

Innovation Contests, Open Innovation, and Multiagent Problem Solving

Christian Terwiesch

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104,
terwiesch@wharton.upenn.edu

Yi Xu

Robert H. Smith School of Business, University of Maryland, College Park, Maryland 20742,
yxu@rsmith.umd.edu

In an innovation contest, a firm (the seeker) facing an innovation-related problem (e.g., a technical R&D problem) posts this problem to a population of independent agents (the solvers) and then provides an award to the agent that generated the best solution. In this paper, we analyze the interaction between a seeker and a set of solvers. Prior research in economics suggests that having many solvers work on an innovation problem will lead to a lower equilibrium effort for each solver, which is undesirable from the perspective of the seeker. In contrast, we establish that the seeker can benefit from a larger solver population because he obtains a more diverse set of solutions, which mitigates and sometimes outweighs the effect of the solvers' underinvestment in effort. We demonstrate that the inefficiency of the innovation contest resulting from the solvers' underinvestment can further be reduced by changing the award structure from a fixed-price award to a performance-contingent award. Finally, we compare the quality of the solutions and seeker profits with the case of an internal innovation process. This allows us to predict which types of products and which cost structures will be the most likely to benefit from the contest approach to innovation.

Key words: innovation; product development; open systems; contests

History: Accepted by Christoph Loch, R&D and product development; received October 26, 2006. This paper was with the authors 8 months for 2 revisions. Published online in *Articles in Advance* July 25, 2008.

1. Introduction

Innovation is at the heart of every R&D process. The detailed mechanics of R&D differ widely by industry, reflecting different cost structures, different success rates, and different market rewards, but the innovation process is remarkably similar across industries. Drug candidates in a pharmaceutical development process, TV shows in an entertainment company, and proposals in a venture capital firm all flow through a conceptually similar innovation process. This process starts with the creation of many innovation opportunities that are then evaluated in a filtering step that selects the most promising opportunity from among the candidates.

Typically, the creation of opportunities, as well as the selection from among the opportunities, happens inside an innovating firm. Inputs from various functions create new opportunities that are then selected based on input from R&D, marketing, and general management. However, there exist a rapidly growing number of innovation processes that rely on the outside world to create opportunities and then select the best from among these alternatives for further development. This approach is often referred to as

open innovation (Chesbrough 2003, von Hippel 2005, Terwiesch and Ulrich 2008).

Open innovation initiatives often rely on the altruism of its community members, their desire to compete for status within the community (Loch et al. 2000), or their self-interest, reflecting their role as a user of the innovation (von Hippel 2005). A remarkable exception to such nonfinancial motives is the innovation contest. In an innovation contest, also known as an innovation tournament, many individuals or teams submit plans or prototypes to an innovating firm. Examples of innovation contests are QVC's product road show (an opportunity for inventors to showcase their ideas and potentially get them included in QVC's assortment), the DARPA Grand Challenge for autonomous robotic vehicles (an open competition in which inventors can enter to win a substantial amount of money if their innovation outperforms others concerning speed, range, or ability to conquer difficult terrain), and the recently launched TV casting show "The Million Dollar Idea" (see Ulrich 2008).

Despite attracting a significant media attention, the importance of these innovation contests has been

rather small relative to the *traditional* innovation process. However, this is currently changing. With a growing trend toward outsourcing and off-shoring innovation-related activities (see Anderson et al. 2006, Eppinger and Chitkara 2006), innovation contests and their applications have expanded from creating *crazy* concepts to solid R&D problem solving in the recent years.

Consider the case of InnoCentive, a company that acts as an intermediary and executes hundreds of innovation contests every year for its clients. At InnoCentive, scientists can register and express what type of scientific problems they are interested in (most of them are in the areas of molecular biology and chemistry). InnoCentive acts as an intermediary between these scientists (also known as *solvers*) and those—typically, a large firm's R&D organizations—that encounter technical or scientific problems as part of their R&D process (also known as *seekers*). InnoCentive works with seekers to formulate a statement describing the problem and the rules of the innovation tournament.¹ Such a statement is made accessible to a pool of 95,000 solvers from around the world. Depending on their availability and interest, a subset of those will start working on the problem and an even smaller subset will actually submit a solution. InnoCentive will provide these solutions to the seeker, who can decide if the solution is useful to him. If it is, the seeker can acquire the intellectual property from the solver; typical rewards are between \$10 k and \$50 k.

Just like in the case of innovation contests executed by DARPA or QVC, the seeking firm obtains several benefits from this form of innovation: (a) it induces competition among solvers; (b) the seeker only pays for successful innovations, but not for the failures (the associated risks of failures are shifted to the solver); (c) the seeker gains access to a broad pool of solvers so problems are solved by those who have the most relevant expertise; (d) there exists an opportunity of wage-rate arbitrage or, more generally, cost savings; (e) there is an increase in the capacity of idea generation and testing.

These potential benefits have led companies with a long R&D tradition, such as Ely Lilly or DuPont, to use InnoCentive's innovation tournaments for a growing portion of their R&D work. InnoCentive's business model of innovation has been praised in the business press (Taylor 2006, Hempel 2006, Huston and Sakkab 2006) and has received several innovation

awards (including the "Business Process Award" from the *Economist's* Innovation Summit and the "Infosys Transformation Award" from The Wharton School).

Despite this growing popularity, little remains known about when such innovation contests should be used and how innovation contests should be executed. In the absence of appropriately designed contracts (rules of the tournament), any form of a decentralized system will lead to inefficiencies reflecting information and incentive problems. Our research questions aim to address these problems:

- What type of innovation problems are most suited to be solved by innovation contests and what problems are better solved internally?
- For a given type of innovation problem, what is the optimal design of the innovation contest? Specifically, what is the optimal award, and how many solvers should the seeker attempt to reach?

Providing answers to these questions is the contribution of this paper. We combine prior research from the field of contests and tournaments (e.g., Moldovanu and Sela 2001) with models of product development and search (e.g., Dahan and Mendelson 2001, Terwiesch and Loch 2004). We provide the following novel results. First, we derive the optimal innovation contest award mechanism (Theorem 1A) and show how the quality of the submissions and the profits for the seeker depend on the number of potential solvers (Theorem 1B). Having many solvers work on an innovation problem will lead to a lower equilibrium effort for each solver, which is undesirable from the perspective of the seeker. Although prior economics research has argued that it is optimal to restrict the number of participants to reduce this effect, we derive an additional benefit of having a large pool of solvers: the seeker can benefit from a larger solver population because he obtains a more diverse set of solutions, which mitigates and sometimes outweighs the effect of underinvestment from each solver (Theorem 1C). Second, the inefficiency of the innovation contest resulting from the solvers' underinvestment can be further reduced by changing the award structure of the innovation contest. Whereas prior research has advocated the use of fixed-price rewards, we show that a performance-contingent award can lead to better solutions, higher seeker-profits, and system efficiency (Theorem 2). Third, we compare innovation contests hosted by the seeker and by an intermediary with the case of an internal innovation process (Theorems 3A and 3B). This allows us to predict which types of products and which cost structures will be the most likely to benefit from the contest approach to innovation. We also show that the seeker primarily benefits from open innovation by obtaining higher performance and not only by obtaining lower costs. These results,

¹ To use an innovation contest as studied in this paper, the seeker has to be able to provide a clear description of the problem. If the problem is highly complex with ill-defined interfaces, it is not suitable for an innovation contest because the resulting coordination costs would be too high (e.g., Ulrich and Ellison 1999, Novak and Eppinger 2001, Mihm et al. 2003).

we believe, are of interest to firms executing innovation contests as well as those participating or considering to participate as either seeker or solver. Moreover, in light of the ongoing discussions concerning the off-shoring of R&D work, we believe that our results should also be of interest to a much broader audience.

The remainder of this article is organized as follows. We first review the relevant literature (§2),² followed by the development of our modeling framework (§3). Sections 4–6 establish our main results (Theorems 1–3), and §7 concludes the paper.

