MANAGEMENT SCIENCE
Vol. 58, No. 3, March 2012, pp. 606–623
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)
http://dx.doi.org/10.1287/mnsc.1110.1418

Managing Delegated Search Over Design Spaces

Sanjiv Erat, Vish Krishnan
Rady School of Management, University of California, San Diego, La Jolla, California 92093
{serat@ucsd.edu, vk@ucsd.edu}

Organizations increasingly seek solutions to their open-ended design problems by employing a contest approach in which search over a solution space is delegated to outside agents. We study this new class of problems, which are costly to specify, pose credibility issues for the focal firm, and require finely tuned awards for meeting the firm's needs. Through an analytical model, we examine the relationship between problem specification, award structure, and breadth of solution space searched by outside agents toward characterizing how a firm should effectively manage such open-ended design contests. Our results independently establish and offer a causal explanation for an interesting phenomenon observed in design contests—clustering of searchers in specific regions of the solution space. The analysis also yields a cautionary finding—although the breadth of search increases with number of searchers, the relationship is strongly sublinear (logarithmic). Finally, from the practical perspective of managing the delegated search process, our results offer rules of thumb on how many and what size awards should be offered, as well as the extent to which firms should undertake problem specification, contingent on the nature (open-endedness and uncertainty) of the design problem solution being delegated to outside agents.

Key words: research and development; open innovation; product design; clustering; search
History: Received March 8, 2010; accepted May 5, 2011, by Kamalini Ramdas, entrepreneurship and innovation. Published online in *Articles in Advance* September 20, 2011.

1. Introduction

Firms have traditionally made substantial investments to build internal research and development (R&D) capabilities. However, the emergence of internetworking technologies and connectivity with skilled professionals worldwide have made it feasible for firms to look beyond the enterprise boundaries in their search for innovations (Eppinger and Chitkara 2006, Anderson et al. 2008, Terwiesch and Ulrich 2009), with even R&D powerhouses, such as P&G and HP, beginning to explore the use of outside resources (Huston and Sakkab 2006). Although some of these "outside" contributions come from contracted suppliers, increasingly such collaborations utilize large pools of freelance problem solvers and employ mechanisms that are variations of the classical *tournament* approach, where the firm poses a problem and offers a reward to the outside problem solver submitting the "best" solution (Terwiesch and Xu 2008, Lakhani et al. 2007).

The tournament approach has to date largely been studied in the economics literature for well-specified problems with convergent solutions. The overwhelming emphasis in this stream, starting with Taylor (1995), has been on understanding how the reward incentivizes the outside problem solver's efforts and consequently determines the value derived from such tournaments. The recent past has seen this tournament approach to innovation applied even to open-ended and early-stage design problems in new product development (NPD). Our focus in this paper is on understanding the expansion of the tournament approach to contexts involving hard-to-specify design problems.

Consider the following two settings, which serve to contrast and illustrate the traditional and emerging environments in which the tournament/contest approach is being applied. The first is a well-specified problem—find a material to replace aluminum—in which the solution-seeking organization ("seeker") precisely spells out the requirements that new material must satisfy; namely, the new material should have heat conductivity >200 W/m K, be heat resistant up to 350 °C, etc. Based on these relatively unambiguous requirements, a group of outside solvers attempt to solve the problem, and the best entry at the time of the tournament deadline receives an award. In this example of a classical research/innovation tournament, the problem is well specified, and upon completion, the solvers know whether they have met the requirements.

Contrast this with the second example, the "Rwanda Rural Electrification Challenge," which illustrates the application of the tournament approach to more open-ended and ill-specified problems (Skipso 2009). It involves the Switzerland nongovernmental organization (NGO) Mabawa, looking to bring electricity to the rural village of Nyamyumba, Rwanda. The challenge aims to generate ideas for cost-efficient, reliable, stand-alone solutions that explore the use of multiple renewable energy sources. The top 10 winning solutions will each get a total cash award of $2,500. In this case, although some very basic conditions have been specified, it is very expensive for the idea-sourcing organization to specify every single requirement, consequently making it difficult for solvers to know if the solution they have fits the seeker's requirements.

The new class of problems to which delegated search is being applied have several unique features: First, the solution-seeking organization may need to spend considerable resources to fully specify the problem. Lack of (adequate) specification of the problem, on the other hand, may deter outside searcher participation. This is because of the additional risk created for searchers of not knowing how closely their search results match the requirements, and whether or not their effort is going to be rewarded.

Second, associated with the design problem, there are multiple solution approaches that form a design space. For example, the Rwanda Challenge specifically seeks the use of multiple renewable energy sources. The various combinations of renewable energy sources (solar, wind, biofuels, etc.) along with local conditions potentially represent solution approaches that could define the design space. There is a key feature of design space, and of search on design spaces, crucial for our development: The quality of a solution obtained by a searcher depends not (just) on the effort or ability of the searcher, but on where it resides on the design space. Indeed, this very fact makes it unlikely that two solutions that employ a common design approach (i.e., are situated on the same point of the design space) would have altogether different solution qualities. For instance, in the Rwanda example, the performance of two solutions, both of which employ solar as the underlying solution approach, would be similarly affected by the price of solar panels, the number of days of sun they get, the efficiency of the infrastructure to move energy, etc. Given such positive correlation (of solution qualities) among solutions that employ the same design approach, the breadth of coverage of the design space is desirable because the design approach that seems best at the outset may turn out to be disappointing; all the solutions using that approach may yield equally poor outcomes.

However, the outside problem solver's (henceforth, searcher's) choice of the design approach is not directly controlled by the firm, but is determined by the incentives (number and size of rewards associated with winning), and more broadly by the very structure of the design space itself. Hence, beyond the need to manage the individual searcher's efforts identified by past literature, employing the tournament approach on design spaces requires the focal firm (henceforth, principal) to manage the breadth of induced search. To illustrate this key idea, we offer a simple numerical example next.

1.1. Importance of Search Breadth and Difficulty Achieving it with Delegated Search

Consider a design problem that may be solved by two distinct design approaches, A and B. We view such design approaches as representing the design space. Assume that the quality of a solution that employs design approach A (B) is ex ante uncertain.

The heart of our model is in viewing this solution quality as being driven not just by idiosyncratic factors, but also by the design approach. We capture this as follows: Assume that there is an identical probability γ that a design approach is "feasible" (i.e., it yields good solutions). In addition, conditional on design approach A (B) being feasible, suppose that a searcher who employs the design approach A (B) obtains a solution quality of \$3 or \$0 (\$2 or \$0) with probability $\frac{2}{3}$ or $1-\frac{2}{3}$, respectively. Of course, if a design approach is not feasible, then the solution quality is always \$0.

Given the above specification, the (unconditional) probability that a searcher employing design approach A (B) obtains solution quality \$3 (\$2) is $\frac{2}{3}\gamma$. In the setting above, consider a principal that is interested in obtaining a *single best solution*. Suppose the principal can attempt only one solution; it should obviously pick design approach A, because the expected payoff from picking A is $\frac{2}{3}\gamma 3+\frac{1}{3}\gamma 0=2\gamma$, and the payoff from picking approach B is $\frac{2}{3}\gamma 2+\frac{1}{3}\gamma 0=\frac{4}{3}\gamma$.

Now suppose that the principal can attempt two solutions.[1] In this case, the question of which design approach (or approaches) it should use in each of its solutions gets more interesting. Specifically, if the principal sticks only to design approach A, then the expected value it obtains is given by

$$(1-\gamma)0+\gamma\left(\tfrac{2}{3}\tfrac{2}{3}3+\tfrac{2}{3}\left(1-\tfrac{2}{3}\right)3+\left(1-\tfrac{2}{3}\right)\tfrac{2}{3}3 + \left(1-\tfrac{2}{3}\right)\left(1-\tfrac{2}{3}\right)0\right)=\tfrac{8}{3}\gamma,$$

where the first term, $(1-\gamma)0$, represents the event that the design approach A is infeasible; the second term, $\gamma\frac{2}{3}\frac{2}{3}3$, represents the event that design approach is feasible, and the additional (idiosyncratic) uncertainty for both solutions is resolved favorably; the third term, $\gamma\frac{2}{3}(1-\frac{2}{3})3$, represents the event that design approach is feasible, and the additional (idiosyncratic)

[1] The choice of whether to attempt one or two solutions depends to a large degree on the costs (Dahan and Mendelson 2001, Loch et al. 2001), an issue that is unnecessary to address for the purpose of this illustrative example.

uncertainty for the first solution is resolved favorably and that for the second solution is resolved unfavorably, etc. Similarly, if the principal sticks only to design approach B, the expected value is given by

$$(1-\gamma)0+\gamma\big(\tfrac{2}{3}\tfrac{2}{3}2+\tfrac{2}{3}\big(1-\tfrac{2}{3}\big)2+\big(1-\tfrac{2}{3}\big)\tfrac{2}{3}2 + \big(1-\tfrac{2}{3}\big)\big(1-\tfrac{2}{3}\big)0\big)=\tfrac{16}{9}\gamma.$$

If the principal attempts one solution from each design approach, then its expected solution quality is

$$\big(\tfrac{2}{3}\gamma\big)\big(\tfrac{2}{3}\gamma\big)3+\big(\tfrac{2}{3}\gamma\big)\big(1-\tfrac{2}{3}\gamma\big)3+\big(1-\tfrac{2}{3}\gamma\big)\big(\tfrac{2}{3}\gamma\big)2 + \big(1-\tfrac{2}{3}\gamma\big)\big(1-\tfrac{2}{3}\gamma\big)0=\tfrac{2}{9}\gamma(15-4\gamma).$$

Hence, the principal should attempt two solutions that employ different design approaches if and only if (iff) $\frac{2}{9}\gamma(15-4\gamma)\geq\frac{8}{3}\gamma$, i.e., iff $\gamma\leq\frac{3}{4}$. Otherwise, the principal benefits from attempting two solutions, both using approach A.

