lower quality demand a stronger emphasis on consolation awards.

Our findings have immediate managerial implications. Although design testing is a complex organizational process, the structure of the optimal compensation schemes is fairly simple. As a result, the optimal contracts derived in this paper should be relatively easy to implement in practice. Regarding the relative sizes of the two payments, we emphasize that the firm must carefully adjust its contracts to reflect the quality of test outcomes. Firms that adopt a one-size-fits-all approach cannot help but sacrifice, eventually, their testing effectiveness.

As for the firm's optimal testing strategy, we find that two parameters critically determine the firm's optimal choice: testing costs and the quality of information. In line with previous research, we show that a parallel testing strategy is optimal only when testing costs are sufficiently small. In contrast, the higher the testing costs, the *more* beneficial a sequential testing strategy becomes. Yet it is an unanswered question just how the firm should set up its sequential testing strategy. Should it mandate a single expert to carry out all test activities, or should it rather assign a different expert to each design test? Our results indicate that the former approach is optimal when the informational quality of the different design tests is relatively homogeneous. When test efficiency is very heterogenous across designs, however, the firm should hire multiple experts because in that case a single expert might artificially keep the testing process alive in order to receive ongoing payments—to the firm's obvious detriment.

These results have clear consequences for practice. We identify two levers the firm can use when designing an effective testing strategy: the order of the different design tests (parallel vs. sequential), and the number of experts to employ. Whereas the former option has been extensively discussed in the academic literature, the latter option has yet to receive serious attention. It is important to recognize that the two levers address two different concerns. The choice of whether to use a parallel or sequential testing strategy depends on the testing *costs*, whereas the ideal number of experts depends on the extent to which the *efficiency* of tests (for the different designs) is heterogeneous.

Finally, we show how the presence of information asymmetry affects the various testing strategies. Overall, our results point to the same conclusion: the delegation of testing leads to a suboptimal testing process whose information asymmetries are significantly more harmful to parallel than to sequential testing strategies. In other words, our findings indicate that parallel testing may be less effective than usually claimed in the academic literature (Dahan and Mendelson 2001, Loch et al. 2001) when the testing process involves information asymmetries. This finding may also explain the practical observation that, even though technological advancements have lowered testing costs in recent years, no significant shift toward parallel testing efforts is evident.

To maintain tractability and develop a parsimonious model, we necessarily made some assumptions about the specific trade-offs inherent to a sequential testing strategy. In particular, we assumed that the design alternatives are sufficiently different that the firm cannot exploit any between-design learning effects. Also, we did not allow the firm to choose previously tested designs for development. Both assumptions clearly lead to an underestimate of the performance of sequential testing strategies, from which it follows that relaxing these assumptions could only strengthen our main message that delegation favors sequential testing. Furthermore, one can readily verify that the optimal contract structures would remain relatively intact even without these assumptions; hence our results are applicable to a wide range of practical scenarios. With regard to the firm's choice of testing strategy, we focused on "polar" cases: fully parallel versus fully sequential testing, and a single expert versus n experts. In reality, firms are free to use any mixture of parallel and sequential testing strategies, and they may also hire any arbitrary number of experts. Whether such hybrid strategies can improve the efficacy of delegated testing processes is an important question for future research. Another interesting research possibility is for empirical studies to examine the relationship between a firm's chosen testing strategy and the severity of its agency issues. Our own theoretical results lead us to conjecture that firms (and industries) with relatively strong agency problems are much more likely to use sequential than parallel testing strategies.

[This page intentionally left blank.]

Appendix A

Proofs of Chapter II

Lemma A.1. *There exists a unique pure-strategy second-round equilibrium for any feedback policy.*

Proof of Lemma A.1. Observe that each Nash equilibrium satisfies $e_{i2} \in [0, \sqrt{A/c}]$ because $u_{i2}(0, e_{j2}) \geq 0 \geq u_{i2}(\sqrt{A/c}, e_{j2}) > u_{i2}(e_{i2}, e_{j2})$ for any e_{j2} and all $e_{i2} > \sqrt{A/c}$. Hence to prove existence and uniqueness of the second-round equilibrium, we can replace the original contest by a modified contest where each solver's effort choice is restricted to $[0, \sqrt{A/c}]$. These two contests have the same Nash equilibria because each Nash equilibrium satisfies $e_{i2} \in [0, \sqrt{A/c}]$, u_{i2} is continuous, and each solver's strategy space is an interval of $\mathbb{R}$. Also, straightforward differentiation shows that u_{i2} is strictly concave in e_{i2} for given e_{j2}. Hence, by Theorem 1.2 in Fudenberg and Tirole (1991), there exists a pure-strategy Nash equilibrium in the original contest.

It remains to show that this equilibrium is also unique. However, once we observe that $\partial^2 u_{i2}/\partial e_{i2}^2 + 2c = -\partial^2 u_{i2}/\partial e_{i2}\partial e_{j2}$, and $\partial \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2})|f_i\right]/\partial e_{i2} \in [-k_e/a^2, k_e/a^2]$ for any feedback f_i, this is just a simple application of Theorems 2 and 6 in Rosen (1965). □

Proof of Proposition 2.1. Without feedback between round one and two, solver i's optimization problem is equivalent to a utility maximization problem with simultaneous decisions on 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 $e_i^{\text{no}} \in \text{argmax}_{e_i} A \cdot \mathbb{E}_{\zeta_1}[G_{\Delta\zeta_2}(2k_e e_i + \zeta_{i1} - 2k_e e_j - \zeta_{j1})] - 2ce_i^2$. Since this is a strictly concave maximization problem, e_i^{no} is given by the following first-order optimality condition:

$$2Ak_e \cdot \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(2k_e e_i^{\text{no}} + \zeta_{i1} - 2k_e e_j^{\text{no}} - \zeta_{j1})\right] = 4ce_i^{\text{no}}. \tag{A.1}$$

By the symmetry of $g_{\Delta\zeta_2}$ around zero, it follows readily that the unique solution to the solvers' optimality conditions is symmetric; that is $e_i^{\text{no}} = e_j^{\text{no}}$. Inserting this information in (A.1) yields $e_i^{\text{no}} = Ak_e/(2c) \cdot \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(\zeta_{i1} - \zeta_{j1})\right] = Ak_e/(2c) \cdot \int_{-a}^{a} g_{\Delta\zeta_2}(z) g_{\Delta\zeta_1}(z) dz = Ak_e/(2c) \cdot \int_{-a}^{a} g_{\Delta\zeta_1}(z)^2 dz = Ak_e/(3ac)$. $\square$

Proof of Proposition 2.2. Given public feedback, solvers perfectly learn v_1 after round one. As such, solver i's second-round equilibrium effort solves $e_{i2}^{\text{pub}} \in \operatorname{argmax}_{e_{i2}} A \cdot G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2}) - ce_{i2}^2$, and the corresponding necessary and sufficient first-order optimality condition is given by

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

By the symmetry of $g_{\Delta\zeta_2}$ around zero, 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 $e_{i1}^{\text{pub}} \in \operatorname{argmax}_{e_{i1}} A \cdot \mathbb{E}_{\zeta_1}\left[G_{\Delta\zeta_2}(k_e e_{i1} + \zeta_{i1} - k_e e_{j1} - \zeta_{j1})\right] - ce_{i1}^2 - \mathbb{E}_{\zeta_1}\left[c(e_{i2}^{\text{pub}})^2\right]$, and the corresponding necessary first-order optimality condition is given by

$$Ak_e \cdot \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(k_e e_{i1}^{\text{pub}} + \zeta_{i1} - k_e e_{j1}^{\text{pub}} - \zeta_{j1})\right] - 2ce_{i1}^{\text{pub}} - \frac{\partial}{\partial e_{i1}}\mathbb{E}_{\zeta_1}\left[c(e_{i2}^{\text{pub}})^2\right] = 0. \tag{A.3}$$

Note that the first two terms in (A.3) 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 i's second-round effort e_{i2}^{pub}. In equilibrium, this indirect effect must be zero. To see this, note that (A.3) 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 (A.2) and (A.3) shows that the unique PBE under public feedback is given by $e_{i2}^{\text{pub}}(\Delta\zeta_1) = Ak_e/(2c) \cdot g_{\Delta\zeta_2}(\Delta\zeta_1) = Ak_e/(2a^2c) \cdot (a - |\Delta\zeta_1|)$, and $e_{i1}^{\text{pub}} = \mathbb{E}_{\zeta_1}\left[e_{i2}^{\text{pub}}(\Delta\zeta_1)\right] = Ak_e/(3ac)$. $\square$

Proof of Proposition 2.3. Given 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 $e_{i2}^{\text{pri}} \in \operatorname{argmax}_{e_{i2}} A \cdot \mathbb{E}_{v_{j1}}[G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2})|v_{i1}] - ce_{i2}^2$, and the corresponding necessary and sufficient first-order optimality condition is given by

$$Ak_e \cdot \mathbb{E}_{v_{j1}}[g_{\Delta\zeta_2}(v_{i1} + k_e e_{i2}^{\mathrm{pri}} - v_{j1} - k_e e_{j2}^{\mathrm{pri}})|v_{i1}] = 2ce_{i2}^{\mathrm{pri}}. \tag{A.4}$$

Lemma A.1 ensures that the second-round equilibrium defined by (A.4) is unique. Next, we establish the uniqueness of the first-round equilibrium. In the first round, solver i's equilibrium effort has to solve $e_{i1}^{\mathrm{pri}} \in \mathrm{argmax}_{e_{i1}} A \cdot \mathbb{E}_{\zeta_1}[G_{\Delta\zeta_2}(k_e(e_{i1} + e_{i2}^{\mathrm{pri}}) + \zeta_{i1} - k_e(e_{j1} + e_{j2}^{\mathrm{pri}}) - \zeta_{j1})] - ce_{i1}^2 - \mathbb{E}_{\zeta_1}\left[c(e_{i2}^{\mathrm{pri}})^2\right]$, and the corresponding necessary first-order optimality condition is given by

$$\begin{aligned} Ak_e \cdot \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(k_e(e_{i1}^{\mathrm{pri}} + e_{i2}^{\mathrm{pri}}) + \zeta_{i1} - k_e(e_{j1}^{\mathrm{pri}} + e_{j2}^{\mathrm{pri}}) - \zeta_{j1}) \cdot \left(1 + \frac{\partial e_{i2}^{\mathrm{pri}}}{\partial e_{i1}} - \frac{\partial e_{j2}^{\mathrm{pri}}}{\partial e_{i1}}\right)\right] - 2ce_{i1}^{\mathrm{pri}} \\ - \mathbb{E}_{\zeta_1}\left[2ce_{i2}^{\mathrm{pri}} \cdot \frac{\partial e_{i2}^{\mathrm{pri}}}{\partial e_{i1}}\right] = 0. \end{aligned} \tag{A.5}$$

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, $\partial e_{j2}^{\mathrm{pri}}/\partial e_{i1} = 0$. Rewriting (A.5) yields $Ak_e\mathbb{E}_{\zeta_1}[g_{\Delta\zeta_2}(k_e(e_{i1}^{\mathrm{pri}} + e_{i2}^{\mathrm{pri}}) + \zeta_{i1} - k_e(e_{j1}^{\mathrm{pri}} + e_{j2}^{\mathrm{pri}}) - \zeta_{j1})] - 2ce_{i1}^{\mathrm{pri}} + \mathbb{E}_{\zeta_1}[(Ak_e g_{\Delta\zeta_2}(k_e(e_{i1}^{\mathrm{pri}} + e_{i2}^{\mathrm{pri}}) + \zeta_{i1} - k_e(e_{j1}^{\mathrm{pri}} + e_{j2}^{\mathrm{pri}}) - \zeta_{j1}) - 2ce_{i2}^{\mathrm{pri}}) \cdot \partial e_{i2}^{\mathrm{pri}}/\partial e_{i1}] = 0$, where the third term is zero because $\mathbb{E}_{\zeta_1}[(Ak_e g_{\Delta\zeta_2}(k_e(e_{i1}^{\mathrm{pri}} + e_{i2}^{\mathrm{pri}}) + \zeta_{i1} - k_e(e_{j1}^{\mathrm{pri}} + e_{j2}^{\mathrm{pri}}) - \zeta_{j1}) - 2ce_{i2}^{\mathrm{pri}})\partial e_{i2}^{\mathrm{pri}}/\partial e_{i1}] = \mathbb{E}_{v_{i1}}[\mathbb{E}_{v_{j1}}[(Ak_e g_{\Delta\zeta_2}(k_e(e_{i1}^{\mathrm{pri}} + e_{i2}^{\mathrm{pri}}) + \zeta_{i1} - k_e(e_{j1}^{\mathrm{pri}} + e_{j2}^{\mathrm{pri}}) - \zeta_{j1}) - 2ce_{i2}^{\mathrm{pri}})\partial e_{i2}^{\mathrm{pri}}/\partial e_{i1}|v_{i1}]] = \mathbb{E}_{v_{i1}}[\partial e_{i2}^{\mathrm{pri}}/\partial e_{i1} \cdot \mathbb{E}_{v_{j1}}[Ak_e g_{\Delta\zeta_2}(k_e(e_{i1}^{\mathrm{pri}} + e_{i2}^{\mathrm{pri}}) + \zeta_{i1} - k_e(e_{j1}^{\mathrm{pri}} + e_{j2}^{\mathrm{pri}}) - \zeta_{j1}) - 2ce_{i2}^{\mathrm{pri}}|v_{i1}]] = \mathbb{E}_{v_{i1}}[\partial e_{i2}^{\mathrm{pri}}/\partial e_{i1} \cdot (Ak_e\mathbb{E}_{v_{j1}}[g_{\Delta\zeta_2}(v_{i1} + k_e e_{i2}^{\mathrm{pri}} - v_{j1} - k_e e_{j2}^{\mathrm{pri}}))|v_{i1}] - 2ce_{i2}^{\mathrm{pri}})] = 0$. 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 (A.4). Thus, the first-order optimality condition of solver i is $Ak_e \cdot \mathbb{E}_{\zeta_1}\left[g_{\Delta\zeta_2}(k_e(e_{i1}^{\mathrm{pri}} + e_{i2}^{\mathrm{pri}}) + \zeta_{i1} - k_e(e_{j1}^{\mathrm{pri}} + e_{j2}^{\mathrm{pri}}) - \zeta_{j1})\right] - 2ce_{i1}^{\mathrm{pri}} = 0$, 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; that is $e_{i1}^{\mathrm{pri}} = e_{j1}^{\mathrm{pri}}$. Moreover, $e_1^{\mathrm{pri}} = \mathbb{E}_{v_{i1}}\left[e_{i2}^{\mathrm{pri}}(v_{i1})\right] = \mathbb{E}_{\zeta_{i1}}\left[e_{i2}^{\mathrm{pri}}(\zeta_{i1})\right]$.