2. Relevant Literature

Our model, presented in §3, combines research on product development processes with research on the economics of contests and tournaments. Over the last 15 years, a number of product development process models were created (e.g., Krishnan and Ulrich 2001, Loch and Kavadias 2007). Although the detailed mathematics of the various models differ, they collectively suggest (a) that product development should be modeled as a stochastic process and (b) that there exist different types of product development problems.

The stochastic aspect of the product development process has been modeled as a search process in an array of binary variables (Loch et al. 2001), sequential draws from a distribution (Terwiesch and Loch 2004, based on work by Weitzman 1979), parallel draws from a distribution (Dahan and Mendelson 2001, based on an application of extreme value models), or a series of Bernoulli trials (e.g., Ha and Porteus 1995). Our model applies the previous work by Dahan and Mendelson in that we view the innovations undertaken by one solver as a set of parallel experiments. The performance outcome of this experimentation is the highest realization of the parallel draws (i.e., the n th order statistics or the extreme value model). This corresponds to a single-period model, where a solver makes the decision on how much to experiment only once.

One of the major accomplishments of the prior product development literature has been to demonstrate that the type of product development problem matters greatly and should influence the process of searching for an optimal solution. Problem types differ along multiple dimensions, including (a) the amount of uncertainty in the overall pay-off function, and hence the solver's ability to predict the outcome of an experiment; and (b) the ability of the solver to learn from one experiment to another. Loch et al.

(2006) discuss different problem types and their implications for managing the associated risks.

The second stream of research that we draw from relates to the design of contests and tournaments (e.g., Glazer and Hassin 1988, Lazear and Rosen 1981). This research has a long tradition in economics and has recently seen a number of applications in operations management (e.g., Deng and Elmaghraby 2005) and marketing, especially in the salesforce domain (e.g., Kalra and Shi 2001, Chen and Xiao 2005). There exist, however, two crucial differences between a salesforce contest and an innovation contest. First, the seeker in an innovation contest is interested in maximizing the value of the highest performance outcome. The seeker in a salesforce contest, in contrast, is interested in maximizing the sum across all outcomes. Put differently, an R&D department prefers 100 bad ideas and 1 outstanding idea over 101 good ideas, whereas a marketing department prefers 101 salespeople with good revenues over 100 salespeople with bad revenues and 1 salesperson with outstanding revenue. Second, participation decisions for solvers are fully voluntary in an innovation contest, whereas salespeople are forced to participate in a salesforce contest.

There exists a small set of papers in the economics literature that have applied contests and tournaments to R&D settings. Taylor (1995) and Fullerton and McAfee (1999) study the optimal design of research tournaments with a sequential stochastic model. These papers focus on the competition among symmetric solvers: the seeker benefits from buyer power as the solvers are competing against each other. They find that the contest suffers from underinvestment in effort by the solvers. To mitigate this effect, these models suggest limiting the pool of solvers, potentially all the way down to two (Fullerton and McAfee 1999). A two-solver contest is sufficient to induce competition while leaving a 50% probability of winning to two symmetric firms.

3. Model Development

We consider an innovation problem in which the performance of the solution can be measured in a one-dimensional space. The assumption of a one-dimensional performance measure is common in product development. This one-dimensional space could reflect a technical specification (e.g., the purity of a material obtained in a chemical reaction) or a consumer's utility measure (e.g., the expressed purchase intent). Similar one-dimensional settings are considered by Dahan and Mendelson (2001) and Terwiesch and Loch (2004).

The performance obtained from a solver is driven by three variables. First, each solver i is endowed with an expertise, β_i , which is a measure of his experience

²An extended literature review is provided in the electronic companion, which is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

and knowledge for a particular problem. For example, everything else being equal, the solution to a chemical engineering contest is more likely to be found by a chemist than by a biologist. This endowed knowledge is available to the solver at no cost.

Second, each solver can enhance the performance of his solution(s) by investing improvement effort, e_i . Such improvement effort corresponds, for example, to conducting a thorough patent search and literature review, or to implementing rigorous quality control systems with high standards. Effort e_i leads to a deterministic improvement $r(e_i)$ of the performance of the solution, where $r(e_i)$ is an increasing and concave function in e_i that measures the performance return on the improvement effort. Let $c_1 e_i$ be the costs associated with the improvement effort of solver i .

Third, problem solving in innovation is often stochastic, which we capture by adding a noise variable, ξ , to the performance. Given this uncertain performance, the solver will most likely engage in a search process by conducting a set of trials and experiments (see, e.g., Loch et al. 2001, Dahan and Mendelson 2001). Let m_i be the number of experiments conducted by solver i . The results of an experiment are captured by the multiple realizations of the random variable, ξ . Following the work by Dahan and Mendelson (2001), we consider the specific case in which the random noise ξ is an independent and identically-distributed Gumbel random variable with mean zero and scale parameter μ . Note that a higher μ increases the variance of a draw. The associated costs are $c_2 m_i$.

Given an expertise, β_i , an improvement effort, e_i , and an experimentation effort, m_i , the performance of the solution is assumed to be of the following additive form:³

$$v_i(\beta_i, e_i, m_i, \xi_i) = \max_j \{v_{ij} = \beta_i + r(e_i) + \xi_{ij}, j = 1, 2, \dots, m_i\}. \quad (1)$$

Note that our performance function, shown in, (1) is rather general, because it includes a baseline performance, a deterministic reward for effort, and a stochastic reward for effort.

The above general performance function nicely blends two important features of an innovation project: heterogeneity in solver expertise (i.e., different solvers have different β_i s) and a stochastic relationship between efforts and performance. Unfortunately,

the analytical tractability of such a general performance function is quite limited. For this reason, we decompose the general performance function (1) into three interesting and tractable special cases based on which of the three terms dominates.

Expertise-based projects ($\xi_{ij} = 0$) have no stochastic influence of the random noise, and thus experimentation is not necessary. Performance is driven by expertise and improvement effort. Thus,

$$v_i(\beta_i, e_i) = \beta_i + r(e_i).$$

Such projects are low-risk projects with little novelty in them, such as converting a computer-aided design drawing into another format or designing a process recipe for a commonly used chemical reaction. This type of concave, deterministic performance function is used in some of the work on software development contracting (e.g., Whang 1992). Solvers are heterogeneous in their expertise. We assume that both the seeker and the solvers have identical beliefs that β_i is distributed across solvers with cumulative density function $F(\beta)$ and probability density function $f(\beta)$. It could be possible that an expert or a solver with higher expertise has better information on the distribution of a competitors' expertise level. We leave that case for future research. We will use a Gumbel distribution with scale parameter λ as a special case for F to illustrate some results in the paper. Thus, the expertise-based project is essentially an auction model with $v_i(\beta_i, e_i)$ as a solver i 's bid. Although the performance is certain for a solver, a seeker still faces uncertainty with respect to the performance of the best solution obtained from a pool of external solvers because of the heterogeneity in endowed solver expertise.

Ideation projects ($\beta_i = \beta, m_i = 1$) are broad and nondetailed innovation problems for which the seeker looks for novel ideas. For example, recent InnoCentive challenges included "a product concept for a child-proof container of medication" or "the design of the next generation binder." Other examples of such projects include design contests for the aesthetics of a new building or the logo for an event such as the FIFA world cup. In these projects, the seeker's taste, which is uncertain for the solver, plays an important role in determining what constitutes a good solution. Hence, in an ideation project, the performance of a solution has a significant noise term that reflects heterogeneity of solvers' solutions in matching the seeker's taste. All solvers are identical in terms of endowed expertise, i.e., $\beta_i = \beta_j = \beta$ for all i and j , that is, all solvers are equally capable for such a broad problem ex ante. As before, solvers can spend effort to increase the quality of their solution (e.g., by building a sophisticated prototype of their idea as opposed to simply submitting a

³Of course, more complex functional forms could be considered. A purely multiplicative form could be converted into an additive form by taking logarithms and then appropriately rescaling the cost functions. A mixture of additive and multiplicative terms (e.g., an interaction term between expertise and deterministic effort) could also be analyzed, but certainly would require numerical solutions.

sketch on paper). Because the noise term will emerge after the seeker reviews all submitted solutions, there is no point for solvers to invest in experimentation. In this case, the performance of a solution is

$$v_i(e_i, \xi_i) = \beta + r(e_i) + \xi_i.$$


Trial-and-error projects ($\beta_i = \beta, e = 0$) are innovation problems with an extremely rugged solution landscape. The solver cannot anticipate the performance of a solution before actually conducting the experiment. However, unlike for ideation problems, the solver can observe the performance of a trial before submitting a solution. Thus, the uncertainty the solver faces is entirely technical in its nature. Given the ruggedness of the solution landscape, every experiment conducted has the same expected performance, i.e., $\beta_i = \beta_j = \beta$ for all i and j . Performance is driven by the experimentation effort, and there exists no learning from one round of experimentation to the other. The solver exerts effort by experimenting (which increases the performance of the best solution stochastically), and there exists no way of obtaining a deterministic return to effort. This case is equivalent to the model presented by Dahan and Mendelson (2001):

$$v_i(m_i, \xi_i) = \max_j \{v_{ij} = \beta + \xi_{ij}, j = 1, 2, \dots, m_i\}.$$

Table 1 summarizes the characteristics, mathematical representations, and examples of these three project types. Figure 1 separates the uncertainty faced by the solver into a technical uncertainty dimension and a market uncertainty dimension.