The preceding analysis offers the basic underlying reason for why maintaining breadth of search on the design space is important for the principal; namely, when design approaches are expected to play an important systematic role in determining solution quality (which they do *not* when γ is large because they are very unlikely to be "infeasible"), then the principal benefits from attempting multiple design approaches and thus searching more broadly. Now consider what occurs when the principal offers a single reward for the best solution and delegates this search process to outside searchers. Suppose there are two searchers, each of who can submit a single solution. Which design approach would these delegated searchers pursue?

Suppose that ties in solution qualities are broken randomly. Because the searchers win a fixed prize, we only need to look at the ex ante probability of winning for each searcher to understand the induced equilibrium search patterns in this example of delegated search.

Specifically, if both searchers pursue the same design approach A or B, then, by symmmetry, they both have identical probability ($\frac{1}{2}$) of winning the prize. Now consider the case where one searcher pursues design approach A while the other pursues design approach B. In this scenario, the searcher pursuing the design approach A obtains the reward in two cases: (i) when he is successful (because his quality would be 3, compared to any quality achievable by the searcher on approach B) or (ii) when both fail, and the tie is broken in his favor; that is, the searcher who employs approach A has probability of winning the award given by

$$\tfrac{2}{3}\gamma+\tfrac{1}{2}\big(1-\tfrac{2}{3}\gamma\big)\big(1-\tfrac{2}{3}\gamma\big)=\tfrac{1}{2}+\tfrac{2}{9}\gamma^2.$$

Table 1 **Likelihood of Winning for the Searchers Contingent on Their Design Approach Choice**

	A	B
A	$\frac{1}{2}$, $\frac{1}{2}$	$\frac{1}{2}+\frac{2}{9}\gamma^2$, $\frac{1}{2}-\frac{2}{9}\gamma^2$
B	$\frac{1}{2}-\frac{2}{9}\gamma^2$, $\frac{1}{2}+\frac{2}{9}\gamma^2$	$\frac{1}{2}$, $\frac{1}{2}$

Hence, the other searcher's probability of winning the award is given by

$$1-\big(\tfrac{1}{2}+\tfrac{2}{9}\gamma^2\big)=\tfrac{1}{2}-\tfrac{2}{9}\gamma^2.$$

Table 1 summarizes the probabilities of winning the award for the two searchers contingent on their choice of design approach. As may be immediately inferred from the table, at (Nash) equilibrium, both searchers will choose design approach A.

The above example illustrates the core issue we deal with in this paper: Delegated search (may) result in decreased search breadth, an interesting and novel inefficiency that has so far not been identified in past literature.[2]

We examine the following related research questions in this paper:

1. What are the drivers of outside agent's search dynamics? How do the awards, problem specification, and structure of the design space influence the number of searchers and breadth of induced search?

2. How should the principal optimally manage the problem specification, and how should the contest awards be structured (i.e., how should the number and size of prizes be decided)?

The first question focuses on quantifying and offering a descriptive characterization of the extent to which search breadth reduces through delegation, and on understanding its primary determinants. In the second part of this paper, the principal's decisions about the extent to which the problem/requirements are to be specified and the number and size of awards are examined in light of their role in determining the breadth of induced search and the ultimate value from the delegated search mechanism. A richer understanding of these decisions promises to expand the scope of delegated search to a much broader class of design problems that goes beyond the classical research/innovation tournaments (Taylor 1995, Che and Gale 2003, Terwiesch and Xu 2008).

In answering these questions, this paper makes several contributions: Our proposed analytical framework explicitly accounts for the unique aspects of

[2] The vast majority of past literature examines a different type of inefficiency, namely, effort distortion in research tournaments.

design problems such as (i) the extent (or completeness) of problem specification and the potential mismatch between searcher- and firm-perceived solution qualities, and (ii) the searchers' choice of design approach (i.e., point in design space to search). It also allows us to parsimoniously capture the search dynamics induced by the awards and problem type, to quantify the above-mentioned novel inefficiency introduced through delegation of the search process, and to characterize its determinants as well as offer normative prescriptions on how best to manage delegated search over design spaces.

Our analysis demonstrates that the key result suggested by the previous numerical example is quite general. Specifically, we find that searchers tend to *cluster* on design approaches that demonstrate higher ex ante potential. Such clustering is likely to occur, irrespective of problem type and whether or not the firm offers single or multiple prizes (as long as there is one largest prize). The critical determinant of clustering is the unequal size of awards and competition among searchers for larger ones.

We reach a cautionary result—although breadth of search increases with number of searchers, the relationship is sublinear in nature (logarithmic or lesser in our model). Thus, looking at only the number of searchers to determine the value of delegated search is likely to significantly overstate the benefit of this type of R&D arrangement. In a similar vein, we find that the relationship between induced search breadth and the award amount is sublinear in nature (logarithmic or lesser in our model); that is, attempting to increase the induced breadth by increasing award amount may be an expensive proposition.

Interestingly, we find that a firm may prefer to not completely specify the problem and instead leave some of the specifications (requirements) ambiguous, even when it is *costless* to specify the problem completely. Such *strategic ambiguity* in the specifications may be desirable because of an interesting trade-off between the extent of specification and the induced search breadth: problems that are poorly specified, and consequently have lower likelihoods of the searcher finding an acceptable solution, will induce fewer searchers to participate. This, in turn, reduces the breadth of search over the solution space. However, well-specified problems increase the searcher's likelihood of finding an acceptable solution and make the participating searcher unwilling to take the added risk of straying outside ex ante promising regions of the solution space. Thus, the coverage of the solution space is greatest when the problem is moderately specified (i.e., neither too incomplete nor too complete).

We begin with a review of the related literature in the next section. In §3 we present a model of delegated search where the principal broadcasts the problem to a population of searchers (also called agents). We solve this model in §4 to characterize the key effects of the award structure (number and size of awards) and problem specification on the number of searchers and the breadth of the induced search. In §5 we derive the optimal award structures and problem specification for the delegated-search mode of problem solving. Last, in §6 we conclude with a discussion of key insights and questions for future research.

2. Review of Relevant Literature

The literature relevant to this paper on innovation processes may be organized into two groups based on whether it views R&D as a *search process* or a repeated *sampling process*.

2.1. R&D as a Search Process

Problem solving, an essential part of R&D, has often been described as "a search through a vast maze of possibilities" (Simon 1969, p. 54). Such a conceptualization emphasizes the probabilistic nature of innovation and captures the reality that the outcome of any innovation process is ex ante uncertain. Levinthal (1997), using a model of NK landscapes developed by Kauffman and Levin (1987), examines how firms adapt/search on strategy landscapes for superior strategies.

Researchers in NPD have extended models of search and problem solving to consider a firm's prototyping and concept development processes. Loch et al. (2001) consider firms that examine multiple alternative designs with the objective of choosing the best one. Erat and Kavadias (2008) examine the implications of learning during sequential testing and find that the ability to learn across designs make it optimal to have an exploratory phase at the beginning of a test cycle. Terwiesch and Loch (2004) examine a firm that conducts searches for others and characterize the effect of problem structure on how the searching firm may set its fees. Sommer and Loch (2004) contrast two different search strategies in innovation—learning or trial-and-error approaches and selectionism or parallel approaches—in the quest for superior performance on rugged landscapes. Mihm et al. (2003) view a project as interdependent component subsystems and thus possessing a performance landscape that is necessarily rugged. Our goal, in contrast to these prior studies, is to understand the induced search dynamics when a firm (principal) delegates the search activity to *multiple searchers* (agents) by offering to reward the best solution using a contest approach. Toubia (2006) considers the question of incentivizing idea generation in groups and identifies interesting information-related positive spillovers in such groups. He finds that because this positive spillover is not internalized by the participating agents, it may result in agents pursuing information-generating exploration strategies to a lesser degree than would be socially optimal.

Two closely related papers to this work empirically examine creative problem solving by outside solvers. Girotra et al. (2010) compare the effectiveness of team process and hybrid approaches for coming up with the best ideas. They find that team processes result in more "similar" ideas, which in some cases may in fact result in poorer ideas overall. Kornish and Ulrich (2011) offered recent evidence that searchers tend to cluster around certain regions of the solution space. Our work complements these papers by offering an analytical framework for understanding these prior findings, especially the clustering of agents in regions of search spaces.

2.2. R&D as a Sampling Process

Models of search over design spaces possess some of the features of classic multiarm bandit problems and are known to be notoriously hard (Auer et al. 2002). Such analytical difficulties have arguably resulted in simplifying the search model of R&D. Specifically, this stream of literature proposes that R&D may be (approximately) viewed as a repeated sampling process where each *independent* sample represents the uncertain outcome of some R&D effort.

Weitzman (1979), in a seminal article, examined sequential sampling of independent alternatives with the goal of choosing the highest outcome. He characterized the optimal policy as one where sampling is terminated only upon obtaining a performance that is higher than a reservation price. Dahan and Mendelson (2001) considered a static parallel sampling model, where a firm may choose the number of alternatives to examine, and then characterized the dependence of the optimal number of alternatives on costs.