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 $e_{i2}^{\mathrm{pri}}(\zeta_{i1}) = e_2^{\mathrm{pri}}(\zeta_{i1})$ and $e_{j2}^{\mathrm{pri}}(\zeta_{j1}) = e_2^{\mathrm{pri}}(\zeta_{j1})$, and that $v^{\mathrm{pri}}(\zeta_{i1}) = \zeta_{i1} + k_e e_2^{\mathrm{pri}}(\zeta_{i1})$ increases in ζ_{i1}. We will demonstrate in retrospective that these claims are true. Together with (A.4), the above properties imply that the equilibrium effort function $e_2^{\text{pri}}(\cdot)$ solves the following integral equation: $Ak_e \cdot \mathbb{E}_{\zeta_{j1}}[g_{\Delta\zeta_2}(\zeta_{i1} + k_e e_2^{\text{pri}}(\zeta_{i1}) - \zeta_{j1} - k_e e_2^{\text{pri}}(\zeta_{j1}))|\zeta_{i1}] = 2ce_2^{\text{pri}}(\zeta_{i1})$, or equivalently,

$$Ak_e^2 \cdot \mathbb{E}_{\zeta_{j1}}[g_{\Delta\zeta_2}(v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}))|\zeta_{i1}] = 2c\left(v^{\text{pri}}(\zeta_{i1}) - \zeta_{i1}\right). \tag{A.6}$$

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 distinguish three cases:

(I) 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, (A.6) is given by

$$\begin{aligned} &\int_{-\frac{a}{2}}^{\zeta_{i1}} (a - v^{\text{pri}}(\zeta_{i1}) + v^{\text{pri}}(\zeta_{j1}))d\zeta_{j1} + \int_{\zeta_{i1}}^{\frac{a}{2}} (a + v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}))d\zeta_{j1} \\ &= 2a\kappa(v^{\text{pri}}(\zeta_{i1}) - \zeta_{i1}), \end{aligned} \tag{A.7}$$

and differentiating both sides with respect to ζ_{i1} leads to the first-order ordinary differential equation $\left(v^{\text{pri}}(\zeta_{i1})\right)' = a\kappa/(\zeta_{i1} + a\kappa)$, with canonical solution $v^{\text{pri}}(\zeta_{i1}) = a\kappa \ln((a\kappa + \zeta_{i1})/\gamma_3)$. 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 $[\zeta_u, \zeta_o] \subset [-a/2, a/2]$.

(II) For $\zeta_{i1} \in [-a/2, \zeta_u]$, (A.6) becomes

$$\begin{aligned} &\int_{-\frac{a}{2}}^{\zeta_{i1}} (a - v^{\text{pri}}(\zeta_{i1}) + v^{\text{pri}}(\zeta_{j1}))d\zeta_{j1} + \int_{\zeta_{i1}}^{v^{-1}(v^{\text{pri}}(\zeta_{i1})+a)} (a + v^{\text{pri}}(\zeta_{i1}) - v^{\text{pri}}(\zeta_{j1}))d\zeta_{j1} \\ &= 2a\kappa(v^{\text{pri}}(\zeta_{i1}) - \zeta_{i1}), \end{aligned} \tag{A.8}$$

and differentiating both sides with respect to ζ_{i1} leads to the differential equation $(v^{\text{pri}}(\zeta_{i1}))'[2\zeta_{i1} + 2a\kappa + a/2 - v^{-1}(v^{\text{pri}}(\zeta_{i1}) + a)] = 2a\kappa$, which is a Bernoulli equation in $v^{-1}(\cdot)$ whose implicit solution is

$$v^{-1}(v^{\text{pri}}) = Ce^{\frac{1}{a\kappa}v^{\text{pri}}} - a\left(\kappa + \frac{1}{4}\right) - \frac{1}{2a\kappa}e^{\frac{1}{a\kappa}v^{\text{pri}}} \int e^{-\frac{1}{a\kappa}v^{\text{pri}}} v^{-1}(v^{\text{pri}} + a)dv^{\text{pri}}. \tag{A.9}$$

(III) In a similar vein, we can show that for $\zeta_{i1} \in [\zeta_o, a/2]$, the implicit solution to (A.6) is given by

$$v^{-1}(v^{\text{pri}}) = C' e^{\frac{1}{a\kappa} v^{\text{pri}}} - a\left(\kappa - \frac{1}{4}\right) - \frac{1}{2a\kappa} e^{\frac{1}{a\kappa} v^{\text{pri}}} \int e^{-\frac{1}{a\kappa} v^{\text{pri}}} v^{-1}(v^{\text{pri}} - a) dv^{\text{pri}}. \tag{A.10}$$

Combining (A.9) and (A.10) allows us to derive closed-form solutions. With (A.10), (A.9) becomes $v^{-1}(v^{\text{pri}}) = Ce^{\frac{1}{a\kappa} v^{\text{pri}}} - a\left(\frac{3}{2}\kappa + \frac{1}{8}\right) - C' \frac{1}{2a\kappa} e^{\frac{1}{\kappa}} e^{\frac{1}{a\kappa} v^{\text{pri}}} v^{\text{pri}} + \left(\frac{1}{2a\kappa}\right)^2 e^{\frac{1}{a\kappa} v^{\text{pri}}} \iint e^{-\frac{1}{a\kappa} v^{\text{pri}}} v^{-1}(v^{\text{pri}}) d(v^{\text{pri}} + a) dv^{\text{pri}}$. 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 $v^{-1}(v^{\text{pri}}) = \gamma_4 e^{\frac{3}{2a\kappa} v^{\text{pri}}} + \gamma_5 e^{\frac{1}{2a\kappa} v^{\text{pri}}} - a\left(\frac{1}{6} + 2\kappa\right)$. In an identical way, we can derive the canonical solution for $\zeta_{i1} \in [\zeta_o, a/2]$: $v^{-1}(v^{\text{pri}}) = \gamma_1 e^{\frac{3}{2a\kappa} v^{\text{pri}}} + \gamma_2 e^{\frac{1}{2a\kappa} v^{\text{pri}}} + a\left(\frac{1}{6} - 2\kappa\right)$. Thus, the canonical solution to (A.6) is given by

$$v^{-1}(v^{\text{pri}}) = \begin{cases} \gamma_4 e^{\frac{3}{2a\kappa} v^{\text{pri}}} + \gamma_5 e^{\frac{1}{2a\kappa} v^{\text{pri}}} - a\left(\frac{1}{6} + 2\kappa\right) & \text{if } v_o - a \leq v^{\text{pri}} < v_u \\ \gamma_3 e^{\frac{1}{a\kappa} v^{\text{pri}}} - a\kappa & \text{if } v_u \leq v^{\text{pri}} \leq v_o \\ \gamma_1 e^{\frac{3}{2a\kappa} v^{\text{pri}}} + \gamma_2 e^{\frac{1}{2a\kappa} v^{\text{pri}}} + a\left(\frac{1}{6} - 2\kappa\right) & \text{if } v_o < v^{\text{pri}} \leq v_u + a, \end{cases} \tag{A.11}$$

where $v_u = v^{\text{pri}}(\zeta_u)$, $v_o = v^{\text{pri}}(\zeta_o)$, $v_o - a = v^{\text{pri}}(-a/2)$, and $v_u + a = v^{\text{pri}}(a/2)$. With the substitution $u = \exp(v/(2a\kappa))$ we can represent $v^{-1}(v^{\text{pri}})$ as a set of cubic equations, which we can solve for $v^{\text{pri}}(\zeta_{i1})$ with standard mathematical tools (Olver et al. 2010, p. 131) to gain

$$v^{\text{pri}}(\zeta_{i1}) = \begin{cases} 2a\kappa \cdot \ln\left(\sqrt{-\frac{4\gamma_5}{3\gamma_4}} \cdot \sin\left(\frac{1}{3} \cdot \sin^{-1}\left(\sqrt{-3\frac{\gamma_4}{\gamma_5}} \cdot \frac{a(1+12\kappa)+6\zeta_{i1}}{4\gamma_5}\right)\right)\right) & \text{if } -\frac{a}{2} \leq \zeta_{i1} < \zeta_u \\ a\kappa \cdot \ln\left(\frac{\zeta_{i1}+a\kappa}{\gamma_3}\right) & \text{if } \zeta_u \leq \zeta_{i1} \leq \zeta_o \\ 2a\kappa \cdot \ln\left(\frac{1}{6} \cdot \sqrt[3]{z(\zeta_{i1})} - 2 \cdot \frac{\gamma_2}{\gamma 1} \cdot \frac{1}{\sqrt[3]{z(\zeta_{i1})}}\right) & \text{if } \zeta_o < \zeta_{i1} \leq \frac{a}{2}. \end{cases} \tag{A.12}$$

It remains to determine the integration constants. From (II), it follows readily that $\gamma_4 = -\gamma_1 n^3$ and $\gamma_5 = \gamma_2 n$. Moreover, (A.6) satisfies all requirements of the Implicit Function Theorem. Therefore, $v^{\text{pri}}(\zeta_{i1})$ is continuously differentiable. From the continuity of $\left(v^{\text{pri}}(\zeta_{i1})\right)'$, it follows that $\gamma_2 n = 2\gamma_3 xy(n^3 x + y)/(n^2 x^2 + y^2)$, and $3\gamma_1 n = 2\gamma_3(ny - x)/(n^2x^2 + y^2)$. Additionally, the continuity of $v^{\text{pri}}(\zeta_{i1})$ implies that $\zeta_u = \gamma_3 x^2 - a\kappa$, $\zeta_o = \gamma_3 y^2 - a\kappa$, $\gamma_3 = 3a(\kappa + 1/6)(n^2x^2 + y^2)/(x^2(3y^2 - n^2x^2 + 4n^3xy))$, 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 (2.6). Finally, the integral equation (A.6) becomes $a/(6\kappa)-a(\kappa-1/6)\ln(y)-a(\kappa+1/6)\ln(x)-a\kappa+2\gamma_1((y/n)^3-x^3)/3+\gamma_3(y^2+x^2)/2=0$, which is the same as (2.7). 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 $\gamma_1,\gamma_2,\gamma_3>0$. Thus, it is obvious that $v^{\text{pri}}(\zeta_{i1})$ increases in ζ_{i1} for $\zeta_{i1}\geq\zeta_u$. For $\zeta_{i1}<\zeta_u$, we have $(v^{-1}(v^{\text{pri}}))'>0$ if $3\gamma_4x^2+\gamma_5=2\gamma_3x>0$, which is true. Therefore, $v^{\text{pri}}(\zeta_{i1})$ increases in ζ_{i1}, which concludes the proof.

Note on our solution methodology: The crucial step is to transform the integral equation (A.6) into an ordinary differential equation (ODE) by differentiating both sides of the equality with respect to ζ_{i1}. Clearly, the unique solution to (A.6) also solves the ODE. However, the ODE may have solutions that do not solve (A.6). To circumvent this problem, we identify the ODE's canonical solution (A.11), which defines the solution to (A.6) up to certain constants. These constants are, in turn, uniquely defined by the properties of (A.6). □

Proof of Corollary 2.1. 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 (2.6)-(2.7) as $\kappa\to\infty$. Inserting $\widetilde{x}$ and $\widetilde{y}$ in (2.6) 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 (2.7) 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 (2.5), 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 2.3, $\lim_{\kappa\to\infty}\widetilde{\gamma}_3=\lim_{\kappa\to\infty}\gamma_3$. Taken together, this implies 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 Proposition 2.4. (i) From Propositions 2.1 and 2.2, 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 (A.4) 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 $\partial e_1^{\text{pri}}/\partial a=-e_1^{\text{pri}}/a+\partial(ae_1^{\text{pri}})/\partial a>-e_1^{\text{pri}}/a$. 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 2.1 - 2.3.

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

Proof of Proposition 2.5. (i) The result follows directly from comparing the firm's expected profits under the two different feedback policies: $\Pi_{\text{best}}^{\text{no}} = \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2} + 2k_e e_i^{\text{no}}\}\right] - A = 2k_e e_1^{\text{no}} + \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2}\}\right] - A = a(2/(3\kappa) + 7/30) - A$; and $\Pi_{\text{best}}^{\text{pub}} = \mathbb{E}[\max_i\{\zeta_{i1} + \zeta_{i2} + k_e e_1^{\text{pub}} + k_e e_2^{\text{pub}}(\zeta_1)\}] - A = 2k_e e_1^{\text{pub}} + \mathbb{E}\left[\max_i\{\zeta_{i1} + \zeta_{i2}\}\right] - A = \Pi_{\text{best}}^{\text{no}}$, where we made use of the well-known fact 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 $\Pi_{\text{best}}^{\text{pri}} = k_e e_1^{\text{pri}} + \mathbb{E}[\max_i\{\zeta_{i1} + \zeta_{i2} + k_e e_2^{\text{pri}}(\zeta_{i1})\}] - A$. 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 (A.7) 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 (A.7) is extended by negative terms compared to the correct solution outlined in Proposition 2.3. 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 $\underline{e}_2(\zeta_{i1})$, $\underline{\Pi}_{\text{best}}^{\text{pri}} = a \cdot (-\kappa^3(\kappa^2 + \frac{1}{4})(e^{\frac{1}{\kappa}} - e^{-\frac{1}{\kappa}}) + \kappa^4(e^{\frac{1}{\kappa}} + e^{-\frac{1}{\kappa}}) + \kappa(\kappa^2 - 2\kappa + \frac{1}{4})\ln(\frac{2\kappa-1}{2\kappa+1}) + \frac{1}{2}\kappa\ln(a^4(2\kappa-1)(2\kappa+1)^3) + \frac{5}{6}\kappa^2 - 2\kappa(1 + \ln(2) + \ln(\gamma_3)) + \frac{1}{12})$, and $\underline{\Pi}_{\text{best}}^{\text{pri}} > \Pi_{\text{best}}^{\text{pub}}$

if and only if $\gamma_3 < \underline{\gamma}_3$, with

$$\begin{aligned}\underline{\gamma}_3 = &\frac{a}{2} e^{-\frac{\kappa^2}{2}\left(\kappa^2+\frac{1}{4}\right)\left(e^{\frac{1}{\kappa}}-e^{-\frac{1}{\kappa}}\right)} e^{\frac{\kappa^3}{2}\left(e^{\frac{1}{\kappa}}+e^{-\frac{1}{\kappa}}\right)} \\ &\cdot \left(\frac{2\kappa-1}{2\kappa+1}\right)^{\frac{1}{2}\left(\kappa^2-2\kappa+\frac{1}{4}\right)} \sqrt[4]{(2\kappa-1)(2\kappa+1)^3} e^{\frac{5}{12}\kappa-1-\frac{3}{40\kappa}-\frac{1}{3\kappa^2}}.\end{aligned} \tag{A.13}$$

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 $\Gamma(x, y) = 2\gamma_3/(a(1+6\kappa)) = (n^2x^2 + y^2)/(x^2(3y^2 - n^2x^2 + 4n^3xy))$, which decreases in x and y. 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:

$$y(x, \Gamma) = \frac{nx}{1 - 3\Gamma x^2} \cdot \left(2n^2\Gamma x^2 - \sqrt{4(1 + n^4)\Gamma^2x^4 - (\Gamma x^2 - 1)^2}\right). \tag{A.14}$$

Inserting (A.14) in (2.6) and (2.7) 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; but before doing so, we derive some important properties.

Let $l_1(x, y)$ be the left-hand side of (2.6), and $l_2(x, y)$ be the left-hand side of (2.7). 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 $x_1(y)$ (resp. $x_2(y)$) the solution to $l_1(x_1(y), y) = 0$ (resp. $l_2(x_2(y), y) = 0$) 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 $x'_1(\Gamma)$ (resp. $x'_2(\Gamma)$) be the solution to $l'_1(x'_1(\Gamma), \Gamma) = 0$ (resp. $l'_2(x'_2(\Gamma), \Gamma) = 0$) 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 $l_2^{\text{Taylor}}(\overline{x}, y(\overline{x}, \underline{\Gamma}_3)) = 11/(3360\kappa^2) - 23/(960\kappa^3) + \mathcal{O}(1/\kappa^4)$. 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 $\kappa > \overline{\kappa}$. □

Proof of Proposition 2.6. (i) Suppose that both solvers have invested arbitrary first-round efforts e_1. Each solver will make his submission decision so as to maximize his expected continuation utility.