Given the project type, the seeker faces two decisions. First, the seeker needs to decide if the problem should be solved internally or whether it should be posted to a broader community in the form of

Figure 1 The Exposures of Different Project Types to Technical Uncertainty and Market Uncertainty



an innovation contest. If the problem is solved internally, the seeker needs to decide upon the two types of effort defined above.

If the problem is posted to a broader community, the general sequence of events is as follows. The seeker needs to determine an award allocation mechanism. This mechanism is announced (together with the problem) to all solvers. Each solver i has a privately known expertise of β_i for the problem. Solvers are risk neutral and face three types of costs: the costs of improvement effort ($c_1 e_i$), the costs of experimentation ($c_2 m_i$), and a fixed cost of participation c_f if the solver elects to work on the problem. Given solver i 's efforts (e_i, m_i) the performance of her best solution, v_i , is determined according to Equation (1) and the solution is submitted to the seeker.

Based on the performance vector of submitted solutions from n solvers, $\mathbf{v} = (v_1, v_2, \dots, v_n)$, the seeker awards the solvers according to the announced mechanism. We assume the seeker's payoff to be a weighted combination of the performance of the best

Table 1 Summary of the Characteristics, Mathematical Representations, and Examples of Expertise-Based Projects, Ideation Projects, and Trial-and-Error Projects

Project type	Characteristics of the project	Action taken by solver	Variables determining performance
Expertise-based project	Engineering tasks with no uncertainty in performance function (well-behaved solution landscape) Example: Modify an existing process design to fit a new production site	Invest effort to enhance the existing expertise relevant to the project	Endowed expertise (β_i) Effort (e) $v_i(\beta_i, e_i) = \beta_i + r(e_i)$
Ideation project	Innovative problems with no clear specifications, leading to uncertainty in the performance function Example: Design next-generation binder	Invest effort to create the best possible presentation	Effort (e) Subjective taste of seeker (market uncertainty) $v_i(e_i, \xi_i) = \beta + r(e_i) + \xi_i$
Trial-and-error project	Solutions to research problems with well-defined goals, yet highly rugged solution landscapes, creating uncertainty in how to improve a solution Example: A pill that reduces grey hair	Experiment by trying out many solutions and then picking the one with the highest performance	Number of experiments (m) Outcome of each trial (technical uncertainty) $v_i(m_i, \xi_i) = \max_j \{v_{ij} = \beta + \xi_{ij}, j = 1, 2, \dots, m_i\}$

solution and the expected average performance of all solutions:

$$V = \rho \max_{i=1, \dots, n} \{v_i\} + (1 - \rho) \frac{\sum_{i=1, \dots, n} v_i}{n},$$

where $0 \leq \rho \leq 1$. This formulation includes the special case in which the seeker is only concerned about the value of the best solution ($\rho = 1$). However, it also lets us explore the case in which the seeker cares about the overall quality of the submitted solutions. For example, the solver might benefit from combinations of the submitted solutions, and therefore favor a high average solution quality. In the extreme case, the solver might only care about the average submitted solution quality ($\rho = 0$). This captures the case of the traditional salesforce contest (see, e.g., Chen and Xiao 2005) or a contest in which the seeker obtains a (potentially negative) payoff for every submitted solution, as is the case in television contests such as American Idol.⁴

Following the common settings in the contest literature, all parameters including β , n , the performance function, the return function on effort, the distribution of solver expertise, and costs are common knowledge to both the seeker and solvers,⁵ except a solver's expertise in an expertise-based project that is known to that solver only. In addition, solvers' efforts are not observable and verifiable to the seeker. The seeker attempts to maximize the expected payoff, V , net of the costs of the award and the cost of internal development effort (note that one of the two is zero). Each solver attempts to maximize the expected net profit consisting of the expected award minus the costs of effort.

4. Open Innovation with Fixed-Price Contest

In a fixed-price contest, the seeker announces a prespecified award with a fixed amount, A , which will be granted to solvers according to a preannounced award allocation structure. The fixed-price

⁴ There exist other functional forms that could capture the seeker's interest in more than the optimal solution. Specifically, it seems plausible that the seeker might be interested in the best m submitted solutions. These cases lead to qualitatively similar results, yet are analytically intractable.

⁵ These assumptions are common in the literature on auctions and contests (e.g., Snir and Hitt 2003 use very similar assumptions). As are all assumptions, they are a simplification of the real world, made out of analytical convenience rather than based on empirical observations. For example, in reality, the exact n will not be known to all parties. The seeker could obtain an estimate of n from InnoCentive. The solvers, however, will have some prior distribution concerning n , which leads to a rational expectation about n . As long as solvers have symmetric priors, we get to the same results that we have now (although the model would be more complex).

contest is the most commonly adopted mechanism in open innovation systems and is the standard contest used by InnoCentive. Theorem 1A shows that instead of splitting a predetermined total award amount into two smaller awards, it is optimal to allocate the entire award to the best solution. All proofs are provided in the online appendix (provided in the e-companion).

THEOREM 1A. *For a given amount of award A , assume that the seeker can allocate it to at most two solvers, with $A_1 \geq A_2 \geq 0$ and $A_1 + A_2 = A$. For ideation projects and trial-and-error projects, it is optimal for the seeker to grant the entire award to the solver with the best solution (i.e., $A_1 = A$ and $A_2 = 0$), although it may or may not be optimal for the seeker to do so for expertise-based projects.*

The above theorem establishes the optimality of the winner-takes-all award structure for an open innovation contest with risk-neutral seeker and solvers. For ideation projects and trial-and-error projects, because solvers are symmetric in endowed expertise, they react to award structures symmetrically. When the solvers are risk neutral, the marginal incentive generated by one more dollar of award for the first prize is higher than the marginal incentives generated by one more dollar of award for one of the lower prizes. As a result, concentrating the award on the first prize will generate the strongest incentives for solvers to exert efforts. For expertise-based projects, because solvers have different endowed expertise, they react differently to reward structures. A winner-takes-all award structure offers stronger incentive to solvers with high endowed expertise to exert effort because they are more likely to win the single award, whereas a multiple-prize award structure is more attractive to solvers with low endowed expertise because they have not much chance to win the first prize. Therefore, whether a winner-takes-all award is optimal depends on the distribution of the solver expertise. In the proof of the theorem, we derive a mild, necessary, and sufficient condition for the winner-takes-all contest to be optimal for expertise-based projects. We will focus on the analysis of the winner-takes-all contest in this paper.

In a winner-takes-all contest, observing the seeker's award A , the n solvers simultaneously make participation and effort decisions to ensure that the expected profit they could earn from entering the contest is at least c_f (i.e., a solver's reservation profit is zero). Each participating solver submits his (best) solution to the seeker for review. The solver who produced the best solution among all n solvers will win the award A , whereas all other solvers will not be awarded anything.