This approach has been extended to examine contractually delegated R&D (see, for instance, Weitzman 1980). Similarly, the literature on research tournaments, typified by what (Taylor 1995, p. 872) termed "digging for golden carrots," has examined settings where a firm, by offering an award for the highest solution, delegates the sampling process to multiple independent outside agents. The delegated sampling framework has since been extended to a number of situations, such as an agent's effort being a choice variable, the possibility of multiple awards (Moldovanu and Sela 2001, Che and Gale 2003), heterogeneous capabilities among agents (Terwiesch and Xu 2008), etc.

Yet, these studies, to (ab)use the "golden carrot" metaphor, have looked primarily at how hard each agent digs for the golden carrot, rather than where they dig! We argue that, with consideration of design contests and the associated solution spaces, the question of "search versus sampling," and more specifically where each agent searches becomes crucial. For instance, if two agents were to examine (search) the same point on the solution space, it seems highly unlikely that their solutions would be independent. Indeed, if anything, we should expect that such points are likely to yield solution qualities that are very similar. Thus, these "delegated sampling" models of research tournaments, which abstract away the possibility that searchers may submit closely related (or even redundant) solutions, seem more applicable for problems where the idiosyncratic searcher's capabilities drive the solution quality. Our work thus complements these past studies and is applicable to contexts where the characteristics of the design space, rather than the searcher's characteristics, drive solution quality.

The current paper differs from past work on innovation tournaments in two important ways. First, we break new ground in formalizing and analytically characterizing how delegated search applies to design contests. Specifically, the search perspective that we adopt allow us characterize the distribution and breadth of the search contingent on the problem type and the award structure. Second, we consider hard-to-specify problems by explicitly accounting for the operational reality that the person finding the solution (the searcher) and the person judging the quality of the solution (the principal) may not always perceive the same solution quality. This naturally leads us to consider the relationship between problem specification, the breadth of induced search, and the resultant value that the principal may obtain through delegated search.

3. Delegated Search: Model

Consider a solution seeker (principal) that wishes to delegate the search of a design problem to a group of outside agents (searchers). The problem in question is assumed to admit multiple design approaches. We conceptualize a design approach (corresponding to a design problem) as constituting the set of implicit and explicit design assumptions, where design approach corresponds to a single point in the design/solution space. Suppose that there are M feasible design approaches for the design problem posed by the principal, and suppose that each searcher may choose which (if any) of these approaches he wishes to pursue. We assume that a searcher may pursue no more than one design approach. This assumption assures mathematical tractability and seems reasonable in situations where either organizational inertia or the time/cost requirements may prevent a searcher from switching between different design approaches. Thus, a searcher who decides to participate in the delegated search chooses a point in the design space to probe.[3]

We assume that when a searcher pursues the design approach i ($\in \{1, \ldots, M\}$), he obtains a solution with *perceived* quality $X_i \in \{0, v_i\}$ with probabilities $1-p_i$ and p_i, respectively. Note that X_i represents the searcher's individual assessment of the solution quality. We assume that both the (potential) searchers and the principal are aware of this design space and have identical beliefs/knowledge (p_i and v_i) about it. This assumption, although stylized, allows us to obtain base-line results on search dynamics. Furthermore, although search dynamics may change with asymmetric information (about solution qualities for instance), such effects reflect information asymmetry rather than anything unique about search over design spaces.

We shall make two simplifying assumptions in our modeling of the solution space: (i) the probability p_i of obtaining the higher quality solution v_i is identical for every design approach, i.e., $p_i = p$, and (ii) the parameter v_i representing the attractiveness of the ith design approach satisfies the functional form $v_i = v\phi^{i-1}$, where $\phi \in [0, 1]$. Thus, the design approaches i ($\in \{1, 2, \ldots, M\}$) are ordered in the decreasing order of attractiveness. The preceding two assumptions, although seemingly limiting, are made *only* for ease of exposition and to focus the narrative on how the award and problem structure drive the dynamics of delegated search; we demonstrate in Online Appendix B (available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1567599) that our key results are robust even in situations where the design approaches have arbitrary attractiveness v_i and arbitrary probability p_i.

The parameter ϕ captures how "skewed" the solution space is; i.e., with higher ϕ, all the design approaches look ex ante very similar, whereas with low ϕ, few of the design approaches ex ante dominate. Past literature has used the term "convergent problems" to represent problems that possess "single accepted solutions," and the term "divergent problems" to represent those that "[do not] have a single solution" (Cameron 1986, p. 548). Because the parameter ϕ captures this very feature, we follow this terminology and call the problems with low ϕ, where a few design approaches dominate, convergent problems, and those where almost all solutions look ex ante acceptable (i.e., high ϕ) divergent problems. (We demonstrate in Online Appendix C, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1567599, that our key results are robust even to alternate parameterization of convergent and divergent problems.)

We posit that a key source of uncertainty in delegated search (and arguably in delegated problem solving in general) is the ambiguity that the searcher has of the principal's actual requirements and of its evaluation process. Thus, our model of delegated search needs to explicitly account for the reality that the principal, which has to finally evaluate the solution, may perceive the solution quality differently from the searcher perceived quality X_i. We capture this difference in perceived solution quality (between the searcher and the principal) by assuming that even if the searcher believes that an offered solution i has performance $X_i = v_i$, there is still a likelihood that the principal evaluated solution quality P_i to be 0. Let such a mismatch occur with probability $1-\lambda$, i.e., $1-\lambda = \Pr(P_i = 0 \mid X_i = v_i) = 1 - \Pr(P_i = v_i \mid X_i = v_i)$.

The extent of indefiniteness in the target laid out by the principal and the (lack of) completeness of its problem specification determine whether the searcher is able to validate the solution quality accurately (Minsky 1961, Smith 1988). For instance, if the principal has specified only the key performance metric of a solution and has left out some other secondary (but important) requirement, it is very likely that there will be a mismatch in the searcher- and principal-perceived solution qualities. Such incomplete problem specification would, thus, make the real solution quality (i.e., solution quality as perceived by the principal) more uncertain from the searcher's perspective. On the flip side, when the principal provides more complete problem specification (i.e., there is less goal indefiniteness), the likelihood of mismatch in the perceived quality is potentially lower. Hence, we interpret the parameter λ as the extent of completeness of specification.

Although the incomplete problem specification and the consequent mismatch between the searcher and principal creates a type of uncertainty that is unique to delegated search, past research in product development and innovation has argued that *any* evaluation of solution/design quality will be prone to error (i.e., is uncertain). Indeed, NPD research has found that it is often difficult, if not impossible, to assess the true quality of any solution upfront, and that the

[3] We assumed that the design approach corresponds to a particular set of design assumptions and maps onto a point in the design space. However, one may also view a design approach as corresponding to specific "search strategy" that is a point on a "meta-landscape of search strategies." Note that the latter interpretation is also consistent with agents whose chosen search strategy dictates examining multiple designs. However, to keep the model and interpretation simple without losing the qualitative insights, we shall consistently employ the former interpretation and view the point chosen by a searcher as a design approach.

full extent of uncertainty in quality is revealed only after market launch. We explicitly account for this by assuming that the principal, after its evaluation, either (a) finds the solution to be mismatched, and thus of solution quality 0, or (b) finds no mismatch (i.e., evaluated quality is v_i), and hence believes that the true solution quality is distributed uniformly between $v_i - \delta$ and $v_i + \delta$. Thus, whereas λ captures the uncertainty due to mismatch that is resolved immediately after the principal evaluates the solution, δ captures the more traditional form of (market) uncertainty considered in prior NPD literature, namely, the extent of uncertainty that principal has even after evaluating a searcher-submitted solution.

The model so far has addressed how a given searcher's solution quality is determined. Our key contention, both in the numerical example offered in §1 and in the literature review in §2, was that given the design space, if two searchers probe the same point on the design space (i.e., undertake the same approach), the performance of their solutions would likely be correlated by a large degree. For ease of exposition, we shall assume in the base model that this correlation is perfect, i.e., the performance of a solution is related *only* to the solution space and not to any idiosyncratic searcher-specific characteristic. More specifically, we shall also assume that the underlying technical uncertainty p, the (specification-related) mismatch uncertainty λ, and the market uncertainty δ are all determined by the solution approach. Thus, any two searchers who are probing the same point on the solution space obtain solutions of identical quality. This assumption shall be subsequently relaxed, and we shall demonstrate that our results hold even when this correlation is not perfect.

We assume in our base model that the principal announces a single fixed prize V to be awarded to the searcher who obtains the best design solution. This assumption of a single prize allows easier comparison to past literature (Taylor 1995) and, more importantly, is consistent with the overwhelming majority of delegated searches in practice. We consider the drivers for and the effect of multiple prizes in §5.1.

Past literature, to our knowledge, has assumed that the principal always has the ability to *unconditionally commit* to awarding the prize to one of the searchers. Still, it must be noted that such an assumption implicitly requires the presence of some commitment device, such as reputation, repeated play, or the presence of trusted middlemen, etc., issues that are unmodelled in past literature. We offer the following example, which suggests that such strong commitment may not always be a reasonable assumption: Intuit, a firm which sells tax preparation software, organized a delegated search (for a particular business problem) among MBA students at a U.S. business school. Although Intuit had promised a sizeable prize, they decided ex post not to accept any of the submitted solutions and not to award the prize.

Consistent with such examples, we explicitly consider a principal that cannot credibly commit to *unconditionally* giving out the prize V; specifically, if the principal's value from acquiring the searcher's solution is less than the prize it announced before, then, ex post, it is not rational for the principal to accept the solution.[4] Thus, the principal, after evaluating the searchers' solutions, may either award the prize V that was announced or not award any prize at all.