Case (a): If both solvers submit their intermediate solutions, they perfectly learn v_1. According to (A.2), each solver will invest a second-round effort of $e_2^{ss}(v_1) = Ak_e/2c \cdot g_{\Delta\zeta_2}(v_{i1} - v_{j1})$. As a result, each solver's expected continuation utility before submission is $u_{ic}^{ss} = A \cdot \mathbb{E}_{v_1}[G_{\Delta\zeta_2}(v_{i1} - v_{j1})] - \mathbb{E}_{v_1}[ce_2^{ss}(v_1)^2]$.

Case (b): If only solver i submits his intermediate solution, both solvers only learn v_{i1}. The solvers' equilibrium second-round effort is $e_2^{sn}(v_{i1}) = Ak_e/2c \cdot \mathbb{E}_{v_{j1}}[g_{\Delta\zeta_2}(v_{i1} - v_{j1})|v_{i1}]$, and each solver's expected continuation utility is $u_{ic}^{sn} = A \cdot \mathbb{E}_{v_1}[G_{\Delta\zeta_2}(v_{i1} - v_{j1})] - \mathbb{E}_{v_{i1}}[ce_2^{sn}(v_{i1})^2]$.

Case (c): If no solver submits his solution, then no feedback is transmitted and according to (A.1), the solvers' equilibrium second-round effort is $e_2^{nn} = Ak_e/2c \cdot \mathbb{E}_{v_1}[g_{\Delta\zeta_2}(v_{i1} - v_{j1})]$, which yields the following expected continuation utility for solver i: $u_{ic}^{nn} = A \cdot \mathbb{E}_{v_1}[G_{\Delta\zeta_2}(v_{i1} - v_{j1})] - c(e_2^{nn})^2$.

By symmetry, case (c) is an equilibrium if and only if $u_{ic}^{nn} > u_{ic}^{sn}$, or equivalently, $\mathbb{E}_{v_{i1}}[e_2^{sn}(v_{i1})^2] > (e_2^{nn})^2$. This is true because Jensen's inequality implies that

$\mathbb{E}_{v_{i1}}[e_2^{sn}(v_{i1})^2] > (\mathbb{E}_{v_{i1}}[e_2^{sn}(v_{i1})])^2 = (e_2^{nn})^2$. To show that case (c) is the unique equilibrium we verify that case (a) is not an equilibrium: $u_{jc}^{sn} > u_{jc}^{ss}$ follows from Jensen's inequality because $\mathbb{E}_{v_{i1}}[e_2^{sn}(v_{i1})^2] = \mathbb{E}_{v_{i1}}[\mathbb{E}_{v_{j1}}[e_2^{ss}(v_1)|v_{i1}]^2] < \mathbb{E}_{v_{i1}}[\mathbb{E}_{v_{j1}}[e_2^{ss}(v_1)^2|v_{i1}]] = \mathbb{E}_{v_1}[e_2^{ss}(v_1)^2]$.

(ii) In the case of private feedback, a solver's submission decision is unobservable to the other solver. Hence it is sufficient to verify that a solver's expected utility from submitting is larger than from not submitting for any mixed strategy of solver j and arbitrary e_1. Let q_j be solver j's probability of submitting, and let e_{j2}^s and e_{j2}^n be his second-round efforts if he submits or not, respectively (Lemma A1 guarantees uniqueness).

Solver i's expected utility when submitting his intermediate solution and receiving feedback v_{i1} is $u_{i2}^s(e_{i2}, v_{i1}) = q_j(A \cdot \mathbb{E}_{v_{j1}}[G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2}^s)|v_{i1}] - ce_{i2}^2) + (1 - q_j)(A \cdot \mathbb{E}_{v_{j1}}[G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2}^n)|v_{i1}] - ce_{i2}^2)$. Let $u_{i2}^s(e_{i2}) = \mathbb{E}_{v_{i1}}[u_{i2}^s(e_{i2}, v_{i1})]$ be his corresponding expected continuation utility, and note that his expected continuation utility from not submitting is $u_{i2}^n(e_{i2}) = q_j(A \cdot \mathbb{E}_{v_1}[G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2}^s)] - ce_{i2}^2) + (1 - q_j)(A \cdot \mathbb{E}_{v_1}[G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2}^n)] - ce_{i2}^2)$. Furthermore, let $e_{i2}^s(v_{i1})$ and e_{i2}^n be solver i's optimal effort choices if he submits or not, respectively, his intermediate solution. Then it is true that $u_{i2}^s(e_{i2}^s(v_{i1})) = \mathbb{E}_{v_{i1}}[u_{i2}^s(e_{i2}^s(v_{i1}), v_{i1})] > \mathbb{E}_{v_{i1}}[u_{i2}^s(e_{i2}^n, v_{i1})] = u_{i2}^n(e_{i2}^n)$, which proves the claim. □

Proof of Proposition 2.7. Let $A_1 = \alpha A$ and $A_2 = (1-\alpha)A$, $\alpha \in [0,1]$, be the awards for the first- and second-round winner, respectively, and note that Proposition 2.6 implies that α must be sufficiently large to incentivize solvers to submit their intermediate solutions. Clearly, if α is not large enough then a milestone award only reduces the overall contest incentives, which cannot be optimal. In contrast, if α is large enough, then solvers submit their intermediate solutions. In this case, following the same steps as in the proof of Proposition 2.2, it is straightforward to show that the unique PBE is symmetric and that the solvers' equilibrium efforts are $e_1^{\text{int}} = (1 + \alpha/2)e_1^{\text{pub}}$ and $e_2^{\text{int}}(\Delta\zeta_1) = (1-\alpha)e_2^{\text{pub}}(\Delta\zeta_1)$. It follows that $e_1^{\text{int}} + \mathbb{E}[e_2^{\text{int}}(\Delta\zeta_1)] = (2-\alpha/2)e_1^{\text{pub}}$, implying that the firm's expected profits decrease in α. Thus the firm always chooses $\alpha = 0$ in optimum. □

Proof of Proposition 2.8. Suppose that after round one, the firm truthfully reveals the solvers' ranking, but not v_1; and wlog suppose that solver i is currently the leader. Hence both solvers know that $\Delta v_1 = v_{i1} - v_{j1} > 0$. Solver i's second-round equilibrium effort solves $e^{\mathrm{r}}_{i2} \in \operatorname{argmax}_{e_{i2}} A \cdot \mathbb{E}_{v_1}[G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2})|\Delta v_1 > 0] - ce^2_{i2}$, and the corresponding necessary and sufficient first-order optimality condition is $Ak_e \cdot \mathbb{E}_{v_1}[g_{\Delta\zeta_2}(v_{i1} + k_e e^{\mathrm{r}}_{i2} - v_{j1} - k_e e^{\mathrm{r}}_{j2})|\Delta v_1 > 0] = 2ce^{\mathrm{r}}_{i2}$. By the symmetry of $g_{\Delta\zeta_2}$ around zero, the unique second-round equilibrium is symmetric, $e^{\mathrm{r}}_{i2} = e^{\mathrm{r}}_{j2}$. In the first round, solver i's equilibrium effort solves $e^{\mathrm{r}}_{i1} \in \operatorname{argmax}_{e_{i1}} A \cdot \mathbb{E}_{\zeta_1}[G_{\Delta\zeta_2}(\zeta_{i1} + k_e e_{i1} - \zeta_{j1} - k_e e_{j1})] - ce^2_{i1} - \mathbb{E}_{\zeta_1}[c(e^{\mathrm{r}}_{i2})^2]$, with necessary optimality condition $Ak_e \cdot \mathbb{E}_{\zeta_1}[g_{\Delta\zeta_2}(\zeta_{i1} + k_e e^{\mathrm{r}}_{i1} - \zeta_{j1} - k_e e^{\mathrm{r}}_{j1})] - 2ce^{\mathrm{r}}_{i1} - \frac{\partial}{\partial e_{i1}} \mathbb{E}_{\zeta_1}[c(e^{\mathrm{r}}_{i2})^2] = 0$. By the very same argument as in the proof of Proposition 2.2, the first-round equilibrium is also unique and symmetric; that is $e^{\mathrm{r}}_{i1} = e^{\mathrm{r}}_{j1}$. Hence, $e^{\mathrm{r}}_1 = Ak_e/(2c) \cdot \mathbb{E}_{\Delta\zeta_1}[g_{\Delta\zeta_2}(\Delta\zeta_1)] = Ak_e/(3ac)$ and $e^{\mathrm{r}}_2 = Ak_e/(2c) \cdot \mathbb{E}_{\Delta\zeta_1}[g_{\Delta\zeta_2}(\Delta\zeta_1)|\Delta\zeta_1 > 0]$. Since $g_{\Delta\zeta_1|\Delta\zeta_1>0}(u) = 2g_{\Delta\zeta_1}(u)$ for $u > 0$ and 0 otherwise, it follows that $e^{\mathrm{r}}_2 = Ak_e/(3ac)$. □

Proof of Proposition 2.9. (i) Upon learning v_1 under a noisy public-feedback policy, solver i chooses his second-round equilibrium effort by solving $e^{\mathrm{noi}}_{i2} \in \operatorname{argmax}_{e_{i2}} qA \cdot G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2}) + (1 - q)A \cdot \mathbb{E}_{v_1}[G_{\Delta\zeta_2}(v_{i1} + k_e e_{i2} - v_{j1} - k_e e_{j2})] - ce^2_{i2}$. Applying Lemma A1 shows that the second-round equilibrium is unique and symmetric. Hence, by the same argument as in the proof of Proposition 2.2, the first-round equilibrium is also unique and symmetric. In particular, equilibrium efforts are given by $e^{\mathrm{noi}}_1 = \mathbb{E}_{\Delta\zeta_1}[e^{\mathrm{noi}}_2(\Delta\zeta_1)] = Ak_e/(3ac)$, and $e^{\mathrm{noi}}_2(\Delta\zeta_1) = Ak_e/(2c) \cdot (qg_{\Delta\zeta_2}(\Delta\zeta_1) + (1 - q)\mathbb{E}_{\Delta\zeta_1}[g_{\Delta\zeta_2}(\Delta\zeta_1)])$. Finally, using the same methodology as in the proofs of Proposition 4A(ii) and 4B(i) shows that Π^{pub}_{avg} and Π^{pub}_{best} are invariant in q.

(ii) In a way identical to the proof of Proposition 2.3, it can be shown that the equilibrium under a noisy private-feedback policy is unique and symmetric. Moreover, $e^{\mathrm{noi}}_1 = \mathbb{E}_{\zeta_{i1}}[e^{\mathrm{noi}}_2(\zeta_{i1})]$, and $e^{\mathrm{noi}}_2(\zeta_{i1})$ solves $q\mathbb{E}_{\zeta_{j1}}[g_{\Delta\zeta_2}(\zeta_{i1} + k_e e^{\mathrm{noi}}_2(\zeta_{i1}) - \zeta_{j1} - k_e e^{\mathrm{noi}}_2(\zeta_{j1}))] + (1 - q)\mathbb{E}_{\zeta_1}[g_{\Delta\zeta_2}(\zeta_{i1} + k_e e^{\mathrm{noi}}_2(\zeta_{i1}) - \zeta_{j1} - k_e e^{\mathrm{noi}}_2(\zeta_{j1}))] = 2ce^{\mathrm{noi}}_2(\zeta_{i1})/(Ak_e)$. Let $e^{pri}_2(\zeta_{i1}, \kappa)$ be the effort function defined in (2.5) for given κ. Then, $e^{\mathrm{noi}}_2(\zeta_{i1}) = e^{pri}_2(\zeta_{i1}, \kappa/q) + (1 - q)Ak_e/(2c) \cdot \mathbb{E}_{\zeta_1}[g_{\Delta\zeta_2}(\zeta_{i1} + k_e e^{pri}_2(\zeta_{i1}, \kappa/q) - \zeta_{j1} - k_e e^{pri}_2(\zeta_{i1}, \kappa/q))]$. Having derived the equilibrium efforts we can follow exactly the same procedure as in the proof of Proposition 4B(iii) to gain the required Taylor series $l^{\mathrm{Taylor}}_2(\overline{x}, y(\overline{x}, \underline{\Gamma}_3)) = 11q/(3360\kappa^2) - 23/(960\kappa^3) + \mathcal{O}(1/\kappa^4)$. Since the first term is positive: for any fixed $q > 0$, there exists a $\overline{\kappa} < \infty$ such that $l^{\mathrm{Taylor}}_2(\overline{x}, y(\overline{x}, \underline{\Gamma}_3)) > 0$ for all $\kappa > \overline{\kappa}$. □

[This page intentionally left blank.]

Appendix B

Proofs of Chapter III

Appendix B contains three different parts. Section B.1 discusses some preliminary technical results which are essential for our analysis. Section B.2 characterizes the Perfect Bayesian equilibria of the different contest formats. Last, Section B.3 provides the detailed proofs for our mathematical results.

B.1 Technical Preliminaries

The Skew-Normal Distribution.

Following Azzalini (1985), we refer to a random variable X with probability density function $\psi(x;\alpha) = 2\Phi(\alpha x)\phi(x)$, $x \in \mathbb{R}$, as a Skew-Normal random variable with parameter $\alpha \in \mathbb{R}$. The cumulative distribution function of X is given by $\Psi(x;\alpha) = \int_{-\infty}^{x} \psi(y;\alpha)dy$. Below we derive some important properties of Skew-Normal random variables.

Lemma B.1. (i) *$\psi(x;0) = \phi(x)$ for all $x \in \mathbb{R}$.*

(ii) *Define $\mathcal{I}_k(\alpha) \equiv \int_{-\infty}^{\infty} \Psi(x;\alpha)^k \psi(x;\alpha)^2 dx$ for $k \geq 0$. Then, $\mathcal{I}_k(\alpha)$ strictly decreases in α for $\alpha < 0$ and strictly increases for $\alpha > 0$.*

(iii) *$(k+1)(k+2)\mathcal{I}_k(0) = \mu^{(k+2)}$ for all $k \geq 0$.*

Proof of Lemma B.1. (i) See Property A in Azzalini (1985).

(ii) Taking the first-order derivative and exploiting properties of the Normal distribution yields

$$\frac{d\mathcal{I}_k(\alpha)}{d\alpha} = \frac{2\alpha}{\pi(1+\alpha^2)} \int_{-\infty}^{+\infty} \phi(\sqrt{2+\alpha^2}x)\Psi(x;\alpha)^k \left(\alpha x \Phi(\alpha x) + \phi(\alpha x)\right) dx. \tag{B.1}$$

Since $\alpha x \Phi(\alpha x) + \phi(\alpha x) > 0$ for all $x, \alpha \in \mathbb{R}$, we have $d\mathcal{I}_k(\alpha)/d\alpha < 0$ for $\alpha < 0$ and $d\mathcal{I}_k(\alpha)/d\alpha > 0$ for $\alpha > 0$.

(iii) $\mathcal{I}_k(0) = \int_{-\infty}^{\infty}(k+1)\Phi(x)^k\phi(x)^2dx/(k+1) = \int_{-\infty}^{\infty} x\phi(x)\Phi(x)^{k+1}dx/(k+1) = \int_{-\infty}^{\infty} x(k+2)\phi(x)\Phi(x)^{k+1}dx/((k+1)(k+2)) = \mu^{(k+2)}/((k+1)(k+2))$. □

A Generalization of Slepian's Inequality.

In this section, we establish a generalization of Slepian's Inequality, whose basic formulation can be found in, e.g., Theorem 2.1.1. in Tong (1980).