Consider a solver i with endowed expertise β_i . Observing the seeker's award A and his own expertise β_i , solver i needs to decide how much improvement effort, e_i , to exert (leading to cost $c_1 e_i$) and the

number of experiments, m_i , to conduct (leading to cost $c_2 m_i$). Following the existing literature on contests, we focus on a symmetric equilibrium throughout the paper. Thus, in determining the effort level, solver i solves

$$\begin{aligned} & \max_{e_i \geq 0, m_i \geq 1} \pi_i(e_i, m_i | n, e, m) \\ & = \delta_i (A \Pr\{\text{solver } i \text{ wins}\} - c_1 e_i - c_2 m_i - c_f) \\ & = \delta_i (A \Pr\{v_i(\beta_i, e_i, m_i, \xi_i) \\ & \quad > \max\{v_j(\beta_j, e, m, \xi_j), \quad j \neq i \text{ and } \delta_j = 1\} \\ & \quad - c_1 e_i - c_2 m_i - c_f\}, \end{aligned}$$

where $\delta_i = 0$ if solver i decides not to participate and $\delta_i = 1$ otherwise. Let $\mathcal{S}^*(A) = \{\delta^*(A), e^*(A), m^*(A)\}$ be the solvers' equilibrium strategy for a given award A . We can then write the seeker's problem as

$$\begin{aligned} & \max_{A \geq 0} \Pi(A | n, \mathcal{S}^*(A)) \\ & = \rho \max\{\delta_k^*(A) v_k(\beta_k, e^*(A), m^*(A), \xi_k), k = 1, 2, \dots, n\} \\ & \quad + (1 - \rho) \left(\frac{\sum_k \delta_k^*(A) v_k(\beta_k, e^*(A), m^*(A), \xi_k)}{n} \right) - A. \end{aligned}$$

For the resulting equilibrium with given number of solvers n , Theorem 1B establishes for each of the three project types the solver entry pattern, the amount of prize the seeker awards, the amount of effort that each solver exerts, and the expected profit the seeker earns. For the solver entry pattern, Theorem 1B derives the expected number of participating solvers in a free-entry fixed-price contest, n^{e*} for expertise-based projects; and the maximum number of solvers a free-entry fixed-price contest can accommodate, n^{i*} and n^{t*} for ideation projects and trial-and-error projects, respectively.⁶

The exposition of our results is much simpler with a specific functional form for the effort function $r(e)$:⁷ $r(e) = \theta \ln e$. Recall that for trial-and-error projects, there exists by definition no improvement efforts, and hence no deterministic improvements ($r(e) = 0$). However, for a trial-and-error project, given a solver's experimentation effort m , the expected performance

⁶ In a free-entry contest, the seeker does not impose any restrictions on solvers' entry to the contest. It is completely up to a solver's own choice whether to enter the contest or not. n^{e*} , n^{i*} , and n^{t*} are the expected or maximum number of solvers that a free-entry contest can accommodate for each type of project, such that every participating solver can earn a nonnegative expected profit. It is possible that n^{e*} , n^{i*} , and n^{t*} are small due to factors such as high fixed cost for solvers. It is also worth noting that a free-entry contest is not necessarily cost-free to solvers. For example, solvers may incur the fixed cost c_f to enter a free-entry contest in our model.

⁷ Our key results are proven in the general case in the online appendix.

of his best solution is $\mu \ln m$ (see the online appendix for the detailed derivation), which is analogous to a logarithmic return function $r(m) = \mu \ln m$. Therefore, for expertise-based projects and ideation projects, θ can be viewed as the return on effort coefficient, and μ can be viewed as the return on effort coefficient for trial-and-error projects. For ideation projects, from the solvers' perspective, μ also measures the stochastic effect (variance) of the seeker's evaluation.

Throughout the paper we use superscripts $\{e, i, t\}$ for the expertise-based project, the ideation project, and the trial-and-error project, respectively.

THEOREM 1B. *In a fixed-price open innovation contest with n solvers, the unique equilibrium is defined as follows:*

(i) *For expertise-based projects, only solvers with endowed expertise that is higher than $\beta_f = F^{-1}((c_f/A^{e*})^{1/(n-1)})$ will participate, where A^{e*} is the optimal award. The improvement effort of a participating solver with expertise $\beta \in [\beta_f, \bar{\beta}]$ is*

$$e^*(\beta) = \frac{A^{e*} F(\beta)^{n-1} - c_f}{c_1}, \quad (2)$$

and the expected number of participating solvers in a free-entry fixed-price open innovation contest is

$$n^{e*} = n(1 - (c_f/A^{e*})^{1/(n-1)}).$$

If $c_f = 0$, the award is $A^{e} = \theta$ and the expected profit for the seeker is*

$$\begin{aligned} \Pi^{e*} & = \rho \int_{\beta_f}^{\bar{\beta}} \beta n F(\beta)^{n-1} f(\beta) d\beta + (1 - \rho) \int_{\beta_f}^{\bar{\beta}} \beta f(\beta) d\beta \\ & \quad + \theta \left(\ln \frac{\theta}{c_1} - \frac{(n-1)(\rho + (1-\rho)n) + n}{n} \right). \end{aligned} \quad (3)$$

(ii) *For ideation projects, the improvement effort of the solver is*

$$e^* = \frac{\theta^2(n-1)}{\mu c_1 n^2}, \quad (4)$$

the award is $A^{i} = \theta$, the expected profit for the seeker is*

$$\Pi^{i*} = \beta + \theta \left(\ln \frac{\theta^2(n-1)}{\mu c_1 n^2} - 1 \right) + \rho \mu \ln n, \quad (5)$$


and the maximum number of participating solvers in a free-entry fixed-price open innovation contest is

$$n^{i*} \approx \frac{\theta(\mu - \theta)}{\mu c_f}.$$

(iii) *For trial-and-error projects, the experimentation effort of the solver is*

$$m^* = \frac{\mu(n-1)}{n^2 c_2},$$

Figure 2 Equilibrium Probability of Winning for a Particular Solver, Effort (e^* or m^*), the Seeker's Optimal Expected Profit as a Function of the Number of Solvers n



Notes. β has a zero mean Gumbel distribution with $\lambda = 2$ for expertise-based projects, $\beta = 1$ for ideation and trial-and-error projects, $\rho = 1$, $c_f = 0.1$, $\mu = 2$, $\theta = 1$, $c_1 = 0.1$, and $c_2 = 0.1$. Figures illustrate only the curvatures, not the scales.

the award is $A^* = \mu$, the expected profit for the seeker is

$$\Pi^{t*} = \beta + \mu \left(\ln \frac{\mu(n-1)}{nc_2} - 1 \right) - (1-\rho)\mu \ln n, \quad (6)$$

and the maximum number of participating solvers in a free-entry fixed-price open innovation contest is

$$n^{t*} = \sqrt{\frac{\mu}{c_f}}$$

For all types of projects, the equilibrium solver effort is increasing in the award amount A and decreasing in effort cost parameters c_1 or c_2 . As effort cost parameter c_1 or c_2 increases, the seeker's expected profit can decrease to zero. In this case, the problem solving for the project does not suit the contest model. Figure 2 illustrates how the equilibrium probability of winning for one solver effort, and the seeker's expected profit change with respect to the size of the solver population, n . An interesting observation is that from the equilibrium solver strategies (Figures 2(d)–2(f)), we can see that there exists a negative externality among solvers in all three projects: for a given award A , the more solvers participate in the open innovation contest, the less effort each solver exerts in equilibrium.

For an expertise-based project, the equilibrium effort e^* defined in (2) is decreasing in n because $F(\beta) \leq 1$. The negative externality effect in an expertise-based project is severe because the equilibrium effort $e(\beta, A)$ is a decreasing power function of n

(see Figure 2(d)). For an ideation project, because the return function r is concave and $n^2/(n-1)$ is increasing in n , the equilibrium effort e^* defined in (4) is also decreasing in n , but at a slower rate. For a trial-and-error project, the equilibrium number of parallel experiments m^* conducted by each firm is decreasing in n at an even slower rate.

The intuition behind this negative externality reflecting an underinvestment in solver effort is that the more solvers participate in the contest, the lower the probability of winning for a particular solver (see Figures 2(a)–2(c)). With lower winning probabilities, the solvers' expected profits decrease, leading to weaker incentives for them to exert higher efforts. This underinvestment in effort leads to an inefficiency in an open innovation system.