We do however assume that the principal or the agent may not renegotiate ex post the size of the prize. The assumption is necessary because without even this limited credibility, agents would have no incentive to pursue search. Similarly, we assume that it is infeasible for the principal to "steal" or "copy" the solution when the searcher offers the solution for evaluation. Both of these assumptions are ubiquitous in the literature on innovation contests (Taylor 1995, Terwiesch and Xu 2008) and although seemingly restrictive, are in fact aligned with practice where (i) the size of the prize, V, is ex ante contractible (however, what is not contractible is whether or not the prize will be awarded at all, because awarding the prize depends on whether the solution is acceptable to the principal), and (ii) the principal cannot copy/steal the solution without the searcher's permission, either because of intellectual property (IP) protections[5] or perhaps because the principal needs, in addition to the solution, some tacit knowledge that resides with the searcher.

After the principal has announced the design problem and associated award and broadcast them to N searchers, each of the searchers decides whether or

[4] In addition, consider the following: With the ability to commit unconditionally to award one prize, at least one searcher will always submit a solution. However, the reality that many "contests" do not attract any searchers shows that the market cannot possibly believe that the award would be given out unconditionally (i.e., irrespective of the actual solution quality). Note that the preceding argument appeals to pure strategy Nash equilibria. Specifically, it is indeed possible that with mixed strategies, even with the ability to commit unconditionally, situations where no one enters can arise. Although mixed strategies provide an alternate explanation, we view such an argument as relatively weak for two reasons: (a) the argument relies on a technical modeling construct, namely, mixed strategies, to rationalize a common phenomenon and thus fails to provide any intuition as to why it occurs, and (b) the argument is relatively fragile in that in any mixed strategy equilibria, the probability that no one enters cannot be very large (specifically, one may show that this probability must be less than c/V).

[5] In many such delegated search settings, the IP for the solution resides with the searchers and is transferred to the principal only for the winning solution (see, for instance, the Innocentive website: http://www.innocentive.com/about/forscientists.html).

Figure 1 Sequence of Decisions, Timing of Uncertainty Resolution, and an Illustration of the Solution Space

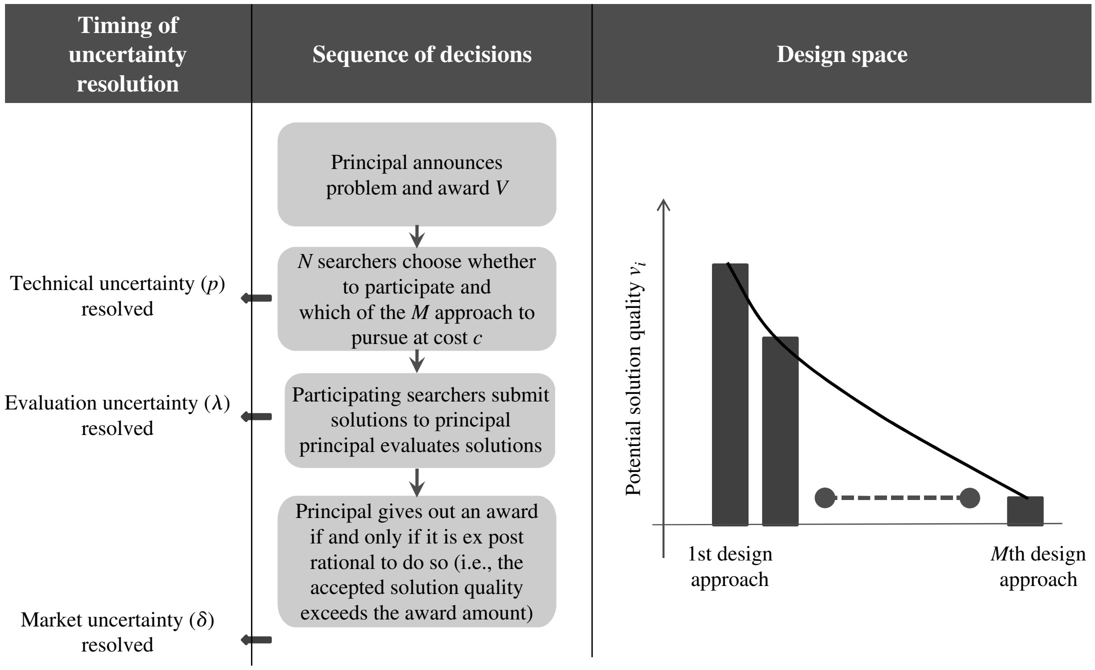

not to participate in the delegated search. If they do enter, they then choose which of the M design approaches to pursue. A searcher who decides to enter the contest, after searching their chosen point on the solution space, say j, will obtain a solution of performance X_j. Suppose that any searcher pursuing a design approach incurs identical cost c. This assumption of equal and identical cost across searchers and design approaches allow us to focus on how the award and problem structure drive the dynamics of delegated search. Needless to say, it is likely that unequal costs may make any design approach more or less attractive and may influence the breadth of search.

Finally, consistent with much of the past literature in innovation tournaments and R&D contests (Taylor 1995, Terwiesch and Ulrich 2009), we assume that the principal can profitably use only one solution; i.e., its payoffs are given by the highest quality solution among the pool of solutions submitted by the searchers.

Figure 1 gives the sequence of decisions in our model of delegated search. In the first stage, the principal sets the fixed reward V and announces them to the searcher population. In the second stage, each of the N searchers simultaneously decide on whether or not they will participate and the solution approach i ($\in \{1, \dots, M\}$) they would undertake to solve the problem. Each searcher who participates incurs a cost c in pursuing their chosen design approach. After the solutions are found, the searchers submit their individual solutions to the principal for evaluation. Finally, the principal evaluates the solutions and, contingent on its ex post rationality constraints, awards the prize V to the searcher who submitted the highest performing solution. Last, we assume that in case of ties between searchers, the principal is equally likely to choose any one of the highest performing solutions.[6]

4. Delegated Search: The Descriptive Characterization

In this section, we characterize the search pattern induced on the solution space and its dependence on the problem type and principal's announced award. At equilibrium, let m be the total number of design approaches attempted by the searchers (i.e., the breadth of search of the solution space), and let n_i be the number of searchers utilizing the design approach i. Lemma 1 completely characterizes the key features of (all[7]) the Nash equilibrium outcomes of

[6] Note that this assumption is identical to assuming that there are secondary metrics that the principal would rely on if solutions have identical qualities, and that these secondary metrics are unknown/uncertain to the searchers. Furthermore, given our assumption of risk-neutral searchers, from a technical perspective, the assumption that ties are broken randomly is equivalent to the assumption that the prize gets split equally in case of ties.

[7] Because we assume symmetric agents making their participation and design approach choices simultaneously, there may in general be multiple pure Nash equilibria. However, these equilibria will be

the "design approach choice" game played by the searchers.

Lemma 1. *For a given V, the number of design approaches attempted by the searchers is*

$$m = \min\left\{\left\lfloor \frac{\ln(V)-\ln(v)}{-\ln(\phi)}\right\rfloor + 1, \left\lfloor \frac{\ln(c)-\ln(V)-\ln(\lambda p)}{\ln(1-\lambda p)} + 1\right\rfloor\right\}, \quad (1)$$

and the number of searchers attempting approach design approach i is

$$n_i = \begin{cases} \left\lfloor \dfrac{V\lambda p(1-\lambda p)^{i-1}}{c} \right\rfloor & \text{for } i = 1, \ldots, m, \\ 0 & \text{for } i = m+1, \ldots, M. \end{cases} \quad (2)$$

The lemma illustrates that the seemingly complex "design approach choice" game that searchers play over the design space has a relatively simple equilibrium structure. Although the proof is somewhat tedious (given in Online Appendix A, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1567599), the intuition is relative simple: the number of searchers employing each design approach is such that the next searcher to pursue the design approach obtains (on expectation) no net value. A number of insights about the structure of the equilibrium may be inferred from the preceding lemma. We summarize these in following proposition.[8]

Proposition 1. (i) *More searchers will pursue design approaches with greater potential; i.e., $\{i < j\} \Rightarrow \{n_i \geq n_j\}$. The breadth of search m is (conditionally) greater than* 1.

(ii) *The breadth of search m and the number of searchers per design approach n_i are monotone decreasing in searcher cost.*

(iii) *The breadth of search m and the number of searchers per design approach n_i increase and then decrease in the award size V.*

Claim (i) states that, at equilibrium, more searchers congregate on more promising design approaches, i.e., searchers *cluster* on design approaches that they believe have greater potential. Intuitively, all else being equal, searchers should attempt the approach with greater potential because a solution from such an approach is more likely to be rewarded by the principal. The result is consistent with prior findings in ideation literature (Kornish and Ulrich 2011)[9] and with past findings in sales contests where rewarding only highest solutions/performances is found to result in agents focusing only on the approach that offers highest upside potential (Gaba and Kalra 1999). Still, unlike in a sales contest, with delegated search over solution spaces, not everyone pursues design approach 1. Why does this occur? The underlying economic reason may be traced to the fact that, in our model, two searchers undertaking the same design approach also have a higher likelihood of ending up with similar solution qualities. Thus, a searcher may find it better to follow a less "crowded" design approach with lower potential on the chance that everyone else following the more crowded, ex ante higher potential design approach ends up, ex post, with lower solution quality.

Thus, Claim (i) illustrates that there are two opposing forces at work in delegated search: agents find it beneficial to (a) pursue higher potential approaches, because such solutions, if successful, are more likely to be rewarded, and (b) move away from crowded design approaches, because success with these approaches implies others pursuing the same the crowded approach succeed as well, reducing the likelihood of obtaining the reward.