Lemma B.2. *Let* $X = (X_1, X_2)$*,* $Y = (Y_1, Y_2)$ *follow a bivariate Normal distribution with marginal distributions* $X_1, Y_1 \sim N(\mu_1, \sigma_1^2)$*,* $X_2, Y_2 \sim N(\mu_2, \sigma_2^2)$ *and correlation* ρ_X*,* ρ_Y*, where* $\rho_Y > \rho_X$*. Then,*

(i) $\mathbb{P}(X_1 \leq u_1, X_2 \leq u_2) < \mathbb{P}(Y_1 \leq u_1, Y_2 \leq u_2)$ *for all* $u_1, u_2 \in \mathbb{R}$*.*

(ii) $\mathbb{P}(X_1 \geq u_1, X_2 \geq u_2) < \mathbb{P}(Y_1 \geq u_1, Y_2 \geq u_2)$ *for all* $u_1, u_2 \in \mathbb{R}$*.*

Proof of Lemma B.2. (i) $\mathbb{P}(X_1 \leq u_1, X_2 \leq u_2) = \mathbb{P}(\overline{X}_1 \leq u_1 - \mu_1, \overline{X}_2 \leq u_2 - \mu_2) < \mathbb{P}(\overline{Y}_1 \leq u_1 - \mu_1, \overline{Y}_2 \leq u_2 - \mu_2) = \mathbb{P}(Y_1 \leq u_1, Y_2 \leq u_2)$ for all $u_1, u_2 \in \mathbb{R}$, where the strict inequality follows from Theorem 2.1.1. in Tong (1980) and the fact that $\overline{X}_j, \overline{Y}_j$, $j = 1, 2$, are centered Normal random variables.

(ii) This result is an immediate consequence of part (i). □

B.2 Derivation of Perfect Bayesian Equilibrium

In this section, we derive symmetric Perfect Bayesian equilibria (PBE) for the most general version of our model; that is, $k \geq 1$ and $\zeta_i = (\zeta_{i1}, \zeta_{i2})$ follows a bivariate Normal distribution with correlation $\rho \in [0, 1)$ and marginal distributions $\zeta_{i1} \sim N(0, \sigma^2)$ and $\zeta_{i2} \sim N(0, k^2\sigma^2)$. We start with the case of technological substitutes and then proceed to technological complements.

B.2.1. PBE for Technological Substitutes

Note that given our assumption of a sufficiently large performance shock (i.e., $\sigma > \sigma_{sub}$), each supplier i always participates in any contest. This is true because supplier i can always guarantee himself a strictly positive expected utility by participating and exerting zero effort. Hence we can derive the suppliers' equilibrium efforts by solving their incentive compatibility (IC) constraints.

System contest.

In a system contest, supplier i's IC constraint is:

$$\begin{aligned}(e_{i1}^{sys}, e_{i2}^{sys}) &\in \operatorname*{argmax}_{e_{i1},e_{i2}} u_i(e_{i1}, e_{i2}) \\ &\equiv A\mathbb{E}_{\zeta_{i1},\zeta_{i2}}\left[\prod_{k\neq i} F_{\zeta_{k1}+\zeta_{k2}}\left(\sum_{j\in\{1,2\}} r(e_{ij}) + \zeta_{ij} - r(e_{kj})\right)\right] - c(e_{i1}) - c(e_{i2}).\end{aligned} \tag{B.2}$$

Using straightforward differentiation and the law of iterated expectations, we find that any symmetric pure-strategy PBE satisfies the following optimality condition for $j = 1, 2$:

$$\frac{c'(e_j^{sys})}{r'(e_j^{sys})} = \frac{A(n-1)\mathcal{I}_{n-2}(0)}{\sigma\sqrt{1+k^2+2k\rho}} = \frac{A\mu^{(n)}}{n\sigma\sqrt{1+k^2+2k\rho}}, \tag{B.3}$$

where $\mu^{(n)}$ is the expected value of the maximum order statistic of n standard Normal random variables. Since c'/r' is strictly increasing, $c'(0) = 0$ and $c'(\infty)/r'(\infty) = \infty$, it follows that (e_1^{sys}, e_2^{sys}) is the unique solution to (B.3). Define $\eta(x) = (r \circ (c'/r')^{-1})(x)$ for all $x \geq 0$. Then the buyer's equilibrium expected profit is

$$\Pi_{sub}^{sys} = 2\eta\left(\frac{A\mu^{(n)}}{n\sigma\sqrt{1+k^2+2k\rho}}\right) + \sigma\mu^{(n)}\sqrt{1+k^2+2k\rho}. \tag{B.4}$$

Component contest.

Given p, supplier i's IC constraint in each component contest $j \in \{1, 2\}$ is:

$$e_{i1}^{cpo} \in \operatorname*{argmax}_{e_{i1}} u_i(e_{i1}) \equiv pA\mathbb{E}_{\zeta_{i1}}\left[\prod_{k\neq i} F_{\zeta_{k1}}\left(r(e_{i1}) + \zeta_{i1} - r(e_{k1})\right)\right] - c(e_{i1}) \tag{B.5}$$

$$e_{i2}^{cpo} \in \operatorname*{argmax}_{e_{i2}} u_i(e_{i2}) \equiv (1-p)A\mathbb{E}_{\zeta_{i2}}\left[\prod_{k\neq i} F_{\zeta_{k2}}\left(r(e_{i2}) + \zeta_{i2} - r(e_{k2})\right)\right] - c(e_{i2}), \tag{B.6}$$

yielding the following optimality conditions for a symmetric PBE:

$$\frac{c'(e_1^{cpo})}{r'(e_1^{cpo})} = \frac{pA(n-1)\mathcal{I}_{n-2}(0)}{\sigma} = \frac{pA\mu^{(n)}}{n\sigma} \tag{B.7}$$

$$\frac{c'(e_2^{cpo})}{r'(e_2^{cpo})} = \frac{(1-p)A(n-1)\mathcal{I}_{n-2}(0)}{k\sigma} = \frac{(1-p)A\mu^{(n)}}{nk\sigma}. \tag{B.8}$$

Again, since c'/r' is strictly increasing, $c'(0) = 0$ and $c'(\infty)/r'(\infty) = \infty$, (e_1^{cpo}, e_2^{cpo}) is the unique symmetric PBE. Given suppliers' equilibrium efforts, the buyer chooses p to maximize expected profits:

$$p^* \in \operatorname*{argmax}_p \Pi^{cpo}(p) = \eta\left(\frac{pA\mu^{(n)}}{n\sigma}\right) + \eta\left(\frac{(1-p)A\mu^{(n)}}{nk\sigma}\right) + \sigma\mu^{(n)}(1+k). \tag{B.9}$$

Assumption 3.2 ensures that η is a strictly concave function, and therefore $\Pi^{cpo}(p)$ is strictly concave in p, implying that p^* is unique. Moreover, the necessary and sufficient first-order condition $\partial\Pi^{cpo}(p)/\partial p = 0$ reveals that $p^*(k=1) = 1/2$ and $\lim_{k\to\infty} p^*(k) = 1$, and we let $\Pi_{sub}^{cpo} \equiv \Pi^{cpo}(p^*)$.

B.2.2. PBE for Complementary Components

As before, the assumption that $\sigma > \sigma_{cml}$ ensures that each supplier i always participates in any contest. Hence we can derive the suppliers' equilibrium efforts by solving their incentive compatibility (IC) constraints.

System contest.

In a system contest, supplier i's IC constraint is:

$$(e_{i1}^{sys}, e_{i2}^{sys}) \in \operatorname*{argmax}_{e_{i1},e_{i2}} u_i(e_{i1}, e_{i2}) \equiv A\mathbb{E}_{\zeta_{i1},\zeta_{i2}}\left[\prod_{k\neq i} F_{s_k}(s_i)\right] - c(e_{i1}) - c(e_{i2}), \tag{B.10}$$

where $s_i = \min\{v_{i1}, v_{i2}\}$, and since $v_{i1} \sim N(r(e_{i1}), \sigma^2)$ and $v_{i2} \sim N(r(e_{i2}), k^2\sigma^2)$ we have

$$f_{s_i}(x) = \sum_{j=1}^{2} \frac{1}{k^{j-1}\sigma}\phi\left(\frac{x - r(e_{ij})}{k^{j-1}\sigma}\right)\Phi\left(\frac{1}{\sigma\sqrt{1-\rho^2}}\left(\frac{\rho(x - r(e_{ij}))}{k^{j-1}} - \frac{x - r(e_{i,3-j})}{k^{2-j}}\right)\right), \tag{B.11}$$

$$F_{s_i}(x) = \Phi\left(\frac{x - r(e_{i1})}{\sigma}\right) + \int_{(x-r(e_{i1}))/\sigma}^{\infty} \Phi\left(\frac{x - r(e_{i2}) - \rho k\sigma u}{k\sigma\sqrt{1-\rho^2}}\right)\phi(u)du \tag{B.12}$$

by eq. (46.77)-(46.78) in Kotz et al. (2000). Using straightforward differentiation and the law of iterated expectations, we find that any symmetric pure-strategy PBE satisfies the following optimality condition for $j = 1, 2$:

$$\begin{aligned}\frac{c'(e_j^{sys})}{r'(e_j^{sys})} &= A(n-1) \\ &\cdot \int_{-\infty}^{\infty} f_s(r(e_j^s) + x)F_s(r(e_j^s) + x)^{n-2}(1 - F_{\zeta_{3-j}|\zeta_j}(r(e_j^s) - r(e_{3-j}^s) + x))f_{\zeta_j}(x)dx.\end{aligned} \tag{B.13}$$

Case A: $k = 1$. If $k = 1$, then there exists a solution to (B.13) with $e_1^{sys} = e_2^{sys}$, and this solution is given by

$$\frac{c'(e_j^{sys})}{r'(e_j^{sys})} = \frac{A(n-1)}{2\sigma} \cdot \mathcal{I}_{n-2}\left(-\sqrt{\frac{1-\rho}{1+\rho}}\right), \tag{B.14}$$

$j = 1, 2$.

Case B: $k \to \infty$. By (B.13), we have the following upper bound on e_2^{sys}:

$$\begin{aligned}\frac{c'(e_2^{sys})}{r'(e_2^{sys})} &\leq A(n-1)\int_{-\infty}^{\infty} f_s(r(e_2^s) + x)f_{\zeta_2}(x)dx && \text{(B.15)} \\ &\leq \frac{A(n-1)}{\sigma^2}\int_{-\infty}^{\infty}\left(\phi\left(\frac{r(e_2^s) - r(e_1^s) + x}{\sigma}\right) + \frac{1}{k}\phi\left(\frac{x}{k\sigma}\right)\right)\frac{1}{k}\phi\left(\frac{x}{k\sigma}\right)dx && \text{(B.16)} \\ &\leq \frac{\sqrt{2}A(n-1)}{k\sigma\sqrt{\pi}}. && \text{(B.17)}\end{aligned}$$

It follows that $\lim_{k\to\infty} c'(e_2^{sys})/r'(e_2^{sys}) = 0$, and thus, given the properties of c and r, $\lim_{k\to\infty} e_2^{sys} = 0$. We now turn to e_1^{sys}:

$$\lim_{k\to\infty} \frac{c'(e_1^{sys})}{r'(e_1^{sys})} = \frac{A(n-1)}{\sigma} \cdot \int_{-\infty}^{\infty} \left[\Phi\left(\frac{\rho y}{\sqrt{1-\rho^2}}\right)\phi(y)\right]^2 \left[\Phi(y) + \int_y^{\infty} \Phi\left(-\frac{\rho u}{\sqrt{1-\rho^2}}\right)\phi(u)du\right]^{n-2} dy \tag{B.18}$$

$$= \frac{A(n-1)}{2^n\sigma} \int_{-\infty}^{\infty} \left[1 + \Psi\left(y; \frac{\rho}{\sqrt{1-\rho^2}}\right)\right]^{n-2} \psi\left(y; \frac{\rho}{\sqrt{1-\rho^2}}\right)^2 dy \tag{B.19}$$

$$= \frac{A(n-1)}{2^n\sigma} \sum_{l=0}^{n-2} \binom{n-2}{l} \mathcal{I}_l\left(\frac{\rho}{\sqrt{1-\rho^2}}\right), \tag{B.20}$$

where the first equality is an application of Lebesgue's Dominated Convergence Theorem. We note that the solution (e_1^{sys}, e_2^{sys}) is unique as $k \to \infty$.

Component contest.

Equilibrium efforts in a component contest (e_1^{cpo}, e_2^{cpo}) are again given by (B.7) and (B.8) because each component contest is run separately, so the technological relationship between components does not affect the suppliers' equilibrium behavior. The buyer's optimal choice of p solves:

$$p^* \in \operatorname*{argmax}_p \Pi^{cpo}(p) = \mathbb{E}[\min\{r(e_1^{cpo}) + \max_i\{\zeta_{i1}\}, r(e_2^{cpo}) + \max_i\{\zeta_{i2}\}\}]. \tag{B.21}$$

By Assumption 3.2, $r(e_j^{cpo})$, $j = 1, 2$, is strictly concave in p, and since concavity is preserved under the pointwise minimization and expectation operators, it follows that $\Pi^{cpo}(p)$ is strictly concave in p. Moreover, the necessary and sufficient first-order condition $\partial\Pi^{cpo}(p)/\partial p = 0$ reveals that $p^*(k = 1) = 1/2$ and $\lim_{k\to\infty} p^*(k) = 1$, and we let $\Pi_{cml}^{cpo} \equiv \Pi^{cpo}(p^*)$.

B.3 Proofs

Proof of Lemma 3.1. (ia) The result follows directly from (B.9) and the discussion thereafter.

(ib) For $k = 1$ and $\rho = 0$, (B.3) implies that $c'(e_1^{sys})/r'(e_1^{sys}) = c'(e_2^{sys})/r'(e_2^{sys}) = A\mu^{(n)}/(\sqrt{2}n\sigma)$, and (B.7)-(B.8) imply that $c'(e_1^{cpo})/r'(e_1^{cpo}) = c'(e_2^{cpo})/r'(e_2^{cpo}) = A\mu^{(n)}/(2n\sigma)$. Since c'/r' is a strictly increasing function, it follows that $e_1^{sys} = e_2^{sys} > e_1^{cpo} = e_2^{cpo}$.

(iia) The result follows directly from (B.9) and the discussion thereafter.

(iib) For $\rho = 0$, as $k \to \infty$, (B.7)-(B.8) together with $p^* = 1$ imply that $c'(e_1^{cpo})/r'(e_1^{cpo}) = A\mu^{(n)}/(n\sigma)$ and $c'(e_2^{cpo})/r'(e_2^{cpo}) = 0$, yielding $e_1^{cpo} > 0$ and $e_2^{cpo} = 0$. Furthermore, (B.3) reveals that $\lim_{k\to\infty} c'(e_1^{sys})/r'(e_1^{sys}) = \lim_{k\to\infty} c'(e_2^{sys})/r'(e_2^{sys}) = 0$, and thus $\lim_{k\to\infty} e_1^{sys} = \lim_{k\to\infty} e_2^{sys} = 0$ by the properties of r and c. □

Proof of Proposition 3.1. For $\rho = 0$, the buyer's expected equilibrium profits in a system and component contest regime are $\Pi_{sub}^{sys} = 2\eta(A\mu^{(n)}/(\sqrt{1+k^2}n\sigma)) + \sqrt{1+k^2}\sigma\mu^{(n)}$ and $\Pi_{sub}^{cpo} = \eta(p^* A\mu^{(n)}/(n\sigma)) + \eta((1-p^*)A\mu^{(n)}/(nk\sigma)) + (1+k)\sigma\mu^{(n)}$, respectively (see (B.4) and (B.9)). Define $\Delta(\sigma, k) = \Pi_{sub}^{sys} - \Pi_{sub}^{cpo}$, which is a continuous function in σ and k for $\sigma > 0$ and $k \geq 1$.