A similar argument can be made for the equilibrium effort (4) for an ideation project with respect to parameter μ , which from the solvers' perspective measures the stochastic effect of the seeker's taste. The equilibrium effort (4) is decreasing in μ : with increasing variance, the effect of effort on the probability of winning is decreasing; the winner is chosen by luck, not by his exerted effort e .

The negative externality among solvers also has an interesting impact on the solver entry pattern. Consistent with intuition, the number of participating solvers is decreasing in the fixed cost c_f irrespective of project type. For a higher fixed cost c_f , a solver needs to earn a higher expected profit to break even. Because the equilibrium expected profit a solver can

earn is decreasing in the number of contestants, n , higher fixed costs lead to fewer participating solvers. For expertise-based projects, the fixed cost excludes solvers with low endowed expertise from entering a contest. The expected number of solvers participating (because solver expertise is a random variable) n^{e*} is decreasing in the fixed cost c_f . For ideation projects and trial-and-error projects, the number of solvers participating in a free-entry fixed-price open innovation is no more than a threshold n^{i*} and n^{t*} , respectively. Both thresholds are also decreasing in the fixed cost c_f .⁸

To mitigate underinvestment in effort caused by the negative externality among solvers, existing literature on R&D tournaments (Taylor 1995, Fullerton and McAfee 1999, Che and Gale 2003) suggests that a free-entry R&D tournament is in general not optimal, and it is necessary to restrict the size of an R&D tournament, potentially all the way down to two (Fullerton and McAfee 1999, Che and Gale 2003). Theorem 1C shows that this result cannot be straightforwardly transferred to innovation contests as described in this paper.

THEOREM 1C. *Consider a fixed-price open innovation contest with n solvers. (i) If the seeker's weight on the performance of the best solution, ρ , is high enough, the free-entry open innovation contest can be optimal for all three types of projects. When the seeker's objective is to maximize the performance of the best solution ($\rho = 1$), the free-entry open innovation contest is always optimal for ideation projects and trial-and-error projects. It is also optimal for expertise-based projects with Gumbel-distributed solver expertise with scale parameter λ if $\lambda \geq \theta/2$. (ii) If the seeker's weight on the performance of the best solution, ρ , is sufficiently low, a free-entry open innovation contest may not be optimal. When the seeker's objective is to maximize the average performance of all solutions ($\rho = 0$), free-entry open innovation contest is not optimal for any of the three types of projects.*

Theorem 1C indicates that whether free-entry open innovation is optimal for the seeker critically hinges on the seeker's objective. If the seeker primarily cares more about the performance of the best solution, free-entry open innovation is optimal. In this case, the seeker obtains a unique benefit of having more participants for ideation and expertise-based projects: higher

solver diversity. This is consistent with an empirical study on 166 scientific problems posted on InnoCentive's website by Lakhani et al. (2007), who find that problem-solving success is related to the ability to attract specialized solvers with diverse scientific interests. When free-entry open innovation is optimal, Theorem 1B and Theorem 1C also imply that the seeker might be better off by subsidizing part of the fixed cost for solvers to encourage entry.

For expertise-based projects, the effect of higher solver diversity is captured by the term $\int_{\beta}^{\beta} \beta n F(\beta)^{n-1} f(\beta) d\beta$ in the seeker's expected profit (3). This expression can be thought of as the expected value of the highest expertise among the n solvers. The more solvers, the higher the best solver's expertise is likely to be. For a Gumbel-distributed noise variable with mean β_0 and scale parameter λ , the expected highest expertise among n solvers is $\beta_0 + \lambda \ln n$. Hence, for an expertise-based project with logarithmic return functions, λ can be viewed as the coefficient of return on having more solvers, whereas θ is the coefficient of return on higher effort. The condition $\lambda \geq \theta/2$ in the Theorem 1C basically says that when the return on diversity is strong enough relative to the return on effort, the diversity benefit overcomes the undesirable effect of underinvestment in solver effort that is associated with having a larger solver pool (see Figure 2(g)).

For an ideation project, this benefit of higher solver diversity is reflected by the term $\mu \ln n$ in the seeker's expected profit (5). Not knowing the seeker's taste, solvers build one single prototype. Having more participants consequently increases the total amount of experimentation, and as a result the best solution provides a better fit to the seeker's taste. For ideation projects, Theorem 1B indicates that the diversity benefit of having more solvers outweighs the negative externality effect of solvers' underinvestment in effort such that free-entry open innovation is optimal (see Figure 2(h)).

Free-entry open innovation is always optimal for trial-and-error projects. This is due to a unique feature of the trial-and-error project: parallel experiments are perfectly cumulative across solvers. From the seeker's point of view, solver A conducting 4 parallel experiments and solver B conducting 6 parallel experiments is equivalent to one solver conducting 10 parallel experiments. This perfect cumulation of effort across solvers is reflected by the term $\mu \ln nm^*$ in the seeker's expected profit (6). Thus, although each solver's equilibrium number of experiments is decreasing in n , the total number of experiments across all solvers, nm^* , is actually increasing in n . As a result, the seeker's optimal expected profit increases as n becomes larger (see Figure 2(i)).

⁸ When $c_f = 0$, we have $n^{e*} = n$, which implies that for any given number of solvers n , all solvers would participate in a fixed-price open innovation contest for an expertise-based project. When $c_f = 0$, we have $n^{i*} = \infty$ and $n^{t*} = \infty$, which also imply that for an ideation project and a trial-and-error project, a fixed-price open innovation contest can accommodate any given number of solvers. Therefore, when we study the case where the number of solvers is a decision variable for the seeker with $c_f = 0$ in §7 (n_s^{e*} , n_s^{i*} , and n_s^{t*}), this decision will not be bounded above.

The next observation further illustrates the benefit of a free-entry open innovation to the seeker. Normally, the sponsor of a contest could increase its payoff by charging contestants a collectable entry fee (e.g., Taylor 1995). In contrast, the next corollary shows that in the innovation contest we study, and a seeker who solely focuses on the performance of the best solution, for ideation projects and trial-and-error projects, it is in the seeker's best interest to charge zero entry fee.

COROLLARY. *Consider the case that the seeker can charge an entry fee from a participating solver. If the seeker's objective is to maximize the performance of the best solution ($\rho = 1$), for ideation projects and trial-and-error projects, the optimal entry fee is zero.*

From a solver's perspective, the effect of an entry fee is the same as the effect of fixed costs. However, from the seeker's perspective, unlike the fixed cost, the seeker collects entry fees paid by participating solvers. One might expect that the seeker could do better by charging such an entry fee. The above result indicates that for ideation projects and trial-and-error projects with $\rho = 1$, the benefit of higher diversity of a large solver pool is so strong that the seeker should charge no entry fee, and thus promote a maximum level of participation from the solvers. For expertise-based projects, corresponding analytical results cannot be obtained. However, straightforward numerical examples demonstrate that the same conclusion holds.⁹

5. Enhancing the Efficiency of Open Innovation

The results of §4 show that for innovation contests as defined in this paper, contrary to the classic economics result for contests, free entry (large n) can be optimal. We now investigate the applicability of another standard result to our specific setting: the optimality of fixed-price awards. Specifically, we study the impact of choosing an alternative award mechanism, a performance-contingent award, on the seeker's profits. One way to implement a performance-contingent award is through a royalty contract. For example, the office product retailer Staples has recently conducted large-scale innovation contests and rewarded successful solvers by allocating them a percentage of the associated profits instead of granting them a fixed reward (see Bulkeley 2006). For ease of exposition, we

consider the case in which the seeker's goal is to maximize the performance of the best solution ($\rho = 1$) and the solver's fixed cost of participation is zero ($c_f = 0$).

In a contest with a performance-contingent award, the winner is awarded a proportion ϕ , of the performance of his solution (where $0 < \phi < 1$). After observing the seeker's proposed award share ϕ , the n solvers simultaneously choose the level of effort to exert. Each participating solver will then submit his solution.