Claim (ii) finds that as the searcher costs become larger, fewer searchers participate in the search for a solution to the design problem. From a managerial perspective, innovations that require large upfront investments may not be suitable candidates for delegated search because such problems attract a relatively lower number of searchers (Lakhani et al. 2007). A second effect arising from an agent's larger upfront investments is that only a smaller part of the solution space will be searched. Thus, our result suggests that with higher searcher costs, fewer searchers participate, and the participating searchers cluster on the small number of design approaches that ex ante show greatest potential.

Claim (iii) finds a curious nonmonotone effect of award size on the breadth of search. As the award size increases from low to moderate values, approaches with lower potential become more attractive to the searchers. Thus, more searchers participate, and even the low potential approaches are attempted. This results in the breadth of the search increasing. However, when the award size becomes very large, all

merely permutations, and the two key features of the equilibria, namely, the number of design approaches (breadth) and the number of searchers per design approach, are identical across all these permutation equilibria and serve to characterize them completely.

[8] Note that the number of design approaches pursued and the number of searchers are both discrete variables. For ease of exposition, we abuse the terminology and use the terms "increasing" and "decreasing" instead of the technically correct terms "nondecreasing" and "nonincreasing," respectively.

[9] However, note that unlike the experimental finding of clustering in prior literature, our results are not driven by the possibility of shared knowledge (Kornish and Ulrich 2011), but emerge because of the searcher's strategic profit-maximizing behavior.

those approaches that yield performances lower than the award are avoided by the searchers. This occurs because the principal will find it rational, ex post, not to accept these solutions (because accepting these solutions yields lower profits to the principal than the profits from not giving out the award). Hence, the search breadth decreases as the award size increases from moderate to very high values. The nonmonotone effect of award size on the breadth of search has an important practical implication, summarized in the next observation.

Observation 1. There exists a threshold i^* such that design approaches $i > i^*$ can never be explored through delegated search irrespective of the award size.

The nonmonotone effect of award size on the search breadth implies that some parts of the solution space are "unreachable" and that there is a limit to which the breadth can be influenced by the award size alone. This inability to infinitely increase coverage using increasing awards stands in contrast to the extant literature on innovation contests (Taylor 1995, Terwiesch and Xu 2008) and stems directly from relaxing the common assumption of principal's credibility. We hasten to add that the specific claim stems from a stylized assumption, namely, that the principal cannot unconditionally commit to award a prize. And to the extent that the assumption is not satisfied, perhaps because of reputational concerns in repeated games, or because the principal uses outside institutions (such as third-party escrows), the principal will indeed be able to expand the breadth of the search space indefinitely by increasing the award amount.[10]

Claim (iii) of Proposition 1 showed that as the size of the award increases (but remains below $v_M = v\phi^{M-1}$), more design approaches will be pursued at equilibrium.[11] Specifically, it may be directly verified from Equation (1) (in Lemma 1) that the (minimum) award that the principal would have to set to induce the searchers to pursue m design approaches is given by

$$V(m) = \frac{c}{\lambda p(1-\lambda p)^{m-1}}. \tag{3}$$

The following corollary summarizes an important structural characteristic of this award function.

Corollary 1. *The award amount is exponentially increasing in the desired breadth of search.*

Thus, to pursue m design approaches in equilibrium, the principal has to set an award that is exponential in m. This finding, in addition to its theoretical importance, highlights an important practical implication: although delegated search may allow the principal to pursue multiple alternatives in parallel, the cost of conducting the search, i.e., the award size, must increase exponentially in number of approaches pursued.

Similar to Equation (3), the (minimum) award that the principal needs to set to induce n searchers to participate may be easily deduced from Equation (2) as

$$\begin{aligned} V(n) &= \arg\min_V \left\{ \sum_i n_i = n \right\} \\ &= \arg\min_V \left\{ \sum_i \left\lfloor \frac{V\lambda p(1-\lambda p)^{i-1}}{c} \right\rfloor = n \right\}. \end{aligned} \tag{4}$$

The following bounds for $V(n)$ allows us to characterize the structure of this award function:

Corollary 2. $nc \le V(n) \le 2nc$.

Thus, to have n searchers participate, the principal needs to offer an award that is linearly bounded in n. Although, our results in Corollary 1 and 2 about the exponential and linear nature of the relationships are doubtless driven by our specific assumption, comparison between the results yields a significantly more valuable insight: increasing the award amount has a disproportionately larger impact on the number of searchers than on the number of design approaches being searched.

Combining the preceding two corollaries, we obtain the next corollary, which demonstrates the relationship between number of searchers who participate and the search breadth.

Corollary 3. *The number of participating searchers n is exponentially larger than the number of design approaches m.*

The result highlights that merely looking at the number of searchers as has been typically done in past literature may result in severely overstating the breadth of solution space that gets covered. To summarize Corollaries 1, 2, and 3, the breadth of search, arguably of great importance to the principal, is costly to increase through awards and may be significantly lower than the number of participating searchers.

The next proposition characterizes the structural characteristics of the search equilibrium contingent on the problem type and degree of specification.

Proposition 2. (i) *As the problem becomes better specified (λ increases), the breadth of search and the number of searchers per design approach increase and then decrease.*

(ii) *The breadth of search increases as the problem admits more divergent solution approaches (ϕ increases).*

[10] We thank the review team for pointing out that reputational concerns may mitigate this effect of higher awards.

[11] We assume from this point on that $V(M) \le v\phi^{M-1}$, i.e., there is an award such that it is feasible for the principal to induce search over the entire solution space if it so chooses. If this condition is violated, then we may redefine M, the number of feasible design approaches, as maximum value of M that satisfies the above inequality.

Claim (i) characterizes the effect of problem specification on the number of searchers and design approaches. The extent of (completeness of) specification of the problem impacts the dynamics of delegated search in two ways: First, with problems that are very well specified (higher λ), a searcher has a greater likelihood of obtaining an acceptable solution. This makes participation in the search more attractive. However, when one has greater likelihood of finding acceptable solutions, the other searchers also have same greater likelihood of obtaining acceptable solutions. This latter effect must, in turn, reduce one's own chance of winning, thus making participation less valuable. The balance of these two opposing forces determines whether more or fewer searchers participate in the search and the breadth of coverage of the design space.

Claim (ii) shows that the nature of the solution space associated with the design problem, namely, whether the problem admits divergent or convergent solutions (i.e., ϕ is high or low), affects the searcher participation and the breadth of search in a relatively intuitive manner. When the problem admits divergent solutions, the searchers will search more broadly, resulting in less (relative) clustering on the more promising approaches. Figure 2 summarizes the insights. As can be observed, for a given award size and searcher cost structure, the breadth of search is greatest when the problem admits divergent design approaches (high ϕ), and the problem is moderately specified (moderate λ).

The goal of the analysis in this section was to offer a descriptive characterization of delegated search. The closed-form results we obtained were doubtless possible because of our simplifying assumption that searchers employing the same design approach obtain solutions of identical qualities (i.e., correlation of solution qualities is 1). Hence, before proceeding to examine the optimal management of delegated search, we relax this assumption of perfect correlation.

Figure 2 **Breadth of Search Contingent on Problem Type**
($c = 1, V = 10, v = 20, p = 1$)

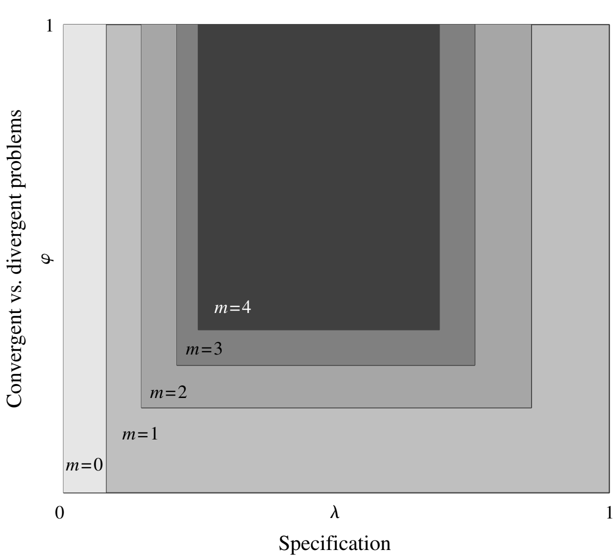

4.1. Case of Searcher-Specific Idiosyncratic Uncertainty

We assumed in the previous section that there is no searcher-specific uncertainty, i.e., because the solution quality depends only on the design approach, every searcher on a given alternative will find the same quality solution. In this section, we assume that in addition to the probability of success, p, associated with each design approach, there is an idiosyncratic searcher-specific probability of success θ; that is, the net probability of success if a searcher probes design approach i is given by $p\theta$. Thus, this extension offers an explicit means to capture idiosyncratic searcher characteristics (such as, perhaps, capability or luck). It may also be noted that by setting $\lambda = 1$, $\theta = \frac{2}{3}$, $M = 2$, $v = 3$, and $\phi = \frac{2}{3}$ and redefining p as γ, this expanded model reduces the numerical example presented in §1.[12] Consequently, even when two searchers probe the exact same design approach, there is still a positive probability that they will end up with solutions of different qualities. Thus,

$$\Pr(X_i = v_i \mid X_j = v_i) = \frac{\Pr(X_i = v_i, X_j = v_i)}{\Pr(X_j = v_i)} = \frac{p\theta^2}{p\theta} = \theta.$$

For ease of exposition of the subsequent results, we define the following.