(i) Given $k = 1$, $\rho = 0$ and the results of Lemma 3.1 the buyer's expected equilibrium profits in a system and component contest regime simplify to $\Pi_{sub}^{sys} = 2\eta(A\mu^{(n)}/(\sqrt{2}n\sigma)) + \sqrt{2}\sigma\mu^{(n)}$ and $\Pi_{sub}^{cpo} = 2\eta(A\mu^{(n)}/(2n\sigma)) + 2\sigma\mu^{(n)}$, respectively. Since η is strictly increasing and $\eta(0) = 0$, it follows that $\lim_{\sigma\to 0} \Delta(\sigma, k = 1) > 0$ and $\lim_{\sigma\to\infty} \Delta(\sigma, k = 1) < 0$. By the Intermediate Value Theorem, there exist thresholds $\underline{\sigma}_{sub}, \overline{\sigma}_{sub}$, with $0 < \underline{\sigma}_{sub} \leq \overline{\sigma}_{sub} < \infty$, such that $\Delta(\sigma, k = 1) > 0$ for all $\sigma \in (0, \underline{\sigma}_{sub})$, and $\Delta(\sigma, k = 1) < 0$ for all $\sigma > \overline{\sigma}_{sub}$. The result now follows from the continuity of $\Delta(\sigma, k)$, the fact that a continuous mapping from a connected subset of a metric space to another metric space yields a connected image set (Ok 2007, p. 220), and by integrating our assumption that $\sigma > \sigma_{sub}$ (note that σ_{sub} may be larger than $\underline{\sigma}_{sub}$, in which case the interval $(\sigma_{sub}, \underline{\sigma}_{sub})$ is empty). To see how $\underline{\sigma}_{sub}$ behaves for large n, note that $\lim_{n\to\infty} \Delta(\sigma, k = 1) = -(2 - \sqrt{2})\sigma \lim_{n\to\infty} \mu^{(n)}$, and therefore $\underline{\sigma}_{sub} \to 0$.

(ii) Fix any $\sigma > \sigma_{sub}$. Then, $\lim_{k\to\infty} \Delta(\sigma, k) < 0$, and thus, by the Intermediate Value Theorem, there exists a threshold $\overline{k}_{sub} < \infty$ such that $\Delta(\sigma, k) < 0$ for all $k > \overline{k}_{sub}$.

(iii) Fix any $k \geq 1$. Then, $\lim_{\sigma\to\infty} \Delta(\sigma, k) < 0$, and thus, by the Intermediate Value Theorem, there exists a threshold $\overline{\sigma}_{sub} < \infty$ such that $\Delta(\sigma, k) < 0$ for all $\sigma > \overline{\sigma}_{sub}$. $\square$

Proof of Lemma 3.2. (ia) The result follows directly from (B.21) and the discussion thereafter.

(ib) For $k = 1$ and $\rho = 0$, (B.7)-(B.8) imply that $c'(e_1^{cpo})/r'(e_1^{cpo}) = c'(e_2^{cpo})/r'(e_2^{cpo}) = A(n-1)\mathcal{I}_{n-2}(0)/(2\sigma)$, and (B.14) implies that $c'(e_1^{sys})/r'(e_1^{sys}) = c'(e_2^{sys})/r'(e_2^{sys}) = A(n-1)\mathcal{I}_{n-2}(-1)/(2\sigma)$. By Lemma AB.1(ii), $\mathcal{I}_{n-2}(-1) > \mathcal{I}_{n-2}(0)$, and therefore $e_1^{sys} = e_2^{sys} > e_1^{cpo} = e_2^{cpo} > 0$.

(iia) The result follows directly from (B.21) and the discussion thereafter.

(iib) For $\rho = 0$, as $k \to \infty$, (B.8) and (B.17) reveal that $e_2^{cpo} = 0$ and $e_2^{sys} = 0$, respectively. By (B.20) and (B.7), we have

$$\lim_{k\to\infty} \frac{c'(e_1^{sys})}{r'(e_1^{sys})} = \frac{A(n-1)}{2^n\sigma} \sum_{l=0}^{n-2} \binom{n-2}{l} \mathcal{I}_l(0) = \frac{A}{n2^n\sigma} \sum_{l=0}^{n-2} \binom{n}{l+2} \mu^{(l+2)} \tag{B.22}$$

$$< \frac{A\mu^{(n)}}{n2^n\sigma} \sum_{l=2}^{n} \binom{n}{l} = \frac{A\mu^{(n)}}{n\sigma}\left(1 - \frac{n+1}{2^n}\right) < \frac{A\mu^{(n)}}{n\sigma} = \lim_{k\to\infty} \frac{c'(e_1^{cpo})}{r'(e_1^{cpo})}. \tag{B.23}$$

Since c'/r' is strictly increasing it follows readily that $\lim_{k\to\infty} e_1^{sys} < \lim_{k\to\infty} e_1^{cpo}$. Last, we note that as $k \to \infty$, (B.22) reveals that the right-hand side of the optimality condition does not depend on e_1^{sys} and e_2^{sys}. Taken together with the properties of c'/r', this implies that (e_1^{sys}, e_2^{sys}) is unique as $k \to \infty$. $\square$

Proof of Proposition 3.2. For $\rho = 0$, the buyer's expected equilibrium profits in a system and component contest regime are $\Pi_{cml}^{sys} = \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\min_{j\in\{1,2\}}\{r(e_j^{sys}) + \zeta_{ij}\}\}]$ and $\Pi_{cml}^{cpo} = \mathbb{E}[\min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{r(e_j^{cpo}) + \zeta_{ij}\}\}]$, respectively. Define $\Delta(\sigma, k) = \Pi_{cml}^{sys} - \Pi_{cml}^{cpo}$, which is a continuous function in σ and k for $\sigma > 0$ and $k \geq 1$.

(i) Given $k = 1$, $\rho = 0$ and the results of Lemma 3.2, the buyer's expected equilibrium profits in a system and component contest regime are $\Pi_{cml}^{sys} = \eta(A(n-1)\mathcal{I}_{n-2}(-1)/(2\sigma)) + \sigma\mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\min_{j\in\{1,2\}}\{X_{ij}\}\}]$ and $\Pi_{cml}^{cpo} = \eta(A(n-1)\mathcal{I}_{n-2}(0)/(2\sigma)) + \sigma\mathbb{E}[\min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{X_{ij}\}\}]$, respectively, where X_{ij} are independent standard normal random variables. Since η is strictly increasing, $\eta(0) = 0$, $\mathcal{I}_{n-2}(-1) > \mathcal{I}_{n-2}(0)$, and $\mathbb{E}[\min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{X_{ij}\}\}] > \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\min_{j\in\{1,2\}}\{X_{ij}\}\}]$, it follows that $\lim_{\sigma\to 0} \Delta(\sigma, k = 1) > 0$ and $\lim_{\sigma\to\infty} \Delta(\sigma, k = 1) < 0$. By the Intermediate Value Theorem, there exist thresholds

$\underline{\sigma}_{cml}, \overline{\sigma}_{cml}$, with $0 < \underline{\sigma}_{cml} \leq \overline{\sigma}_{cml} < \infty$, such that $\Delta(\sigma, k = 1) > 0$ for all $\sigma \in (0, \underline{\sigma}_{cml})$, and $\Delta(\sigma, k = 1) < 0$ for all $\sigma > \overline{\sigma}_{cml}$. The result now follows from the continuity of $\Delta(\sigma, k)$, the fact that a continuous mapping from a connected subset of a metric space to another metric space yields a connected image set (Ok 2007, p. 220), and by integrating our assumption that $\sigma > \sigma_{cml}$ (note that σ_{cml} may be larger than $\underline{\sigma}_{cml}$, in which case the interval $(\sigma_{cml}, \underline{\sigma}_{cml})$ is empty). To see how $\underline{\sigma}_{cml}$ behaves for large n, note that $\lim_{n\to\infty} \Delta(\sigma, k = 1) = -\sigma \lim_{n\to\infty}[\mathbb{E}[\min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{X_{ij}\}\}] - \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\min_{j\in\{1,2\}}\{X_{ij}\}\}]]$, and therefore $\underline{\sigma}_{cml} \to 0$.

(ii) Fix any $\sigma > \sigma_{cml}$. For $k \to \infty$, Lemma 3.2 implies that $\Pi^{cpo}_{cml} = \mathbb{E}[\min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{r(e^{cpo}_j) + \zeta_{ij}\}\}] > \mathbb{E}[\min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{r(e^{sys}_j) + \zeta_{ij}\}\}] > \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\min_{j\in\{1,2\}}\{r(e^{sys}_j) + \zeta_{ij}\}\}] = \Pi^{sys}_{cml}$. Thus, $\lim_{k\to\infty} \Delta(\sigma, k) < 0$, and by the Intermediate Value Theorem, there exists a threshold $\overline{k}_{cml} < \infty$ such that $\Delta(\sigma, k) < 0$ for all $k > \overline{k}_{cml}$.

(iii) Fix any $k \geq 1$, and define $\underline{r} = \min_{j\in\{1,2\}}\{r(e^{cpo}_j)\}$ and $\overline{r} = \max_{j\in\{1,2\}}\{r(e^{sys}_j)\}$. It follows that $\Pi^{cpo}_{cml} \geq \underline{r} + \sigma\mathbb{E}[\min\{\max_{i\in\{1,\ldots,n\}}\{X_{i1}\}, \max_{i\in\{1,\ldots,n\}}\{kX_{i2}\}\}]$ and $\Pi^{sys}_{cml} \leq \overline{r} + \sigma\mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\min\{X_{i1}, kX_{i2}\}\}]$, where X_{ij} are independent standard normal random variables. Since $\mathbb{E}[\min\{\max_{i\in\{1,\ldots,n\}}\{X_{i1}\}, \max_{i\in\{1,\ldots,n\}}\{kX_{i2}\}\}] > \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\min\{X_{i1}, kX_{i2}\}\}]$, and $\underline{r}$ as well as $\overline{r}$ are bounded in value, it follows that $\lim_{\sigma\to\infty} \Delta(\sigma, k) < 0$, and thus, by the Intermediate Value Theorem, there exists a threshold $\overline{\sigma}_{cml} < \infty$ such that $\Delta(\sigma, k) < 0$ for all $\sigma > \overline{\sigma}_{cml}$. □

Proof of Proposition 3.3. (i) This result follows immediately from (B.9) and the fact that $\mu^{(n)}$ and thus also p^* do not depend on ρ.

(ii) To prove the result, we need to show that the buyer's equilibrium expected profit $\Pi^{cpo}_{cml}(p^*(\rho); \rho)$ increases in ρ, where $p^*(\rho)$ is the optimally chosen p for given ρ, and equilibrium efforts are given by (B.7)-(B.8). We do so by verifying that $\Pi^{cpo}_{cml}(p^*(\rho_1); \rho_1) \leq \Pi^{cpo}_{cml}(p^*(\rho_1); \rho_2) \leq \Pi^{cpo}_{cml}(p^*(\rho_2); \rho_2)$ for any fixed $0 \leq \rho_1 < \rho_2 < 1$. Clearly, the last inequality follows from the optimality of p^*, so it remains to prove the first inequality.

For any given p, $\Pi^{cpo}_{cml}(p; \rho) = \mathbb{E}[\min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{v^{cpo}_{ij}\}\}]$, where $v^{cpo}_i = (v^{cpo}_{i1}, v^{cpo}_{i2})$ follows a bivariate Normal distribution with correlation ρ and marginal distributions $v^{cpo}_{i1} \sim N(r(e^{cpo}_1), \sigma^2)$ and $v^{cpo}_{i2} \sim N(r(e^{cpo}_2), k^2\sigma^2)$ for all i. To conclude the proof, we need to show that $\Pi^{cpo}_{cml}(p; \rho)$ increases in ρ for any fixed p. This is true if the random variable $V(\rho) \equiv \min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{v^{cpo}_{ij}\}\}$ first-order stochastically increases in ρ; a property that we verify in the next step. Let $V_j \equiv \max_{i\in\{1,\ldots,n\}}\{v^{cpo}_{ij}\}$ for $j = 1, 2$.

For any $y \in \mathbb{R}$ and fixed p, $\mathbb{P}(V(\rho) > y) = \mathbb{P}(V_1 > y) - \mathbb{P}(V_1 > y|V_2 < y)\mathbb{P}(V_2 < y)$. Clearly, only the term $\mathbb{P}(V_1 > y|V_2 < y)$ depends on ρ, and thus $V(\rho)$ stochastically increases in ρ if and only if this term decreases in ρ. Note that $\mathbb{P}(V_1 > y|V_2 < y) = 1 - \mathbb{P}(v_{i1}^{cpo} < y, v_{i2}^{cpo} < y)^n / \mathbb{P}(v_{i2}^{cpo} < y)^n$, where the denominator is independent of ρ and $\mathbb{P}(v_{i1}^{cpo} < y, v_{i2}^{cpo} < y)$ increases in ρ by Lemma B.2(i). Hence $\mathbb{P}(V_1 > y|V_2 < y)$ decreases in ρ. □

Proof of Proposition 3.4. (i) By (B.4), the buyer's expected equilibrium profits are $\Pi_{sub}^{sys}(\rho) = 2\eta(A\mu^{(n)}/(n\sigma\sqrt{1+k^2+2k\rho})) + \sigma\mu^{(n)}\sqrt{1+k^2+2k\rho}$. Taking the first-order derivative of $\Pi_{sub}^{sys}(\rho)$ with respect to ρ leads to $d\Pi_{sub}^{sys}(\rho)/d\rho = k\mu^{(n)}(n\sigma^2(1+k^2+2k\rho) - 2A\eta'(A\mu^{(n)}/(n\sigma\sqrt{1+k^2+2k\rho})))/(n\sigma(1+k^2+2k\rho)^{3/2})$, and the result follows by comparing this expression to zero.

(ii) For given ρ, the buyer's expected equilibrium profits are $\Pi_{cml}^{sys}(\rho) = \mathbb{E}[\max_{i\in\{1,\dots,n\}}\{\min\{v_{i1}^{sys}, v_{i2}^{sys}\}\}]$, where $(v_{i1}^{sys}, v_{i2}^{sys})$ follows a bivariate Normal distribution with marginal distributions $v_{i1}^{sys} \sim N(r(e_1^{sys}), \sigma^2)$, $v_{i2}^{sys} \sim N(r(e_2^{sys}), k^2\sigma^2)$ and correlation ρ for all $i \in \{1,\dots,n\}$. To prove the claim, we will show that for $k = 1$ and σ small, $\Pi_{cml}^{sys}(\rho)$ decreases in ρ, and that $\Pi_{cml}^{sys}(\rho)$ increases in ρ for $k \to \infty$ or $\sigma \to \infty$. The result follows then directly from the continuity of $\Pi_{cml}^{sys}(\rho)$ and the Intermediate Value Theorem.

Case A: $k = 1$. By (B.14) we have $e_1^{sys} = e_2^{sys}$, and hence $\Pi_{cml}^{sys}(\rho) = r(e_1^{sys}) + \sigma\mathbb{E}[\max_{i\in\{1,\dots,n\}}\{\min\{X_{i1}, X_{i2}\}\}]$, where (X_{i1}, X_{i2}) follows a standard bivariate Normal distribution with correlation ρ for all $i \in \{1,\dots,n\}$. In addition, (B.14) together with Lemma B.1(ii) implies that equilibrium efforts decrease in ρ, and thus $dr(e_1^{sys})/d\rho < 0$. Therefore, $\lim_{\sigma\to 0} d\Pi_{cml}^{sys}(\rho)/d\rho < 0$.