Consider a solver i with endowed expertise β_i . Given the seeker's award share ϕ and his own expertise β_i , solver i needs to decide how much improvement effort, e_i , and experimentation effort m_i to exert. Suppose all other solvers exert efforts e and m . Then, the general problem for solver i in a contest with a performance-contingent award can be written as

$$\begin{aligned} & \max_{e_i \geq 0, m_i \geq 1} \pi_i(e_i, m_i) \\ & = \phi(\max\{v_j(\beta_j, e, m, \xi_j), \forall j\}) \times \Pr\{v_i(\beta_i, e_i, m_i, \xi_i) \\ & > \max\{v_j(\beta_j, e, m, \xi_j), j \neq i\} - c_1 e_i - c_2 m_i. \end{aligned}$$

Let $e^*(\phi)$ and $m^*(\phi)$ be the solver's equilibrium strategy for a given award share ϕ .

Given the solvers' equilibrium strategy, the seeker's problem is to choose the award share ϕ to maximize the expected payoff, which is just $1 - \phi$ percent of the expected performance of the best solution generated from the open innovation contest. Therefore, the seeker's problem is

$$\begin{aligned} & \max_{0 < \phi < 1} \Pi_p(A_p, \phi) \\ & = (1 - \phi) \max\{v_k(\beta_k, e^*(\phi), m^*(\phi), \xi_k), k = 1, 2, \dots, n\}. \end{aligned}$$


The following theorem shows that with the logarithmic return function $r(e) = \theta \ln e$, a contest with a performance-contingent award is more profitable to the seeker than fixed-price awards.

THEOREM 2. *For ideation projects and trial-and-error projects, the seeker makes a higher expected profit in an open innovation contest with a performance-contingent award than in an open innovation contest with a fixed-price award. For expertise-based projects, the seeker may or may not benefit from an open innovation contest with a performance-contingent award.*

For both ideation and trial-and-error projects, a contest with a performance-contingent award enhances the efficiency of open innovation (Figure 3). In such a contest, the amount of award a solver potentially can win depends on the realized performance of the solver's solution, which is stochastically increasing in his effort. From a solver's point of view, a performance-contingent award presents two incentives to increase effort: Exerting higher effort not only

⁹ However, it is worth noting that our results only establish the optimality of free entry for fixed-price contests. There could exist more complicated mechanisms (with or without entry fee) such as the one we will discuss in the next section that can do better than the free-entry open innovation contest with fixed-price award.

Figure 3 The Seeker's Expected Profits with Performance-Contingent Award and Fixed-Price Award as Functions of the Number of Solvers n for Ideation Projects and Trial-and-Error Projects



Note. $\mu = 2$, $\theta = 1$, $c_1 = 0.1$, $c_2 = 0.1$, and $\beta = 1$.

increases the probability of winning, but also the amount of award that is received in the case of winning. The second incentive is missing in a fixed-price contest. Thus, a performance-contingent award creates a stronger incentive compared to a fixed-price award. Figure 3 also indicates that the benefits of a performance-contingent award diminish as the number of solvers becomes larger. With a larger number of solvers, both the negative externality effect and the diversity benefit are so strong that all other factors become negligible. This observation suggests that for ideation projects and trial-and-error projects, using a simple fixed-price award would not sacrifice much profit for the seeker in large contests. However, for small contests with a limited number of solvers, it is advisable to use performance-contingent awards to enhance the profitability of open innovation contests.

A performance-contingent award, however, may not work for expertise-based projects. Recall that unlike ideation projects and trial-and-error projects, in an expertise-based project, solvers are differentiated in endowed expertise, and the design process is not influenced by random noise. A solver with relatively low endowed expertise still has a chance to win, but predicts that the award that he potentially could win is limited because it is proportional to his low endowed expertise plus his equilibrium effort. As a result, solvers with relatively low endowed expertise become more conservative in exerting effort compared to a fixed-price contest. Of course, solvers with relatively high endowed expertise become more aggressive compared to a fixed-price contest. Therefore, the power of a contest with a performance-contingent award for an expertise-based project depends on the distribution of solvers' endowed expertise $F(\beta)$.

6. Internal R&D vs. Open Innovation Systems

In this section, we contrast the profitability of internal R&D with two different types of innovation contests.

One is an innovation contest that is administrated by the seeker (as done by QVC or Staples), and another one is an innovation contest that is administrated by an independent intermediary (as done by InnoCentive). To make the comparisons tractable, we consider the case in which the seeker's goal is to maximize the performance of the best solution ($\rho = 1$), and the solver's fixed cost of participation is zero ($c_f = 0$).

We first examine the seeker's optimal internal R&D strategy for each type of project. Let c_{s1} and c_{s2} be the seeker's costs-of-improvement effort and experimentation effort, respectively. For an expertise-based project, let β_s be the seeker's endowed expertise on the project. If the seeker chooses to conduct the project internally, her problem can be written as

$$\max_{e \geq 0, m \geq 1} \Pi_I(e, m) = v(\beta_s, e, m, \xi) - c_{s1}e - c_{s2}m,$$

where the performance function $v(\beta_s, e, m, \xi)$ is defined in (1).

When administrating an open innovation contest, the seeker needs to develop a pool of participating solvers through advertising, invitations, and other means. There is a cost associated with such solver pool development activities. We assume that to develop a solver network with n solvers, the seeker incurs a cost of $c_s n$, which is linear in the size of the solver pool n . The seeker's problem for a self-administrated open innovation system is

$$\max_{n \geq 2} \Pi_S(n) = \Pi^*(n) - c_s n,$$

where $\Pi^*(n)$ is the seeker's optimal expected profit in a fixed-price open innovation contest with n solvers, which is defined in Theorem 1B. Solving the above problem for a logarithmic return function $r(e) = \theta \ln e$ and Gumbel-distributed solver expertise with scale parameter λ for expertise-based projects, the optimal size of the solver pool the seeker would develop for expertise-based, ideation, and trial-and-error projects are

$$n_S^{e*} = \frac{\lambda + \sqrt{\lambda^2 - 4c_s\theta}}{2c_s}, \tag{7}$$

$$n_S^{i*} \approx \frac{\mu - \theta}{c_s}, \tag{8}$$

and

$$n_S^{t*} = \sqrt{\frac{\mu}{c_s}}, \tag{9}$$

respectively. Note that the seeker's marginal value of obtaining an additional solution is decreasing in our model. For the seeker, each solution he receives corresponds to a draw of some distribution. Because of the properties of the extreme value distribution, the extra gain from this draw decreases.

Table 2 Conditions for Each Innovation Mechanism to Be Optimal for Each Project Type

	Internal R&D	Self-administrated open innovation	Intermediary-administrated open innovation
Expertise-based	$\theta < \theta_{SO}^e$ and $c_1 > f^e(\theta)$ or $\theta \geq \theta_{SO}^e$ and $c_1 > g^e(\theta)$	$\theta < \theta_{SO}^e$ and $c_1 \leq f^e(\theta)$	$\theta \geq \theta_{SO}^e$ and $c_1 \leq g^e(\theta)$
Ideation	$\theta < \theta_{SO}^i$ and $c_1 > f^i(\theta)$ or $\theta \geq \theta_{SO}^i$ and $c_1 > g^i(\theta)$	$\theta < \theta_{SO}^i$ and $c_1 \leq f^i(\theta)$	$\theta \geq \theta_{SO}^i$ and $c_1 \leq g^i(\theta)$
Trial-and-error	$\mu < \mu_{SO}^t$ and $c_2 > c_2^t$ or $\mu \geq \mu_{SO}^t$ and $c_2 > g^t(\mu)$	$\mu \geq \mu_{SO}^t$ and $c_2 \leq g^t(\mu)$	$\mu < \mu_{SO}^t$ and $c_2 \leq c_2^t$

If the seeker decides to use an open innovation contest with n_o solvers that is administrated by an independent intermediary, the seeker needs to pay a fixed fee, p , to use the intermediary’s service. For example, InnoCentive charges a fixed fee to allow a seeker to run a fixed number of contests on its website. In this case, the seeker’s optimal expected profits for the three types of projects are $\Pi_O^* = \Pi^*(n_o) - p$, where $\Pi^*(n_o)$ is defined in Theorem 1B for each type of project.