Definition 1. The search distribution of any outcome O of the game is the vector $(k_1, k_2, \ldots, k_M)$ that summarizes the number of searchers that are probing the different design approaches in the outcome O; that is, if $(k_1, k_2, \ldots, k_M)$ is the search distribution associated with outcome O, then in the outcome O there are k_i searchers probing the design approach i.

Definition 2. Let $\psi_0 = 0$. For a given search distribution $(k_1, k_2, \ldots, k_M)$, recursively define p_i and ψ_i ($i \in \{1, 2, \ldots, M\}$) as follows:

$$\psi_i = p\lambda(1 - (1-\theta)^{k_i}),$$

$$p_i = \psi_i \prod_{j=0}^{i-1} (1 - \psi_j).$$

[12] Given our assumption of binomial distribution of solution qualities, once a searcher finds v_i from solution approach i, the marginal value from additional searchers probing the same solution approach is 0. Still, given the ex ante uncertainty, the expected value of the best solution from an approach is higher when more searchers pursue the approach; i.e., on expectation, the marginal value of additional searchers on a given solution approach is always positive. We thank the associate editor for noting this, a subtle point that may possible have implications to contexts where searchers can observe each other's efforts/performances.

Definition 3. Recursively define $(n_1, n_2, \ldots, n_M)$ as follows: For $i = 1, \ldots,$ let

$$n_i = \begin{cases} 0 & \text{if } \dfrac{p_i(n_1, n_2, \ldots, 1)}{1} < \dfrac{c}{V} \\ & \text{or if } i > \left\lfloor 1 + \dfrac{\ln(v) - \ln(V)}{-\ln\phi} \right\rfloor, \\ n & \text{such that } \dfrac{p_i(n_1, n_2, \ldots, n)}{n} \geq \dfrac{c}{V} \\ & \text{and } \dfrac{p_i(n_1, n_2, \ldots, n+1)}{n+1} < \dfrac{c}{V}. \end{cases}$$

Note that the recursive definition of n_i yields a unique vector $(n_1, n_2, \ldots, n_M)$ because

$$\frac{p_i}{n_i} = p\lambda \frac{(1-(1-\theta)^{n_i})}{n_i} \prod_{j=0}^{i-1} (1-\psi_j)$$

is a decreasing function of n_i (because $(1-x^n)/n$ is decreasing function of n when $x < 1$). Finally, define $m = \sum_{i=1}^{M} 1_{\{n_i > 0\}}$, and $n = \sum_{i=1}^{M} n_i$. As before, m and n represent the breadth of search and number of searchers, respectively.

For ease of exposition, we shall assume that the "entry + search − alternative choice" game has (at least) one equilibrium in pure strategies. Online Appendix A shows sufficient conditions when this assumption is satisfied.

Lemma 2. *An outcome is a* (*pure strategy*) *equilibrium iff its search distribution is* $(n_1, n_2, \ldots, n_M)$.

Corollary 4. (i) *More searchers pursue the design approach with greater potential*; $\{i < j\} \Rightarrow \{n_i \geq n_j\}$.

(ii) m *is* (*no greater than*) $O(\ln V)$, *where* V *is the award amount*.

(iii) m *is* (*no greater than*) $O(\ln n)$, *where* n *is the number of participating searchers*.

The corollary demonstrates that the main structural results derived earlier are surprisingly robust even when searcher-specific idiosyncratic uncertainty is considered. The intuition for the three claims mirrors the intuition offered earlier for Claim (i) of Proposition 1, Corollary 1, and Corollary 3, respectively. The next corollary mirrors Proposition 2 and verifies the robustness of the structural insights with respect to the role of problem type.

Corollary 5. (i) *There exists* θ_L *and* θ_H *such that* $m = 1$ *when* $\theta < \theta_L$ *or when* $\theta > \theta_H$.

(ii) *There exists* λ_L *and* λ_H *such that* $m = 1$ *when* $\lambda < \lambda_L$ *or when* $\lambda > \theta_H$.

(iii) m *is nondecreasing in* ϕ.

Although we cannot fully characterize the dependence of the search breadth on the specification, the corollary does verify that the main structural insight from earlier, namely, that search breadth is greatest for intermediate specification, is valid even in this extension. In addition, Claim (i) demonstrates a very similar effect of searcher-specific idiosyncratic uncertainty θ. Specifically, for very low θ, a searcher is unlikely to find the solution and would thus focus only on the more attractive solution approach (thus reducing the induced breadth). For very high θ, on the other hand, those searchers who focus on more attractive solution approaches have a high likelihood of obtaining the solution, thus reducing the value of probing ex-ante less attractive solution approaches.

It is interesting to note an ancillary benefit of Claim (i). Specifically, we conceptualized the specification parameter λ as determining the mismatch between a solution approach's particulars and the principal's needs. Hence, we modeled this parameter as a characteristic of the solution approach. However, Claim (i) shows that even if we had modeled λ as an idiosyncratic searcher-specific parameter, our result that broadest search occurs for intermediate specification remains valid.[13]

We offer a summation of the key descriptive characteristics of delegated search below before proceeding to analyze in §5 the optimal management of delegated search:

(A) Clustering is likely to occur wherein more searchers pursue design approaches with higher ex ante potential.

(B) The breadth of coverage of the design space is proportional to the logarithm of the number of participating searchers.

(C) The breadth of coverage of the design space is proportional to the logarithm of the award size (for awards less than a threshold).

(D) The breadth of coverage of the design space is greatest when the problem admits divergent design approaches and is moderately specified.

5. Managing Delegated Search

The descriptive results from the previous section illustrated the role of award size and problem specification on the search breadth. In this section, we analyze how the principal, contingent on the problem type, should optimally manage delegated search; specifically, we first investigate how the optimal single award should be set by the principal to maximize its profits. Subsequently, we also examine its drivers for offering multiple awards (§5.1), as well as the question of what resources it should devote to specifying the problem (§5.2).

[13] We thank the anonymous reviewer who pointed out that such mismatches may also arise from simple searcher-specific misunderstanding of the requirements, and that such mismatches are more likely to be idiosyncratic.

For completeness, we state below the optimal award as the solution of an optimization problem. Define

$$\Pi(m) = \left\{ \sum_{i=1}^{m} \lambda p(1-\lambda p)^{i-1} v\phi^{i-1} - V(m)(1-(1-\lambda p)^m) \right\},$$
where $V(m)$ is defined in Equation (3).

Then, the principal optimal expected payoff is $\max_m \Pi(m)$, and the optimal prize is $V(m^*)$, where $m^* = \arg\max_m \Pi(m)$. The next proposition completely characterizes the optimal search breadth m^* induced contingent on the problem type.

Proposition 3. (i) *m^* is increasing in ϕ.*
(ii) *m^* is decreasing in c.*
(iii) *m^* is increasing and then decreasing with λ.*

The proposition shows that the effect of problem type and degree of specification mirrors that discussed in Proposition 2. Thus, even when the award is set optimally by the principal, our previously found results regarding search breadth and its dependence on the problem type and specification prove robust; i.e., the optimal (induced) search breadth is highest for divergent problems that are moderately specified (see Figure 3 for a graphical illustration). In addition, the searcher cost c has an intuitive effect on the optimal induced breadth. Specifically, when the searcher cost c decreases, the principal benefits by inducing a broader search by setting the appropriate award.

The dependence of actual optimal award size set by the principal contingent on the problem parameters may be easily obtained by combining the preceding proposition with Equation (3). We state these sensitivity results as Corollary 6.

Corollary 6. (i) *The optimal award is increasing in searcher's cost (except at finitely many points where it exhibits a a discontinuous decrease).*
(ii) *The optimal award is monotone decreasing in λ when $m^* = 1$. For a given m^* (>1), optimal award is decreasing and then increasing in λ. (At the boundary, the optimal award is greater when the breadth is greater.)*
(iii) *The optimal award is increasing in ϕ.*

The effect of searcher's cost on the award size is addressed in Claim (i). Intuitively, higher searcher's cost introduces inefficiencies (because it results in lower breadth of search), and consequently the principal needs to set a higher award. Still, at the boundary points when searcher cost increases so much that the principal may choose to forgo enhanced search breadth, the award may decrease. Claim (ii) states that when the problem becomes more and more well specified (i.e., $1-\lambda$; the mismatch probability decreases), the searchers perceive less risk that their solution will not be acceptable to the principal. Consequently, the principal may reduce the optimal award size and still induce the searchers to participate. However, Claim (ii) also states that such a monotonic relation between specification and award size fails to exist when the optimal search breadth is large ($m^* \geq 2$). The intuition for this may be explained as follows: when the (optimal) search breadth is large, then increased specification makes searchers more prone to undertaking only the most promising design approach (i.e., clustering). Thus, to prevent such narrowing of the search space, the principal may have to increase the award size when the extent of specification is higher. Last, Claim (iii) demonstrates that when the problem admits divergent design approaches (i.e., multiple equally good solution approaches), the principal benefits from having the searchers pursue a variety of design approaches and should optimally set a higher award.

Next, we extend our baseline analysis of optimal single award along two dimensions. First, we examine the question of when the principal would wish to offer multiple awards. Second, given the importance of problem specification in determining the search breadth (highlighted by Proposition 2), we also examine the firm's decision of how many resources to expend toward specifying the problem. We simplify our previous model of the solution space and assume, from this point on, that there are only two design approaches, i.e., $M = 2$. This assumption enables mathematical tractability and allows ease of exposition while at the same time parsimoniously capturing the essential aspects of solution spaces.