Case B: $k \to \infty$ or $\sigma \to \infty$. $\Pi_{cml}^{sys}(\rho) = \mathbb{E}[\max_{i\in\{1,\dots,n\}}\{\min\{v_{i1}^{sys}, v_{i2}^{sys}\}\}]$ increases in ρ if the random variable $W(\rho) \equiv \max_{i\in\{1,\dots,n\}}\{\min\{v_{i1}^{sys}, v_{i2}^{sys}\}\}$ stochastically increases in ρ, or equivalently, $\mathbb{P}(W(\rho) < u) = (1 - \mathbb{P}(v_{i1}^{sys} \geq u, v_{i2}^{sys} \geq u))^n$ decreases in ρ for all $u \in \mathbb{R}$. This is true if $\mathbb{P}(v_{i1}^{sys} \geq u, v_{i2}^{sys} \geq u) = \mathbb{P}(\overline{v}_{i1} \geq u - r(e_1^{sys}), \overline{v}_{i2} \geq u - r(e_2^{sys}))$ increases in ρ for all $u \in \mathbb{R}$, where $\overline{v}_{i1}$ and $\overline{v}_{i2}$ are centered Normal random variables. For fixed e_1^{sys} and e_2^{sys}, this probability increases in ρ by Lemma B.2(ii); and for fixed ρ, it increases obviously in e_1^{sys} and e_2^{sys}. Since both sensitivities point in the same direction, it remains to verify that e_1^{sys} and e_2^{sys} weakly increase in ρ as $k \to \infty$ (resp. $\sigma \to \infty$). For $k \to \infty$, this is true because e_1^{sys} increases in ρ by (B.20) and Lemma B.1(ii), and

$e_2^{sys} = 0$ by (B.17). For $\sigma \to \infty$, $e_1^{sys} = e_2^{sys} = 0$ by (B.13), thereby concluding the proof. $\square$

Proof of Proposition 3.5. (i) The proof follows by considering Π_{sub}^{sys} and Π_{sub}^{cpo} as given in (B.4) and (B.9), respectively, and applying exactly the same argument as in the proof of Proposition 3.1.

(iia) For $k = 1$ and σ small, Propositions 3.3(ii) and 3.4(ii) imply that Π_{cml}^{cpo} increases in ρ whereas Π_{cml}^{sys} decreases in ρ. Hence to verify that $\Pi_{cml}^{sys}(\rho^{sys}) \geq \Pi_{cml}^{cpo}(\rho^{cpo})$ for any $\Delta\rho \geq 0$ it is sufficient to establish that $\Pi_{cml}^{sys}(\rho = 1) \geq \Pi_{cml}^{cpo}(\rho = 1)$. By (B.7), (B.8) and (B.14) it follows that for $k = 1$ and $\rho = 1$, $e_1^{sys} = e_2^{sys} = e_1^{cpo} = e_2^{cpo}$. As a result, $\Pi_{cml}^{sys}(\rho = 1) = r(e_1^{sys}) + \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\zeta_{i1}\}] = r(e_1^{cpo}) + \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\zeta_{i1}\}] = \Pi_{cml}^{cpo}(\rho = 1)$, which proves the claim.

(iib) For $\sigma \to \infty$, (B.7), (B.8) and (B.14) imply that $e_1^{sys} = e_2^{sys} = e_1^{cpo} = e_2^{cpo} = 0$. It follows that $\Pi_{cml}^{sys}(\rho^{sys}) = \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\min\{\zeta_{i1}, \zeta_{i2}\}\}]$ and $\Pi_{cml}^{cpo}(\rho^{cpo}) = \mathbb{E}[\min_{j\in\{1,2\}}\{\max_{i\in\{1,\ldots,n\}}\{\zeta_{ij}\}\}]$. Clearly, if $\rho^{sys} = \rho^{cpo}$ then $\Pi_{cml}^{cpo}(\rho^{cpo}) \geq \Pi_{cml}^{sys}(\rho^{sys})$; and if $\rho^{sys} = 1$ and $\rho^{cpo} = 0$ then $\Pi_{cml}^{sys}(\rho^{sys}) = \mathbb{E}[\max_{i\in\{1,\ldots,n\}}\{\zeta_{i1}\}] > \Pi_{cml}^{cpo}(\rho^{cpo})$. The result now follows immediately from the continuity of all involved functions in conjunction with the fact that a continuous mapping from a connected subset of a metric space to another metric space yields a connected image set (Ok 2007, p. 220). $\square$

[This page intentionally left blank.]

Appendix C

Proofs of Chapter IV

Proof of Proposition 4.1. To prove the claim, we first derive the firm's optimization problem P and then solve P to determine the optimal compensation scheme and the firm's expected profits.

The Optimization problem: By the revelation principle, we restrict attention to the optimal contract that induces high-effort testing and truth telling by all experts. This, however, requires several incentive constraints to be satisfied. To derive those, we need to ensure that high-effort testing and truth telling is indeed optimal for each expert $i \in \mathcal{I}$, given the assumptions that all other experts exert high effort and report truthfully, and that the firm chooses the ex post optimal design alternative.

After having received all recommendations, the firm chooses the design alternative for development that offers the highest ex post expected net contribution and for which there is a good recommendation (in case there is no good recommendation at all, the firm develops none of the designs). Given $r_i = g$, design i's ex post expected net contribution is $q_i(v_i - u_{ig}) - (1 - q_i)u_{ib} - \sum_{k \neq i} u_{ka}$. Constraint (4.5) orders the designs according to their maximum ex post expected net contribution and thus ranks them according to their relative attractiveness to the firm; represented by the index j in $y_i^{(j)}$.

We now derive the incentive compatibility constraints for design i that is the jth most attractive alternative. Given $e_i = h$ and upon receiving a good signal ($s_i = g$), expert i receives an expected utility of $\pi_{ij}^{gg} = (q_i u_{ig} + (1 - q_i)u_{ib})/2^{j-1} + (1 - 1/2^{j-1})u_{ia}$ when making a good recommendation ($r_i = g$), and $\pi_{ij}^{bg} = u_{ia} + u_{it}/2^{n-1}$ when making a bad recommendation ($r_i = b$). Similarly, given $e_i = h$ and upon receiving a bad signal ($s_i = b$), expert i receives an expected utility of $\pi_{ij}^{gb} = ((1 - q_i)u_{ig} + q_i u_{ib})/2^{j-1} + (1 - 1/2^{j-1})u_{ia}$ when making a good recommendation ($r_i = g$), and $\pi_{ij}^{bb} = u_{ia} + u_{it}/2^{n-1}$ when making a bad recommendation ($r_i = b$). Also, given truth telling, expert i's expected utility from exerting high effort is $\pi_{ij}(h) = (q_i u_{ig} + (1 - q_i)u_{ib})/2^j + (u_{ia} + (1/2)^{n-j}u_{it})/2^j + (1 - 1/2^{j-1})u_{ia} - c$, and $\pi_{ij}(l) = (u_{ig}/2 + u_{ib}/2)/2^j + (u_{ia} + (1/2)^{n-j}u_{it})/2^j + (1 - 1/2^{j-1})u_{ia}$ from exerting low effort. The incentive compatibility constraints follow from setting

$\pi_{ij}^{gg} \geq \pi_{ij}^{bg}$, $\pi_{ij}^{bb} \geq \pi_{ij}^{gb}$, and $\pi_{ij}(h) \geq \pi_{ij}(l)$ for all $i, j \in \mathcal{I}$, and multiplying these inequalities with $y_i^{(j)}$. Finally, because wages must be non-negative, we require $u_{ig}, u_{ib}, u_{ia}, u_{ia}+u_{it} \geq 0$, and (4.6) follows from noting that each design i must be assigned to exactly one attractiveness rank, and each rank j is hold by exactly one design.

The firm's expected profit consists of the expected market value of the chosen design net of development costs and the experts' expected wages. Given $y_i^{(j)} = 1$, the firm develops design i with probability $1/2^j$ and receives an expected value of $q_i v_i - K$. The expected wage payments to expert i are u_{ig} with probability $q_i/2^j$, u_{ib} with probability $(1-q_i)/2^j$, u_{ia} with probability $1 - 1/2^j$, and u_{it} with probability $1/2^n$. Summing over $i, j \in \mathcal{I}$ gives the firm's expected profit Π.

(i) Suppose the designs in $\mathcal{I}$ can be ordered such that $q_i v_i \geq q_{i+1}v_{i+1} + 2^{i+1}c[q_i/(2q_i - 1) - 2q_{i+1}/(2q_{i+1}-1)]^+$ for all $i \in \mathcal{I}\backslash\{n\}$. To solve the optimization problem P, we first derive the solution of a relaxed variant of P by dropping constraints (4.5), and then show that this solution is also feasible—and thus optimal—in P.

Given the structure of P without (4.5), maximizing the firm's expected profit is equivalent to separately minimizing the wage payments associated with each design $i \in \mathcal{I}$ whenever $y_i^{(j)} = 1$. Obviously, (4.4) implies that $u_{ig} > u_{ib}$, which allows us to rewrite (4.2) and (4.3) as $q_i u_{ig} + (1-q_i)u_{ib} \geq u_{ia} + 2^{j-n}u_{it} \geq (1-q_i)u_{ig} + q_i u_{ib}$. It follows that wage payments for design i with relative attractiveness j are minimized when $u_{ia} + 2^{j-n}u_{it} = (1-q_i)u_{ig} + q_i u_{ib}$, and u_{ig} and u_{ib} are chosen as low as possible. By (4.4) and (4.7), these minimal payments are $u_{ig} = 2^{j+1}c/(2q_i - 1)$ and $u_{ib} = 0$. Moreover, (4.1) reveals that the firm prefers paying u_{it} over u_{ia}; therefore $u_{ia} = 0$ and $u_{it} = 2^{n+1}(1-q_i)c/(2q_i - 1)$. Inserting these payments into (4.1) and using (4.6) gives $\Pi_P = \sum_{j=1}^n \sum_{i=1}^n y_i^{(j)}(q_i v_i/2^j) - \sum_{i=1}^n 2c/(2q_i - 1) - \sum_{j=1}^n K/2^j$. By the assumed ordering, we have $q_i v_i \geq q_{i+1}v_{i+1}$, and it follows that in optimum $y_i^{(i)} = 1$ for all $i \in \mathcal{I}$, and $y_i^{(j)} = 0$ for all $i \neq j$. Moreover, this candidate optimal solution satisfies (4.6) and is thus feasible.

It remains to show that the solution also satisfies (4.5). However, this is obvious because we can rewrite this condition by $q_i v_i - (2^{i+1}q_i c/(2q_i - 1)) \geq q_{i+1}v_{i+1} - (2^{i+2}q_{i+1}c/(2q_{i+1}-1))$, which is true by assumption.

(ii) This result follows directly from (i). □

Proof of Proposition 4.2. The optimization problem M can be derived in a similar way to the proof of Proposition 4.1. In particular, for each $i \in \mathcal{I}$, $\pi_i^{gg} = q_i u_{ig} + (1-q_i)u_{ib}$, $\pi_i^{bg} = \pi_i^{bb} = u_{ia} + P(s_j = b\ \forall j > i)\delta^{n-i}u_{it}$, $\pi_i^{gb} = (1-q_i)u_{ig} + q_i u_{ib}$, $\pi_i(h) = (q_i u_{ig} + (1 -$

$q_i)u_{ib})/2 + (u_{ia} + P(s_j = b\ \forall j > i)\delta^{n-i}u_{it})/2 - c$, and $\pi_i(l) = (u_{ig}/2 + u_{ib}/2)/2 + (u_{ia} + P(s_j = b\ \forall j > i)\delta^{n-i}u_{it})/2$. Since for design $i = n$, u_{na} and u_{nt} are paid simultaneously, we only require $u_{na} + u_{nt} \geq 0$ to ensure non-negative wages.

As for the firm's profits, the firm develops design i with probability $P(r_i = g,\ r_j = b\ \forall j < i)$ and receives a discounted expected value of $\delta^{i-1}(q_i v_i - K)$. The expected wage payments to expert i are $\delta^{i-1}u_{ig}$ with probability $P(\Theta_i = G | s_i = g)P(s_i = g,\ s_j = b\ \forall j < i)$, $\delta^{i-1}u_{ib}$ with probability $P(\Theta_i = B | s_i = g)P(s_i = g,\ s_j = b\ \forall j < i)$, $\delta^{i-1}u_{ia}$ with probability $P(s_i = b,\ s_j = b\ \forall j < i)$, and $\delta^{n-1}u_{it}$ with probability $P(s_i = b,\ s_j = b\ \forall j \neq i)$. Summing over $i \in \mathcal{I}$ gives the firm's expected profit Π.

(i) Given the structure of M, maximizing the firm's expected profit is equivalent to separately minimizing the wage payments associated with each design i. Note that (4.11) implies that $u_{ig} > u_{ib}$, which allows us to rewrite (4.9) and (4.10) as $q_i u_{ig} + (1 - q_i)u_{ib} \geq u_{ia} + 2^{i-n}\delta^{n-i}u_{it} \geq (1 - q_i)u_{ig} + q_i u_{ib}$. It follows readily that $u_{ia} + 2^{i-n}\delta^{n-i}u_{it} = (1 - q_i)u_{ig} + q_i u_{ib}$, and u_{ig} and u_{ib} should be chosen as low as possible. By (4.11) and (4.12), these minimal payments are $u_{ig} = 4c/(2q_i - 1)$ and $u_{ib} = 0$. Moreover, the firm is indifferent between paying u_{ia} or u_{it}, so without loss of optimality we can choose $u_{ia} = 4(1 - q_i)c/(2q_i - 1)$ and $u_{it} = 0$. Finally, $u_{ia}/u_{ig} = 1 - q_i < 1/2$ because $q_i > 1/2$.

(ii)-(iii) Given the optimal contract, we can rewrite the firm's expected profit as $\Pi_M = \sum_{i=1}^{n}(\delta^{i-1}/2^i)(q_i v_i - K - 4c/(2q_i - 1))$. Since $(\delta^{i-1}/2^i)$ is decreasing in i, the firm maximizes Π_M by testing the designs in decreasing order of $q_i v_i - 4c/(2q_i - 1)$. $\square$

Proof of Proposition 4.3. Define the expert's expected continuation utility before testing design $i \in \mathcal{I}$ by $\hat{\pi}_{i-1} = (q_i u_{ig} + (1 - q_i)u_{ib} + u_{ia} - 2c + \delta\hat{\pi}_i)/2$, with $\hat{\pi}_n = 0$. With this definition, the derivation of S is identical to that of M as given in the proof of Proposition 4.2. In particular, for each $i \in \mathcal{I}$, $\pi_i^{gg} = q_i u_{ig} + (1 - q_i)u_{ib}$, $\pi_i^{bg} = \pi_i^{bb} = u_{ia} + \delta\hat{\pi}_i$, $\pi_i^{gb} = (1 - q_i)u_{ig} + q_i u_{ib}$, $\pi_i(h) = (q_i u_{ig} + (1 - q_i)u_{ib})/2 + (u_{ia} + \delta\hat{\pi}_i)/2 - c$, and $\pi_i(l) = (u_{ig}/2 + u_{ib}/2)/2 + (u_{ia} + \delta\hat{\pi}_i)/2$, and limited liability enforces non-negative wage payments.