As discussed above, both innovation contest approaches (self-administrated and intermediary administered) suffer from a loss of efficiency caused by the underinvestment in the solver effort. However, relative to internal R&D, there exists a second form of inefficiency, which is the classical “double marginalization” effect in decentralized systems. With internal R&D, the seeker’s investment in R&D costs directly generates performance. In contrast, in an innovation contest, the seeker and the solvers form a decentralized system. The seeker’s investment in the award size cannot directly generate performance. Instead, it needs to induce the solvers to exert efforts that in turn generate performance. Thus, the attractiveness of open innovation relative to internal R&D depends on whether or not it can offer sufficient benefits to overcome the two efficiency losses. Theorem 3A characterizes the seeker’s optimal R&D mechanism.


THEOREM 3A. Let $r(e) = \theta \ln e$ be the solvers’ effort function and consider Gumbel-distributed solver expertise with scale parameter λ for expertise-based projects. Given

the seeker expertise, β_s , the effort cost, c_{s1} , the number of solvers at the intermediary, n_o , and the price charged by the intermediary, p , the conditions for each innovation mechanism to be optimal for each project type are provided in Table 2. The definitions for θ_{SO}^e , θ_{SO}^i , μ_{SO}^t , c_2^t , $f^e(\theta)$, $f^i(\theta)$, $g^e(\theta)$, $g^i(\theta)$, and $g^t(\mu)$ are provided in the online appendix.

The seeker’s optimal choices of R&D mechanism for different project types are illustrated in Figure 4. As indicated in Theorem 3A and shown in Figure 4, if the external solver’s effort cost is lower than a certain level, either self-administrated or intermediary-administrated open innovation is a better choice than internal R&D. With low effort costs, external solvers would exert higher effort in equilibrium, which will create sufficient benefits to offset the inefficiency of the open innovation system.

Theorem 3A and Figure 4 also reveal a counterintuitive difference between self-administrated open innovation system and intermediary-administrated open innovation system. For expertise-based projects and ideation projects, self-administrated contests are preferred to intermediary-administrated contests when the return on effort coefficient is *small*. In contrast, for trial-and-error projects, self-administrated contests are preferred to intermediary-administrated contests when the return on effort coefficient is *large*. As we discussed above, with a larger solver pool, each solver exerts less effort in equilibrium due to the negative externality among solvers. With a high return-on-effort coefficient, this negative effect of lower effort

Figure 4 The Seeker’s Optimal R&D Mechanism



Notes. β has a zero mean Gumbel distribution with $\lambda = 2$ for expertise-based projects, $\beta = 1$ for ideation and trial-and-error projects, $\mu = 2$ for ideation projects, $c_{1s} = 0.1$, $c_{2s} = 0.1$, $c_s = 0.05$, $\beta_s = 1$, and $n_o = 5$. θ is the return on improvement effort coefficient for expertise-based projects and ideation projects. μ is the scale parameter for Gumbel distributions for all projects.

is amplified for expertise-based projects and ideation projects. Therefore, it is optimal for the seeker to reduce the size of the self-administrated open innovation system. This is reflected in the optimal sizes of the solver pool for these two types of projects specified in Equations (7) and (8), which are both *decreasing* in the return on effort coefficient θ .

However, for trial-and-error projects, effort is perfectly cumulative across solvers. Although each solver exerts less effort in a larger open innovation system, the total amount of effort exerted by all solvers is actually higher (see detailed discussion in §4). With a high return-on-effort coefficient, there hence exists a stronger incentive for the seeker to develop a larger solver pool, which strengthens the attractiveness of a self-administrated contest. This is reflected in the optimal size of the solver pool specified in Equation (9), which is *increasing* in the return-on-effort coefficient μ .

In addition to providing access to a network of solvers, an intermediary-administrated contest can offer other benefits beyond what is captured in our model. This includes the fact that the identity of the seeker is kept private (in some cases, the seeker prefers the outside world not to know that he is working on a particular problem), the benefits of establishing a trustworthy third party that can broker the intellectual property rights between the seeker and the solvers, and the development of the required technical and organizational infrastructure.


THEOREM 3B. Consider a fixed-price intermediary-administrated open innovation contest with n_o solvers and a logarithmic return function $r(e) = \theta \ln e$. (i) For expertise-based projects, if the upper bound on the support of F , β , is high enough, there exists a solver pool size, \bar{n}^e , such that the seeker strictly prefers open innovation to internal R&D when $n_o \geq \bar{n}^e$. Otherwise, there exists a $\bar{\beta}_s$ such that the seeker strictly prefers internal R&D to open innovation when the solver's own expertise $\beta_s \geq \bar{\beta}_s$. (ii) For ideation projects, there exists \bar{n}^i such that the seeker strictly prefers open innovation to internal R&D when $n_o \geq \bar{n}^i$. (iii) For trial-and-error projects, if $c_{s2} > c_2$, there exists a solver pool size, \bar{n}^t , such that the seeker strictly prefers open

innovation to internal R&D when $n_o \geq \bar{n}^t$. Otherwise, the seeker always prefers internal R&D to open innovation.

For an expertise-based project, the benefit of an intermediary-administrated open innovation system is that there potentially could be a “genius” in the solver pool with a much higher endowed expertise than the seeker’s own endowed expertise β_s . For ideation projects, solvers are differentiated as captured by the random noise variable ξ in the performance function. An intermediary-administrated open innovation benefits the seeker by providing a higher diversity in solutions. For trial-and-error projects, an intermediary-administrated open innovation offers no additional benefits relative to internal R&D other than potential cost savings. Hence, without a cost advantage (i.e., $c_{s2} \leq c_2$), open innovation is never a better choice for the seeker. This conclusion seems surprising given that we have shown that free-entry open innovation is always optimal for a trial-and-error project due to its perfect effort cumulation. However, just because of this perfect effort cumulation, the seeker is able to completely replicate the benefit of multiple experiments internally.

Figure 5 summarizes the insights of the above discussions by illustrating how the seeker’s optimal innovation process (internal versus open) changes with respect to relevant parameters for each type of project. Consistent with intuition, as the seeker’s own costs of effort increase, the necessary number of solvers (i.e., \bar{n}^e , \bar{n}^i , and \bar{n}^t) that leads the seeker to choose an intermediary-administrated open innovation over internal R&D decreases for all project types. For both, expertise-based projects and ideation projects, the seeker prefers open innovation to internal R&D even when there exists a cost advantage over the external solvers (i.e., $c_{1s} < c_1$), whereas this can never happen for trial-and-error projects (when $c_{2s} \leq c_1$). Thus, we conclude that expertise-based projects and ideation projects are more suitable for an intermediary-administrated open innovation model than are trial-and-error projects.

Figure 5 The Seeker’s Optimal Choice Between Internal R&D and Intermediary-Administrated Open Innovation System with n_o Solvers



Note. β has a zero mean Gumbel distribution with $\lambda = 2$ for expertise-based projects, $\beta = 1$ for ideation and trial-and-error projects, $\mu = 2$, $\theta = 1$, $c_1 = 0.1$, $c_2 = 0.1$, and $\beta_s = 2$.

COROLLARY. Consider a logarithmic return function $r(e) = \theta \ln e$. Using a fixed-price open innovation contest, for expertise-based and trial-and-error projects, the seeker benefits from solutions with better performances rather than from lower costs. For ideation projects, the seeker benefits from both better solutions and from lower costs.

For the logarithmic return function $r(e) = \theta \ln e$, the seeker's total R&D costs are θ , $\theta + \mu$, and μ for expertise-based, ideation, and trial-and-error projects, respectively. The optimal awards the seeker offers in fixed-price open innovation contests for expertise-based, ideation, and trial-and-error projects are θ , μ , and μ , respectively. Regardless of whether R&D is carried out internally or externally, the seeker spends the same amount for expertise-based projects and trial-and-error projects. For ideation projects, the seeker spends less with open innovation compared to internal R&D. In an ideation project, because each solver must submit one prototype, the seeker obtains n_0 prototypes for free in the open innovation contest, whereas in internal R&D, the seeker incurs costs for each prototype built. The corollary shows that open innovation should not be viewed primarily as a mechanism to achieve cost savings. Instead, open innovation leads to better performance. This point should be kept in mind in the ongoing discussion about the impact of open innovation on our economy. Even if some R&D work is shifted to other regions because of the locations of the solvers, the local economy (and society) would still benefit from obtaining better products (which might also translate into more non-R&D jobs).