5.1. Managing the Number and Size of Awards

So far we have assumed that the principal offered a single award V. Now we relax that assumption and let the principal announce two awards, V_1 and V_2 ($V_2 \leq V_1$), when it broadcasts the problem to outside agents; that is, the principal announces that the award V_1 shall be given (contingent on the principal's ex post rationality constraints) to the searcher who submits the best solution, and V_2 to the searcher who submits the second-best solution.[14]

The principal's ex post rationality constraints imply that, as in §3, after the searchers have submitted their solutions for evaluation, the principal assesses the quality of each of the design solutions. And the principal may, depending on which choice gives it greater payoffs, choose to accept one, two, or none of the solutions and consequently award either only V_1, V_1, and V_2, or no awards at all, respectively.[15]

Recall that the principal's estimate of the quality of the submitted solutions, based on our model, may

[14] Note that because there are only two design approaches, the principal will at best accept only two solutions. Hence, a third prize is simply not credible, and assuming a first prize V_1 and a second prize V_2 is adequate in our model.

[15] The constraint, as in §3, implies that the principal cannot commit to unconditionally giving the awards.

be either 0 or v for solutions that employed design approach 1, or 0 or ϕv for those that employed design approach 2. Furthermore, the principal has evaluation (market) uncertainty in that it believes that the true solution quality, when its estimate is v (ϕv), is distributed uniformly between $v-\delta$ and $v+\delta$ ($[\phi v-\delta, \phi v+\delta]$).

The principal obtains expected value v from accepting a single solution that uses approach 1 and that has an estimated quality v. Similarly, it obtains expected value ϕv from accepting a single solution that uses approach 2 and that has expected quality ϕv. Now consider the case where the principal accepts two solutions, one employing approach 1 and other employing approach 2. We assume in this case that the principal, after all uncertainty is resolved, will select the design approach that yields the highest payoff. Hence, the expected value to the principal from accepting both of the solutions is the expectation of the *realized* highest quality among the two accepted solutions. Note that given the ex ante uncertainty that the principal has (regarding the realization of evaluation/market uncertainty δ), it cannot with certainty ex ante select only the solution that has highest payoff. Thus, although the principal needs only one solution, with sufficient uncertainty δ (about its evaluation), it does have a positive marginal value from also accepting the solution that it evaluates to be of lower quality.

The next lemma characterizes the principal's optimal actions once the solutions have been submitted by the searchers; i.e., whether to accept one, two, or no solutions (and correspondingly make one, two, or no awards). For ease of exposition, we consider the case $V_1 < \phi v$, which ensures that it will be always be ex post rational for the principal to give at least one award, allowing us to focus on the second prize, V_2.

Lemma 3. (i) *If the submitted solutions employ approaches* 1 and 2, *and they have principal estimated quality v and ϕv, respectively, the principal*

- (*randomly*) *accepts one of the solutions employing approach* 1, *and awards a single prize V_1 to the corresponding searcher, and,*
- *if* $V_2 \le \max\{0, (2\delta - v + v\phi)^3/(24\delta^2\})$, (*randomly*) *accepts one of the solutions employing approach* 2, *and awards V_2 to the corresponding searcher.*

(ii) *If Case* (i) *is not true, and if at least some of the submitted solutions have a principal estimated quality that is nonzero, the principal* (*randomly*) *accepts one of these solutions* (*with nonzero quality*) *and awards V_1 to the corresponding searcher.*

(iii) *If neither Case* (i) *or* (ii) *is true, the principal accepts none of the solutions and awards no prizes.*

The lemma demonstrates that principal awards the second prize V_2 only when the submitted solutions span the full solution space. Intuitively, given our assumption that solutions from same approach have (perfectly) correlated qualities, the principal does not benefit from multiple solutions if these multiple solutions use the same design approach. The lemma offers an interesting constraint on the principal's credibility when offering a second prize, and thus on the use of the second prize V_2 summarized in the next observation.

Observation 2. The principal's announcement of a second prize V_2 is credible only if

$$V_2 \le \max\left\{0, \frac{(2\delta - v + v\phi)^3}{24\delta^2}\right\}.$$

The second prize is never credible (and thus never offered) when (i) the residual uncertainty δ (i.e., principal uncertainty after it estimates the solution quality) is low and/or (ii) the problem is more convergent (i.e., low ϕ). The underlying intuition for these effects is identical: when residual uncertainty is low or when problem is more convergent, then the additional benefit from accepting the poorer solution (i.e., the second solution) is low. Searchers anticipate this and, consequently, will not find announcement of any sizable second prize credible, thus ensuring that the principal will not offer the second prize.

The proposition highlights an important constraint on awards in delegated search, namely, the inability of the principal to commit to awarding a prize constrains both the number and size of prizes that may be offered. When the problem is more convergent and/or when residual (market or technical) uncertainty is low, the principal cannot use more than one prize. Although the vast majority of prior literature argues that awarding a "single prize seems consistent with the general intuition about the efficiency of rewarding only the best" (Moldovanu and Sela 2001, p. 543), some recent work suggests that second prizes may prove useful as a means to induce effort from low capability agents (Szymanskia and Valletti 2005) or to modify the risk behavior among agents (Gaba and Kalra 1999). Our findings complement these findings by identifying that, for convergent problems, multiple prizes may simply not be credible, whereas for divergent problems, multiple prizes are feasible.

Next, we completely characterize the size of the optimal awards that the principal sets contingent on the breadth of coverage it wishes to induce.

Lemma 4. (i) *The principal, to induce search over both the alternatives, will optimally set awards*

$$[V_2^*, V_1^*] = \left[\min\left\{\frac{c}{\lambda p}, \max\left\{0, \frac{(2\delta - v + v\phi)^3}{24\delta^2}\right\}\right\}, \frac{c - V_2^*(\lambda p)^2}{\lambda p(1-\lambda p)}\right].$$

Table 2 Sensitivity of Search Breadth and Optimal Awards with Problem Type

	δ	ϕ	λ
Search breadth m^*	↑	↑	↑ ↓
First prize V_1^*	↓	↓	↓ if $m^* = 1$, ↓ ↑ ↓ when $m^* = 2$
Second prize V_2^*	↑	↑	0 when $m^* = 1$, ↓ when $m^* = 2$

(ii) *The principal, to induce search over only the most promising alternative, will optimally set awards*

$$[V_2^*, V_1^*] = \left[0, \frac{c}{\lambda p}\right].$$

The lemma demonstrates that multiple awards are necessary only when the principal wishes to broaden the search. Intuitively, multiple awards allow the principal to induce a broader search, and if a broader search is not warranted, then the principal can forgo using multiple awards. The next proposition fully characterizes the dependence of breadth of search and award structure on the problem structure.

Proposition 4. *Table* 2 *summarizes the sensitivity of optimal search breadth and optimal award on the problem structure.*

Figure 3 graphically illustrates the proposition. The first column of Table 2 shows that both the search breadth and the size of the second prize are increasing in the ex post evaluation uncertainty δ. Intuitively, when the (market) uncertainty δ is high, it is optimal for the principal to defer choosing the final solution until later and to accept multiple solutions. Hence, the principal finds it optimal to broaden the induced search and will use a larger second prize to do so. The first prize, on the other hand, may decrease because the principal now has the additional lever of second prize and thus no longer needs to solely utilize the first prize to induce broader search. Thus, increasing δ results in a lower optimal first prize, a higher optimal second prize, and increased optimal search breadth.

Figure 3 Breadth of Search and Sensitivity of Optimal Award Structure Contingent on Problem Type ($v = 20, c = 1, p = 1$)

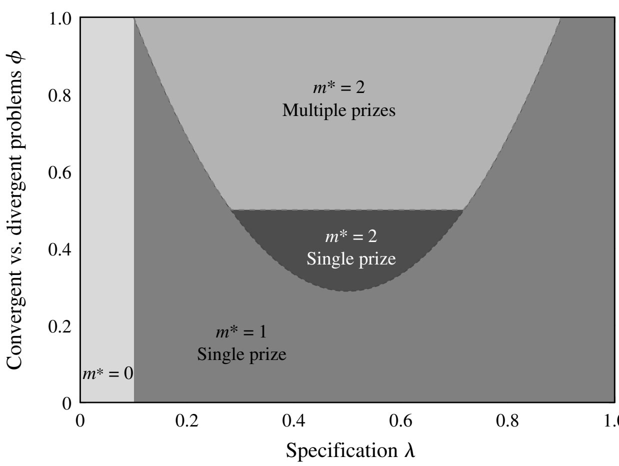

The second column of Table 2 offers the sensitivity of search breadth and award structure on the parameter ϕ. When the problem admits divergent design approaches (i.e., high ϕ), it is optimal to broaden the search and to use higher second prizes to do so. As argued in the previous paragraph, with greater ϕ, the ability to have a higher second prize to gain increased breadth allows the principal to reduce the size of the first prize.

The third column of Table 2 illustrates that the effect of λ is somewhat more complex: For both very low and very high λ, it is optimal to narrow the search and thus not employ second prizes, but for intermediate specifications, the firm benefits from broadening the search and utilizing second prizes. The second prize, when employed, is decreasing in the specification λ. Intuitively, as the problem gets better specified, the searchers will find limited risk that their evaluation of solution quality does not match with the principal's evaluation, and thus even smaller second prizes can induce search. The first prize is monotone decreasing in λ when $m^* = 1$ (because in this case the second prize is not used and the result identified in Proposition 6 applies). When $m^* = 2$, for low levels of λ, as λ increases, inability to set high enough second prizes results in the first prize initially decreasing and then increasing in λ, an effect identical to the one identified in Claim (ii) of Corollary 6. Still, as λ increases further, the lower second prize proves to be adequate to induce breadth, and the first prize decreases with λ.