As for the firm's profits, the firm develops design i with probability $P(r_i = g,\ r_j = b\ \forall j < i)$ and receives a discounted expected value of $\delta^{i-1}(q_i v_i - K)$. The expected wage payments to expert i are $\delta^{i-1}u_{ig}$ with probability $P(\Theta_i = G | s_i = g)P(s_i = g,\ s_j = b\ \forall j < i)$, $\delta^{i-1}u_{ib}$ with probability $P(\Theta_i = B | s_i = g)P(s_i = g,\ s_j = b\ \forall j < i)$, and $\delta^{i-1}u_{ia}$ with probability $P(s_i = b,\ s_j = b\ \forall j < i)$. Summing over $i \in \mathcal{I}$ gives the firm's expected profit Π.

(i) Using the definition of $\hat{\pi}_0$, we can rewrite the firm's expected profit as $\Pi_S = \sum_{i=1}^{n}(\delta^{i-1}/2^i)(q_i v_i - K - 2c) - \hat{\pi}_0$. Thus, maximizing Π_S is equivalent to minimizing $\hat{\pi}_0$, which we do in the following. As a first step, we derive the minimum feasible $\hat{\pi}_{i-1}$ for given fixed $\hat{\pi}_i$.

Case (a): $\delta\hat{\pi}_i < 4(1-q_i)c/(2q_i-1)$. The optimal payments are $u_{ig} = 4c/(2q_i-1)$, $u_{ib} = 0$, and $u_{ia} = 4(1-q_i)c/(2q_i-1) - \delta\hat{\pi}_i$, and it follows that $\hat{\pi}_{i-1} = (3-2q_i)c/(2q_i-1)$.

Case (b): $4(1-q_i)c/(2q_i-1) \leq \delta\hat{\pi}_i \leq 4q_i c/(2q_i-1)$. The optimal payments are $u_{ig} = 4c/(2q_i-1)$, $u_{ib} = 0$, and $u_{ia} = 0$, and it follows that $\hat{\pi}_{i-1} = c/(2q_i-1) + \delta\hat{\pi}_i/2$.

Case (c): $\delta\hat{\pi}_i > 4q_i c/(2q_i-1)$. The optimal payments are $u_{ig} = \delta\hat{\pi}_i/q_i$, $u_{ib} = 0$, and $u_{ia} = 0$, and it follows that $\hat{\pi}_{i-1} = \delta\hat{\pi}_i - c$

Taken together, Cases (a)-(c) imply that $\hat{\pi}_{i-1}$ is non-decreasing in $\hat{\pi}_i$ for all $i \in \mathcal{I}$. As such, minimizing $\hat{\pi}_0$ is equivalent to separately minimizing $\hat{\pi}_i$ for each $i \in \mathcal{I}$, starting with $\hat{\pi}_n = 0$ and using Cases (a)-(c) for backwards induction. Thus, the optimal contract satisfies $u_{ig} = 4c/(2q_i-1) + [\delta\hat{\pi}_i/q_i - 4c/(2q_i-1)]^+$, $u_{ib} = 0$, and $u_{ia} = [4(1-q_i)c/(2q_i-1) - \delta\hat{\pi}_i]^+$ for all $i \in \mathcal{I}$.

(ii) If the designs in $\mathcal{I}$ can be ordered such that $q_i v_i \geq q_{i+1}v_{i+1}$, $q_i \geq q_{i+1}$ and $(1-q_i)4c/(2q_i-1) \leq \delta\hat{\pi}_i \leq 4q_i c/(2q_i-1)$ for all $i \in \mathcal{I}\backslash\{n\}$, then Cases (a) and (b) imply that $q_i u_{ig} + (1-q_i)u_{ib} + u_{ia} < q_{i+1}u_{i+1g} + (1-q_{i+1})u_{i+1b} + u_{i+1a}$ for all $i \in \mathcal{I}\backslash\{n\}$. By (4.13) and the assumption that $q_i v_i \geq q_{i+1}v_{i+1}$ it follows readily that it is optimal to test the designs in increasing order of i.

(iii) This result follows directly from inserting Proposition 4.3(i) in (4.13) and rearranging terms. □

Proof of Proposition 4.4. (i) By Proposition 4.2(iii), $\Pi_M = \sum_{i=1}^{n}(\delta^{i-1}/2^i)(q_i v_i - K - 4c/(2q_i-1))$, which reveals that the sign of the net profit contribution of each design $i \in \mathcal{N}$ is independent of the number and identity of the other designs to be tested. As a result, the firm finds it optimal to include all designs $i \in \mathcal{N}$ into the testing set $\mathcal{I}_M$ for which $q_i v_i - K - 4c/(2q_i-1) \geq 0$.

(ii) Consider the optimization problem P. By (4.2)-(4.4), we have $u_{ig} \geq 2^{j+1}c/(2q_i-1)$ and $u_{ia} + 2^{j-n}u_{it} \geq (1-q_i)u_{ig}$ for all $i, j \in \mathcal{I}$ such that $y_i^{(j)} = 1$. Hence, the profit contribution of design i with relative attractiveness j is $\Pi_i^{(j)} \leq (q_i v_i - K - q_i u_{ig} - u_{ia} - 2^{j-n}u_{it})/2^j \leq (q_i v_i - K - u_{ig})/2^j \leq (q_i v_i - K)/2^j - 2c/(2q_i-1)$. A necessary condition for $i \in \mathcal{I}_P$ is that $(q_i v_i - K)/2^j - 2c/(2q_i-1) \geq 0$ for some $j \in \mathcal{I}_P$. However, this can only be true if $q_i v_i - K - 4c/(2q_i - 1) \geq 0$. Comparing this condition with $\mathcal{I}_M$ completes the proof.

(iii) By Proposition 4.2 (ii) and (iii), it is optimal to have design n in the optimal set of designs for testing since it is profitable to test and eventually develop this design due to the stated condition $q_n v_n - K - 4c/(2q_n - 1) > 0$. Furthermore, all designs j with $q_j \geq q_n$ and $v_j \geq v_n$ belong to this optimal set because then, we also have $q_j v_j - K - 4c/(2q_j - 1) > 0$. $\square$

Proof of Proposition 4.5. (i) Suppose $\delta < 1$. For $c = 0$, we have $\mathcal{I}_P = \mathcal{I}_M = \mathcal{I}_S = \mathcal{N}$ and designs are tested in decreasing order of $q_i v_i$. By Propositions 4.1-4.3, it follows readily that $\Pi_P^* > \Pi_M^* = \Pi_S^*$. Thus, by continuity of the expected profits in c, there exists $\underline{c} > 0$ such that $\Pi_P^* > \max\{\Pi_M^*, \Pi_S^*\}$ for all $c < \underline{c}$.

(ii) Let $\mathcal{I}_P$ be the optimal set of designs to be tested under a parallel testing strategy, and assume that the designs in $\mathcal{I}_P$ can be ordered such that $q_i v_i \geq q_{i+1} v_{i+1} + 2^{i+1} c[q_i/(2q_i - 1) - 2q_{i+1}/(2q_{i+1} - 1)]^+$ for all $i \in \mathcal{I}_P \backslash \{n\}$. Then, by Proposition 4.1(iii), $\Pi_P^* = \sum_{i \in \mathcal{I}_P}((q_i v_i - K)/2^i - 2c/(2q_i - 1))$. Now, if the firm fixes the identity and ordering of designs, but instead uses a multi-expert sequential testing strategy, then $\Pi_M(\mathcal{I}_P) = \sum_{i \in \mathcal{I}_P} \delta^{i-1}(q_i v_i - K - 4c/(2q_i - 1))/2^i$. By comparing the different profits, we have $\Pi_M(\mathcal{I}_P) > \Pi_P^*$ if $c > \bar{c} \equiv \sum_{i \in \mathcal{I}_P}((1 - \delta^{i-1})(q_i v_i - K)/2^i) / \sum_{i \in \mathcal{I}_P}((2(1 - (\delta/2)^{i-1})/(2q_i - 1))$. Moreover, since $\mathcal{I}_P$ need not be optimal under a multi-expert sequential testing strategy, it follows that if $c > \bar{c}$, then $\Pi_P^* < \Pi_M(\mathcal{I}_P) \leq \Pi_M^* \leq \max\{\Pi_M^*, \Pi_S^*\}$.

(iii) Let $\mathcal{I}_M$ be the optimal set of designs to be tested under a multi-expert sequential testing strategy, with $\mathcal{I}_M$ optimally ordered according to Proposition 4.2(ii). By Proposition 4.2(iii) and 4.3(iii), we have $\Pi_M^* = \sum_{i \in \mathcal{I}_M}(\delta^{i-1}/2^i)(q_i v_i - K - 4c/(2q_i - 1))$ and $\Pi_S(\mathcal{I}_M) = \sum_{i \in \mathcal{I}_M} \delta^{i-1}(q_i v_i - K - \max\{4q_i c/(2q_i - 1), \delta\hat{\pi}_i, 4c/(2q_i - 1) - \delta\hat{\pi}_i\})/2^i \leq \Pi_S^*$. Clearly, a sufficient condition for $\Pi_S(\mathcal{I}_M) \geq \Pi_M^*$ is that $\delta\hat{\pi}_i \leq 4q_i c/(2q_i - 1)$ for all $i \in \mathcal{I}_M$.

By (4.17), we have $\delta\hat{\pi}_n = 0$ and $\delta\hat{\pi}_{i-1}$ increases in $\delta\hat{\pi}_i$ for all $i \in \mathcal{I}_M$. Moreover, if $\delta\hat{\pi}_i = 4q_i c/(2q_i - 1)$, then $\delta\hat{\pi}_{i-1} = \delta(4q_i c/(2q_i - 1) - c)$. Thus, by induction, if $\delta\hat{\pi}_i \leq 4q_i c/(2q_i - 1)$, then $\delta\hat{\pi}_{i-1} \leq \delta(4q_i c/(2q_i - 1) - c)$, and $\delta(4q_i c/(2q_i - 1) - c) \leq 4q_{i-1} c/(2q_{i-1} - 1)$ if $q_i \geq \underline{q}_{i-1}$. Finally, it is easy to show that $\underline{q}_i \leq q_i$ and $\underline{q}_i \leq 5/6$. $\square$

Proof of Proposition 4.6. (i) Note that $\mathcal{I}_{\text{seq}}^{\text{fb}} = \{i \in \mathcal{N} \mid q_i v_i - K - 2c \geq 0\}$. Comparing this with $\mathcal{I}_M$ as given in Proposition 4.4(i) immediately yields $\mathcal{I}_M \subseteq \mathcal{I}_{\text{seq}}^{\text{fb}}$. We next show that for any $i \in \mathcal{I}_S$, $q_i v_i - K - 4q_i c/(2q_i - 1) \geq 0$; implying that $\mathcal{I}_S \subseteq \mathcal{I}_{\text{seq}}^{\text{fb}}$.

Consider an arbitrary design $i \in \mathcal{I}_S$. By Proposition 4.3(i), if the firm receives a good recommendation for this design, the expected value of developing it is given by $q_i v_i - K - q_i u_{ig} - (1-q_i)u_{ib} \leq q_i v_i - K - 4q_i c/(2q_i - 1)$. Obviously, the firm only develops design i if the development generates a nonnegative expected value; i.e., $q_i v_i - K - 4q_i c/(2q_i - 1) \geq q_i v_i - K - q_i u_{ig} - (1 - q_i)u_{ib} \geq 0$. Suppose to the contrary that there exists a design $i \in \mathcal{I}_S$ such that $q_i v_i - K - q_i u_{ig} - (1 - q_i)u_{ib} < 0$. Obviously, the firm would never develop this design as the firm's outside option has zero, and thus greater value. In equilibrium, the expert anticipates the firm's development decision, and as a result, it is impossible for the firm to motivate the expert to exert high testing efforts. Hence, since design i will never be tested anyways, it is optimal for the firm to erase it from the set of designs to be tested.

(ii) We prove the claim by example. Consider a setting with three design alternatives and the following parameters: $q_1 = 0.55$, $q_2 = 0.64$, $q_3 = 1$, $v_1 = 100$, $v_2 = 85$, $v_3 = 53$, $K = 50$, and $c = 0.6$. In this case, the optimal set of designs to be tested under first-best conditions is $\mathcal{I}_{\text{par}}^{\text{fb}} = \{1, 2\}$, leading to an expected profit of $\Pi_{\text{par}}^{\text{fb}} = 2.4$. In contrast, under delegation, we have $\mathcal{I}_P = \{3\}$ with an expected profit of $\Pi_P^* = 0.3$. It follows that $\mathcal{I}_P \cap \mathcal{I}_{\text{par}}^{\text{fb}} = \emptyset$.

(iii) For brevity, let $|\mathcal{I}_M| = n_M$, $|\mathcal{I}_P| = n_P$, $|\mathcal{I}_{\text{seq}}^{\text{fb}}| = n_{\text{seq}}$, and $|\mathcal{I}_{\text{par}}^{\text{fb}}| = n_{\text{par}}$. With symmetric test efficiencies (i.e., $q_i = q$ for all $i \in \mathcal{N}$), under any sequential testing strategy it is always optimal to test designs in decreasing order of v_i, and under any parallel testing strategy the designs attractiveness decreases in v_i. It follows from Propositions 4.1-4.4 that $\mathcal{I}_P \subseteq \mathcal{I}_{\text{par}}^{\text{fb}}, \mathcal{I}_M \subseteq \mathcal{I}_{\text{seq}}^{\text{fb}}$, and consequently, $n_P \leq n_{\text{par}}, n_M \leq n_{\text{seq}}$. Without loss of generality, we relabel the designs such that $v_i \geq v_{i+1}$ for all $i \in \mathcal{N}$. Given these preliminaries, we prove the result by showing that for any $n_P \geq 0$, $\Pi_{\text{seq}}^{\text{fb}} \geq \Pi_{\text{par}}^{\text{fb}}$ implies $\Pi_M^* \geq \Pi_P^*$.

Case (a): $n_P = 0$. Since it always holds that $\Pi_P^* = 0 \leq \Pi_M^*$, the claim is trivially satisfied.

Case (b): $n_P = 1$. By Proposition 4.1(iii) and 4.2(iii), we have $\Pi_P^* = (qv_1 - K)/2 - 2c/(2q - 1) = \Pi_M(n = 1) \leq \Pi_M(n_M) = \Pi_M^*$, where the inequality follows from the optimality of n_M.

Case (c): $n_P \geq 2$. Define $\Delta\Pi(x, y) = \Pi(x) - \Pi(y)$. With this notation, we can rewrite the firm's first-best expected profits as $\Pi_{\text{seq}}^{\text{fb}}(n_{\text{seq}}) = \Pi_M(n_M) + \Delta\Pi_M(n_{\text{seq}}, n_M) + \sum_{i=1}^{n_{\text{seq}}}(c(\delta/2)^{i-1}(3-2q)/(2q-1))$, and $\Pi_{\text{par}}^{\text{fb}}(n_{\text{par}}) = \Pi_P(n_P) + \Delta\Pi_{\text{par}}^{\text{fb}}(n_{\text{par}}, n_P) + \sum_{i=1}^{n_p}(c(3-2q)/(2q-1))$. Hence $\Pi_{\text{seq}}^{\text{fb}} \geq \Pi_{\text{par}}^{\text{fb}}$ is equivalent to $\Pi_M(n_M) \geq \Pi_P(n_P) + \Delta\Pi_{\text{par}}^{\text{fb}}(n_{\text{par}}, n_P) -$

$\Delta\Pi_M(n_{\text{seq}}, n_M) + c(3-2q)/(2q-1)(n_P - (1-(\delta/2)^{n_{\text{seq}}})/(1-(\delta/2)))$. The right-hand side of this inequality is larger than $\Pi_P(n_P)$, which proves the claim. To see this, note that by optimality of n_{par} and n_M, we have $\Delta\Pi^{\text{fb}}_{\text{par}}(n_{\text{par}}, n_P) \geq 0$ and $\Delta\Pi_M(n_{\text{seq}}, n_M) \leq 0$, and finally, $n_P - (1 - (\delta/2)^{n_{\text{seq}}})/(1 - (\delta/2)) \geq n_P - 2 \geq 0$ because $n_P \geq 2$ by assumption.