7. Discussion and Conclusion

The promise of open innovation is appealing: increase your capacity to innovate by tapping into a network of knowledge transcending organizational boundaries. However, as we have shown, not all innovation problems are suited equally well to this type of process. Unlike in the case of internal innovation, solvers participating in open innovation contests have to fear that their problem-solving effort might not be financially rewarded. This leads them to underinvest in effort, and an inefficiency in the market. The seeker (or, if applicable, the intermediate) organizing the innovation contest needs to be conscious of this effect and to design the reward system taking into account the type of the innovation problem (see Theorem 1A). In addition to choosing an appropriate reward, a key question for the open innovation system relates to the number of potential solvers. Although economists have argued that contests should be limited to two solvers to minimize the underinvestment effect while still benefiting from competition, we show that for an

innovation contest the benefits of diversity can outweigh or at least mitigate the negative effect of underinvestment. This can make large, fully open contests profit maximizing to the seeker (Theorem 1B).

To further increase the efficiency of innovation contests, mechanisms beyond the performance-contingent award can be conceived. For example, one could design a multiround contest, in which the first round is played with a large pool of contestants who make relatively little investment. This will identify skillful (and/or lucky) solvers who then could be allowed to play in a limited ("private") second-round contest. In this second round, the limited pool of solvers will drastically increase the probability of any participating solver winning the award and hence overcome some of the underinvestment problem. Furthermore, it is interesting to extend our one-dimensional innovation model into more complex innovation models such as innovation with unforeseeable uncertainty, whose R&D processes are often characterized as open ended search for the unknown unknowns (e.g., Sommer and Loch 2004). For an open ended search problem, open innovation contests could enable the seeker not only to improve performance along a known dimension, but also to see if there exist solutions/ideas of which the seeker is not even aware. Alternative contest mechanisms and innovation models are thus fruitful areas of future research.

Future empirical research could analyze how innovation contests are operated in practice as well as how (and if) they are replacing internal innovation and development processes. Based on our findings, we certainly expect a growing popularity of this form of innovation, with applications going well beyond the current focus on biology or chemistry. To take an example "close to home," consider the academic research process that leads to publications in a journal such as *Management Science*, and imagine how an author might benefit from relying on the help of an experienced solver when searching for the proof of a difficult theorem!

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Acknowledgments

The authors thank the department editor, Christoph Loch, the associate editor, and three anonymous reviewers for their helpful comments. Support for this work was provided to the second author through a summer research grant from the School of Business Administration at University of Miami.

References

- Anderson, E. G., A. Davis-Blake, S. S. Erzurumlu, N. R. Joglekar, G. G. Parker. 2006. The effects of outsourcing, offshoring, and distributed product development organization on coordinating the NPD process. C. Loch, S. Kavadias, eds. *Handbook of New Product Development Management*. Butterworth Heinemann-Elsevier, Burlington, MA, 259–290.
- Bulkeley, W. M. 2006. Got a better letter opener? Staples solicits inventive ideas from the public for products it can brand, sell exclusively. *Wall Street Journal* (July 13).
- Che, Y., I. Gale. 2003. Optimal design of research contests. *Amer. Econom. Rev.* **93**(3) 646–671.
- Chen, F., W. Xiao. 2005. Salesforce incentives and information screening: Individualistic vs. competitive schemes. Working paper, Columbia University, New York.
- Chesbrough, H. W. 2003. *Open Innovation: The New Imperative for Creating and Profiting from Technology*. Harvard Business School Press, Cambridge, MA.
- Dahan, E., H. Mendelson. 2001. An extreme-value model of concept testing. *Management Sci.* **47**(1) 102–116.
- Deng, S., W. Elmaghraby. 2005. Parallel sourcing with a tournament. *Productions Oper. Management* **14**(2) 252–276.
- Eppinger, S. D., A. R. Chitkara. 2006. The new practice of global product development. *Sloan Management Rev.* **47**(4) 22–30.
- Fullerton, R. L., R. P. McAfee. 1999. Auctioning entry into tournaments. *J. Political Econom.* **107**(3) 573–605.
- Glazer, A., R. Hassin. 1988. Optimal contests. *Econom. Inquiry* **26**(1) 133–143.
- Ha, A., E. L. Porteus. 1995. Optimal timing of reviews in concurrent design for manufacturability. *Management Sci.* **41**(9) 1431–1447.
- Hempel, J. 2006. Crowdsourcing—Milk the masses for inspiration. *Bus. Week* (September 25).
- Huston, L., N. Sakkab. 2006. Connect and develop: Inside Procter & Gamble's new model for innovation. *Harvard Bus. Rev.* **84**(3) 58–67.
- Kalra, A., M. Shi. 2001. Designing optimal sales contests: A theoretical perspective. *Marketing Sci.* **20**(2) 170–193.
- Krishnan, V., K. T. Ulrich. 2001. Product development decisions: A review of the literature. *Management Sci.* **47**(1) 1–21.
- Lakhani, K. R., L. B. Jeppesen, P. A. Lohse, J. A. Panetta. 2007. The value of openness in scientific problem solving. Working paper, Harvard Business School, Boston.
- Lazear, E. P., S. Rosen. 1981. Rank order tournaments as optimum labor contracts. *J. Political Econom.* **89**(5) 841–864.
- Loch, C. H., S. Kavadias. 2007. *Handbook of New Product Development Management*. Butterworth Heinemann-Elsevier, Burlington, MA.
- Loch, C. H., A. DeMeyer, M. T. Pich. 2006. *Managing the Unknown: A New Approach to Managing High Uncertainty and Risk in Projects*. John Wiley & Sons, Hoboken, NJ.
- Loch, C. H., B. A. Huberman, S. T. Stout. 2000. Status competition and performance in work groups. *J. Econom. Behav. Organ.* **43** 35–55.
- Loch, C. H., C. Terwiesch, S. Thomke. 2001. Parallel and sequential testing of design alternatives. *Management Sci.* **47**(5) 663–678.
- Mihm, J., C. Loch, A. Huchzermeyer. 2003. Problem-solving oscillations in complex engineering projects. *Management Sci.* **49**(6) 733–750.
- Moldovanu, B., A. Sela. 2001. The optimal allocation of prizes in contests. *Amer. Econom. Rev.* **91**(3) 542–558.
- Novak, S., S. D. Eppinger. 2001. Sourcing by design: Product complexity and the supply chain. *Management Sci.* **47**(1) 189–204.
- Snir, E., L. M. Hitt. 2003. The emerging knowledge economy: Exchange in Internet spot markets for IT services. *Management Sci.* **49**(11) 1504–1520.
- Sommer, S. C., C. Loch. 2004. Selectionism and learning in projects with complexity and unforeseeable uncertainty. *Management Sci.* **50**(10) 1334–1347.
- Taylor, C. R. 1995. Digging for golden carrots: An analysis of research tournaments. *Amer. Econom. Rev.* **85**(4) 872–890.
- Taylor, W. C. 2006. Here's an idea: Let everyone have ideas. *New York Times* (March 26).
- Terwiesch, C., C. H. Loch. 2004. Collaborative prototyping and the pricing of custom-designed products. *Management Sci.* **50**(2) 145–158.
- Terwiesch, C., K. T. Ulrich. 2008. *Innovation*. Forthcoming.
- Ulrich, K. T. 2008. *Design: The Creation of Artifacts in Society*. Forthcoming.
- Ulrich, K. T., D. J. Ellison. 1999. Holistic customer requirements and the design-select decision. *Management Sci.* **45**(5) 641–658.
- von Hippel, E. 2005. *Democratizing Innovation*. MIT Press, Cambridge, MA.
- Weitzman, M. L. 1979. Optimal search for the best alternative. *Econometrica* **47**(3) 641–654.
- Whang, S. 1992. Contracting for software development. *Management Sci.* **38**(3) 307–324.