Last, we prove that the converse statement is not always true. We do this by example. Consider a scenario with the following parameters: $N = 4$, $v_1 = 10$, $v_2 = 8$, $v_3 = 6$, $v_4 = 4$, $q = 1$, $\delta = 0.8$, $K = 2$, $c = 0.2$. Then $\Pi^*_P = 3.6 \leq \Pi^*_M = 3.82$, but $\Pi^{\text{fb}}_{\text{par}} = 4.2 \geq \Pi^{\text{fb}}_{\text{seq}} = 4.14$. □

[This page intentionally left blank.]

Bibliography

Adam K (2001) Learning while searching for the best alternative. *J Econ Theory* 101(1):252–280.

Ales L, Cho SH, Körpeoğlu E (2016) Innovation tournaments with multiple contributors, Working Paper.

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

Allen T (1977) *Managing the Flow of Technology* (MIT Press, Cambridge, MA).

Aoyagi M (2010) Information feedback in a dynamic tournament. *Game Econ Behav* 70(2):242–260.

Armstrong M (2000) Optimal multi-object auctions. *Rev Econ Stud* 67(3):455–481.

Arozamena L, Cantillon E (2004) Investment incentives in procurement auctions. *Rev Econ Stud* 71(1):1–18.

Avery C, Hendershott T (2000) Bundling and optimal auctions of multiple products. *Rev Econ Stud* 67(3):483–497.

Azzalini A (1985) A class of distributions which includes the normal ones. *Scand J Stat* 12(2):171–178.

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

Beil D, Wein L (2003) An inverse-optimization-based auction mechanism to support a multiattribute RFQ process. *Manage Sci* 49(11):1529–1545.

Bhaskaran S, Krishnan V (2009) Effort, revenue, and cost sharing mechanisms for collaborative new product development. *Manage Sci* 55(7):1152–1169.

Billinger S, Stieglitz N, Schumacher T (2014) Search on rugged landscapes: An experimental study. *Organ Sci* 25(1):93–108.

Bimpikis K, Ehsani S, Mostagir M (2016) Designing dynamic contests, Working Paper.

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

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

Buenen M, Muthukrishnan G (2016) World quality report 2016-17. Technical report, Capgemini, Sogeti and HPE.

Burton T, Jenkins N, Sharpe D, Bossanyi E (2011) *Wind Energy Handbook* (John Wiley & Sons, Chicester, United Kingdom), 2nd edition.

Cabral L, Cozzi G, Denicolò V, Spagnolo G, Zanza M (2006) Procuring innovations. Dimitri N, Piga G, Spagnolo G, eds., *Handbook of Procurement*, chapter 19 (Cambridge University Press, Cambridge, UK).

Cachon G, Zhang F (2006) Procuring fast delivery: Sole sourcing with information asymmetry. *Manage Sci* 52(6):881–896.

Chaturvedi A, Beil D, Martínez-de-Albéniz V (2014) Split-award auctions for supplier retention. *Manage Sci* 60(7):1719–1737.

Chaturvedi A, Martínez-de-Albéniz V (2011) Optimal procurement design in the presence of supply risk. *Manuf Serv Oper Manag* 13(2):227–243.

Che YK (1993) Design competition through multidimensional auctions. *RAND J Econ* 24(4):668–680.

Che YK, Gale I (2003) Optimal design of research contests. *Am Econ Rev* 93(3):646–671.

Chen F (2007) Auctioning supply contracts. *Manage Sci* 53(10):1562–1576.

Chen R, Roundy R, Zhang R, Janakiraman G (2005) Efficient auction mechanisms for supply chain procurement. *Manage Sci* 51(3):467–482.

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

Clark K, Fujimoto T (1989) Lead time in automobile development: Explaining the japanese advantage. *J Eng Technol Manage* 6(1):25–58.

Clark K, Fujimoto T (1991) *Product Development Performance: Strategy, Organization and Management in the World Auto Industry* (Harvard Business School Press, Cambridge, MA).

Dahan E, Mendelson H (2001) An extreme-value model of concept testing. *Manage Sci* 47(1):102–116.

Dasgupta S, Spulber D (1990) Managing procurement auctions. *Information Economics and Policy* 4(1):5–29.

Ederer F (2010) Feedback and motivation in dynamic tournaments. *J Econ Manage Strat* 19(3):733–769.

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

Elmaghraby W (2000) Supply contract competition and sourcing policies. *Manuf Serv Oper Manag* 2(4):350–371.

Elmaghraby W, Keskinocak P (2004) Combinatorial auctions in procurement. *The Practice of Supply Chain Management: Where Theory and Application Converge*, chapter 15 (Springer, Boston, MA).

Erat S, Gneezy U (2016) Incentives for creativity. *Exp Econ* 19(2):269–280.

Erat S, Kavadias S (2008) Sequential testing of product designs: Implications for learning. *Manage Sci* 54(5):956–968.

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

Ethiraj S, Levinthal D (2004) Modularity and innovation in complex systems. *Manage Sci* 50(2):159–173.

EU (2004) Directive 2004/18/EC of the european parliament and of the council. *Official J EU* 134:114–240.

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

Fullerton RL, McAfee RP (1999) Auctioning entry into tournaments. *J Polit Econ* 107(3):573–605.

Gavious A, Moldovanu B, Sela A (2002) Bid costs and endogenous bid caps. *RAND J Econ* 33(4):709–722.

Gerardi D, Maestri L (2012) A principal–agent model of sequential testing. *Theor Econ* 7(3):425–463.

Gershkov A, Perry M (2009) Tournaments with midterm reviews. *Game Econ Behav* 66(1):162–190.

Girotra K, Terwiesch C, Ulrich K (2010) Idea generation and the quality of the best idea. *Manage Sci* 56(4):591–605.

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

Green JR, Stokey NL (1983) A comparison of tournaments and contracts. *J Polit Econ* 91(3):349–364.

Greene B (2006) The universe on a string. *The New York Times* 20 October 2006.

Gromb D, Martimort D (2007) Collusion and the organization of delegated expertise. *J Econ Theory* 137(1):271–299.

Gross D (2017) Performance feedback in competitive product development. *RAND J Econ* 48(2):438–466.

Gürtler O, Münster J, Nieken P (2013) Information policy in tournaments with sabotage. *Scand J Econ* 115(3):932–966.

Gurvich I, Van Mieghem J (2015) Collaboration and multitasking in networks: Architectures, bottlenecks, and capacity. *Manuf Serv Oper Manag* 17(1):16–33.

Halac M, Kartik N, Liu Q (2016) Contests for experimentation. *J Polit Econ* forthcoming.

Hausch D (1986) Multi-object auctions: Sequential vs. simultaneous sales. *Manage Sci* 32(10):1599–1610.

Hörner J, Samuelson L (2013) Incentives for experimenting agents. *RAND J Econ* 44(4):632–663.

Hu B, Duenyas I, Beil D (2013) Does pooling purchases lead to higher profits? *Manage Sci* 59(7):1576–1593.

Jaruzelski B, Hirsh E (2016) The 2015 global innovation 1000: Automotive industry findings. Technical report, PricewaterhouseCoopers.

Kauffman S (1993) *The Origins of Order* (Oxford University Press, New York, NY).

Kauffman S, Weinberger E (1989) The NK model of rugged fitness landscapes and its application to maturation of the immune response. *J Theor Biol* 141(2):211–245.

Kavadias S, Sommer S (2009) The effects of problem structure and team diversity on brainstorming effectiveness. *Manage Sci* 55(12):1899–1913.

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

Konrad K, Kovenock D (2009) Multi-battle contests. *Game Econ Behav* 66(1):256–274.

Kornish L, Ulrich K (2011) Opportunity spaces in innovation: Empirical analysis of large samples of ideas. *Manage Sci* 57(1):107–128.

Körpeoğlu E, Cho SH (2017) Incentives in contests with heterogeneous solvers. *Manage Sci* Forthcoming.

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

Kotz S, Balakrishnan N, Johnson N (2000) *Continuous Multivariate Distributions* (John Wiley & Sons, New York, NY), 2nd edition.

Krohn S, Morthorst P, Awerbuch S (2009) The economics of wind energy. Technical report, The European Wind Energy Association.

Laffont JJ, Tirole J (1993) *A Theory of Incentives in Procurement and Regulation* (MIT Press).

Lazear EP, Rosen S (1981) Rank-order tournaments as optimum labor contracts. *J Polit Econ* 89(5):841–864.

Lewis T (2012) A theory of delegated search for the best alternative. *RAND J Econ* 43(3):391–416.

Li C (2013) Sourcing for supplier effort and competition: Design of the supply base and pricing mechanism. *Manage Sci* 59(6):1389–1406.

Li C, Debo L (2009) Second sourcing vs. sole sourcing with capacity investment and asymmetric information. *Manuf Serv Oper Manag* 11(3):448–470.

Li C, Wan Z (2016) Supplier competition and cost improvement. *Manage Sci* 63(8):2460–2477.

Lichtenberg F (1988) The private r and d investment response to federal design and technical competitions. *Am Econ Rev* 78(3):550–559.

Loch C, Terwiesch C, Thomke S (2001) Parallel and sequential testing of design alternatives. *Manage Sci* 47(5):663–678.

Manso G (2011) Motivating innovation. *J Financ* 66(5):1823–1860.

Manwell J, McGowan J, Rogers A (2009) *Wind Energy Explained: Theory, Design and Application* (John Wiley & Sons, Chicester, United Kingdom), 2nd edition.

March J (1991) Exploration and exploitation in organizational learning. *Organ Sci* 2(1):71–87.

March J, Simon H (1958) *Organizations* (Blackwell, Cmabridge, MA).

Marinovic I (2015) The credibility of performance feedback in tournaments. *J Econ Man Strat* 24(1):165–188.

Mihm J (2010) Incentives in new product development projects and the role of target costing. *Manage Sci* 56(8):1324–1344.

Mihm J, Loch C, Huchzermeier A (2003) Problem–solving oscillations in complex engineering projects. *Manage Sci* 49(6):733–750.

Mihm J, Loch C, Wilkinson D, Huberman A (2010) Hierarchical structure and search in complex organizations. *Manage Sci* 56(5):831–848.

Mihm J, Schlapp J (2017) Sourcing innovation: On feedback in contests. *Manage Sci* forthcoming.

Moldovanu B, Sela A (2001) The optimal allocation of prizes in contests. *Am Econ Rev* 91(3):542–558.

Moldovanu B, Sela A (2006) Contest architecture. *J Econ Theory* 126(1):70–96.

Nalebuff BJ, Stiglitz JE (1983) Prizes and incentives: Towards a general theory of compensation and competition. *Bell J Econ* 14(1):21–43.

Ok E (2007) *Real Analysis with Economic Applications* (Princeton University Press, Princeton, NJ).

Olver F, Lozier D, Boisvert R, Clark C (2010) *NIST Handbook of Mathematical Functions* (Cambridge University Press, New York, NY).

Oraiopoulos N, Kavadias S (2014) The path-dependent nature of r&d search: Implications for (and from) competition. *Prod Oper Manag* 23(8):1450–1461.

Palfrey T (1983) Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica* 51(2):463–483.

Pich M, Loch C, De Meyer A (2002) On uncertainty, ambiguity, and complexity in project management. *Manage Sci* 48(8):1008–1023.

Pisano G (2015) You need an innovation strategy. *Harvard Bus Rev* 6(44–54).

Rivkin J, Siggelkow N (2003) Balancing search and stability: Interdependencies among elements of organizational design. *Manage Sci* 49(3):290–311.

Rob R (1986) The design of procurement contracts. *Am Econ Rev* 76(3):378–389.

Roels G (2014) Optimal design of coproductive services: Interaction and work allocation. *Manuf Serv Oper Manag* 16(4):578–594.

Rogerson W (1989) Profit regulation of defense contractors and prizes for innovation. *J Polit Econ* 97(6):1284–1305.

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

Schlapp J, Oraiopoulos N, Mak V (2015) Resource allocation decisions under imperfect evaluation and organizational dynamics. *Manage Sci* 61(9):2139–2159.

Scotchmer S (2004) *Innovation and Incentives* (MIT Press, Cambridge, MA).

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

Siegel R (2010) Asymmetric contests with conditional investments. *Am Econ Rev* 100(5):2230–2260.

Siggelkow N, Rivkin J (2005) Speed and search: Designing organizations for turbulence and complexity. *Organ Sci* 16(2):101–122.

Simon H (1969) *The Sciences of the Artificial* (1996, MIT Press, Cambridge, MA), 3rd edition.

Slack N, Lewis M (2015) *Operations Strategy* (Pearson, Harlow, UK), 4th edition.

Sommer S, Loch C (2004) Selectionism and learning in projects with complexity and unforeseeable uncertainty. *Manage Sci* 50(10):1334–1347.

Sommer S, Loch C (2009) Incentive contracts in projects with unforeseeable uncertainty. *Prod Oper Manag* 18(2):185–196.

Sommer S, Loch C, Dong J (2009) Managing complexity and unforeseeable uncertainty in startup companies: An empirical study. *Organ Sci* 20(1):118–133.

Taylor CR (1995) Digging for golden carrots: An analysis of research tournaments. *Am Econ Rev* 85(4):872–890.

Terwiesch C, Loch C (2004) Collaborative prototyping and the pricing of custom-designed products. *Manage Sci* 50(2):145–158.

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

The Economist (2007) The Airbus A380—The giant on the runway. *The Economist* 13 October 2007:79–82.

Thomke S (1998) Managing experimentation in the design of new products. *Manage Sci* 44(6):743–762.

Thomke S (2003) *Experimentation Matters: Unlocking the Potential of New Technologies for Innovation* (Harvard Business School Press, Cambridge, MA).

Thomke S, Bell D (2001) Sequential testing in product development. *Manage Sci* 47(2):308–323.

Tong Y (1980) *Probability Inequalities in Multivariate Distributions* (Academic Press, New York, NY).

Ulbricht R (2016) Optimal delegated search with adverse selection and moral hazard. *Theor Econ* 11(1):253–278.

Van Mieghem J, Allon G (2015) *Operations Strategy* (Dynamic Ideas, Belmont, MA), 2nd edition.

Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J Finance* 16(1):8–37.

Wan Z, Beil D (2009) RFQ auctions with supplier qualification screening. *Oper Res* 57(4):934–949.

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

Weitzman M (1979) Optimal search for the best alternative. *Econometrica* 47(3):641–654.

Wheelwright S, Clark K (1992) *Revolutionizing Product Development* (The Free Press, New York, NY).

Wooten J, Ulrich K (2017) Idea generation and the role of feedback: Evidence from field experiments with innovation tournaments. *Prod Oper Manage* 26(1):80–99.

World Bank (2016) The World Bank: Procurement regulations for IPF borrowers. July 2016.

Yang Z, Aydın G, Babich V, Beil D (2009) Supply disruptions, asymmetric information, and a backup production option. *Manage Sci* 55(2):192–209.

Yang Z, Aydın G, Babich V, Beil D (2012) Using a dual-sourcing option in the presence of asymmetric information about supplier reliability: Competition vs. diversification. *Manuf Serv Oper Manag* 14(2):202–217.

Yildirim H (2005) Contests with multiple rounds. *Game Econ Behav* 51(1):213–227.

Yücesan E (2013) An efficient ranking and selection approach to boost the effectiveness of innovation contests. *IIE Trans* 45(7):751–762.