In sum, there are two key insights that emerge from the proposition: First, the structure of optimal search breadth is remarkably similar to Proposition 2 (Figure 2). Thus, our key finding, namely, that the principal induces the greatest breadth for problems that have divergent design approaches and are moderately specified, appears to be relatively robust. Second, multiple awards are useful for moderately specified problems that have divergent design approaches.

5.2. Managing the "Problem Specification"

We now consider a second important managerial lever: problem specification. Recall that in our model, the parameter λ was interpreted as the extent (completeness) of the problem specification, because $1 - \lambda$ is the probability that there is a mismatch between the searcher's perception of the solution quality and the principal's perception of the solution quality. We assume in this section that the principal may choose the extent of resources it commits to problem specification. We capture this by assuming that there is a nondecreasing cost $C_p(\lambda)$ incurred by the principal when it specifies the problem upto λ. We assume throughout this section that the principal uses only a single award V_1. This assumption allows mathematical tractability and, as we argued earlier, is aligned with much of practice (and prior literature). For completeness, we offer next lemma, which characterizes the net payoff to the principal and the equilibrium of the delegated search, when the extent of specification is a decision variable.

Lemma 5. *Let*

$$\Pi(\lambda) = \lambda p v - c + \max\left\{\lambda p(1-\lambda p)v\phi - \frac{c}{1-\lambda p}\right\}. \quad (5)$$

Then, the principal specifies the problem up to λ^*, *where* $\lambda^* = \arg\max_{\lambda\in[0,1]}\{\Pi(\lambda) - C_p(\lambda)\}$.

The lemma lays out the optimization problem associated with the principal's optimal choice of resources devoted to problem specification. The functions $\Pi(\lambda)$ and $C_p(\lambda)$ are the revenues and costs when the principal specifies the problem to λ. Although we may intuitively expect the revenue function to be increasing in the specification (i.e., a more well-specified problem will yield greater revenue), the following example (shown in Figure 4) demonstrates that increasing the specification can lead to lower revenue! We state this as the following observation before explaining the underlying intuition.

Observation 3. The revenue $\Pi(\lambda)$ is not always increasing in λ.

The result, although surprising, has an interesting underlying intuition: As the problem becomes more specified (as λ increases), the breadth of design approaches may decrease as the searchers start clustering on only the most promising design approaches. To overcome this clustering, the principal would need to increase the award amount by a significant degree. Our results indicate that this increase in award amount may be large enough that it fails to offset any additional value gained by preventing the clustering. An interesting practical consequence of this nonmonotone behavior is stated as the following proposition.

Figure 4 **Revenues as a Function of Specification λ ($v = 64$, $c = 1$, $\phi = 0.9$, $p = 0.85$ for the Example)**

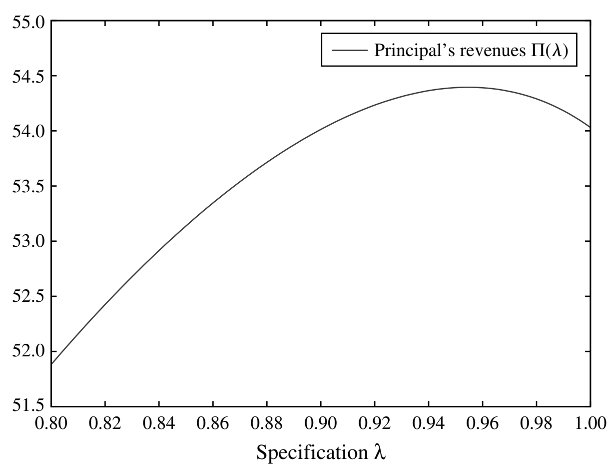

Proposition 5. *Even when* $C_p(\lambda) = 0$ (*i.e., no cost in specifying the problem completely*), *the principal may optimally choose not to specify the problem completely.*

The proposition is a direct result of the possibility of decreased revenues (with increasing specification) identified in the previous observation. Thus, even under the very generous assumption that it is *costless* to increase the fidelity of the specification, our results clearly demonstrate that the principal may not be well served by defining and communicating all the requirements as completely as possible to the searchers. Thus, the previously identified problem of clustering may result in the principal introducing some *strategic ambiguity* in the problem specification to induce searchers to spread out over the design space.

Intuitively, the optimal specification should decrease when the cost of specification is greater. However, the effect of searcher cost c on the optimal specification is more subtle and is characterized next.

Proposition 6. *There are thresholds* c_1, c_2, *and* c_3 *such that optimal specification* $\lambda^*(c)$ *is* (i) *decreasing in* c *for* $c < c_1$, (ii) *is equal to a constant* $\hat{\lambda}$ *for* $c_1 < c < c_2$, *and* (iii) *is equal to "0" for* $c > c_2$.

When searcher cost is very low, it is relatively easy to induce search over the entire solution space, and the principal is hence focused on offering as complete a specification as possible to reduce the potential for mismatch. As the searcher cost increases, however, the searchers start focusing on only the most promising approach unless this approach is less appealing. This may be accomplished by reducing the specification λ^*. Still, as the searcher costs become greater, the loss due to higher mismatch likelihood starts to dominate. Consequently, the principal will reduce the optimal award, narrowing the search; and with this narrower breadth, concerns about clustering become minimal, thus resulting in the optimal specification being mainly determined by the specification cost ($C_p(\cdot)$). Finally, for very high searcher costs, delegated search is infeasible, and the principal does not expend any resources on problem specification.

6. Conclusions

Motivated by the growing practice of organizations looking to outside agents to search for solutions to their complex problems, we develop a formal analytical model that helps understand how such delegated search should be managed. Our modeling framework captures several unique aspects of such contexts, including (i) the extent (or completeness) of problem specification and the potential mismatch between searcher- and firm-perceived solution qualities, and (ii) the delegating firm's lack of credibility when it announces awards associated the problem. We identify the following robust characteristics of the delegated search dynamics summarized below:

• *Clustering*. Clustering is likely to occur wherein more searchers pursue design approaches with higher ex ante potential.

• *Breadth of search*. The breadth of search is proportional to (i) the logarithm of the number of searchers and (ii) the logarithm of the award.

This issue of clustering and the associated difficulty in expanding the search breadth highlight the challenges in utilizing delegated search as a tool to explore a number of options in parallel in a cost-effective manner. Still, depending on the problem type, the appropriate choice of awards and the problem specification can counteract these difficulties and allow the firm to profitably manage delegated search.

Our results summarized in Table 2 and Figure 3 point to the effect of problem type on the optimal size and structure of awards. It appears that the issue of clustering in delegated search may be partially mitigated by increasing the number of awards. This allows us to identify a key reason for clustering, namely, a "winner-take-all" payoff structure resulting from a single award, which leads to more searchers exploring the ex ante promising alternatives instead of spreading out (as the principal would wish for). Multiple awards are preferable when (i) the principal has significant ex post (market) uncertainty of solution quality even after examining the solution, (ii) the problem is modestly specified, and/or (iii) the problem admits divergent solution approaches. Under all these three conditions, the principal greatly benefits from increased breadth of search and may optimally employ (possible large) multiple awards.

We also examined the question of how what resources a firm should devote toward specifying the problem (before broadcasting it to the searchers), and we find that specifying a problem completely may be suboptimal. This result holds true even if the problem specification may be completely costless. Our model allows us to identify that underspecifying a problem (i.e., leaving it somewhat ambiguous) may result in increased breadth of search, which may be beneficial for the focal firm. On the flip side, specifying a problem fully may result in searchers exploring only the most promising options and disregarding large portions of the solution space.

The primary identified issue of clustering, and more specifically, the concavity of search breadth with awards, suggests a number of fruitful extensions. We briefly discuss some of these next. As we have noted before, some of our results are consistent with recent empirical work in this area that identifies clustering of searchers. Still, although past research indicates the existence of clustering, it remains unclear as to its causes, and whether or not incentives can affect it. Our theoretical model makes predictions about how incentives and problem specification can alter clustering, and these predictions may be examined in experimental settings to gain greater understanding of the causal mechanisms underlying clustering.

Our focus in this paper is on developing an analytically tractable framework for understanding and improving the effectiveness of delegated search over design spaces. Thus, we took it as given that the firm employs delegated search and focused on the "how" of managing delegated search. Although a broader and more comprehensive investigation would be required to understand "when" a firm should conduct delegated search, our model (and results) offers a useful starting point for such a study. For instance, the ability to more effectively guide internal search (compared to delegated search, where the searchers choose the design approach) may have different implications for when and for which problem types delegated search may be preferred over in-house R&D. Thus, given the difficulty in increasing search breadth through the mechanism of delegated search, divergent design spaces (where obtaining breadth is important) might be more suited for in-house R&D (see Online Appendix D, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1567599, for a preliminary treatment). Similarly, cost differentials and principal's risk preferences may dictate the optimal choice between delegation and in-house R&D.

A recent paper examined the optimal design of R&D tournaments (Moldovanu and Sela 2006), asking if it is optimal to consider one grand R&D contest or a number of minicontests. They argue that minicontests (or subcontests) may be used to ensure that only high ability agents enter the "finals" where competition between high ability searchers ensures high effort levels from the searchers. The concavity result that we noted earlier suggests an alternate role for subcontests. Specifically, splitting a grand design contest (or delegated search, as we termed it) into a number of small design contests by subdividing the solution space may offer a mechanism design approach to managing delegated search. The contest approach to innovation is being used in the industry to tackle far more open-ended problems than what we have modeled. The retailer Staples has begun organizing

"Invention Quest" contests in which customers are invited to submit new office product ideas, and the best ones are launched and distributed by the firm (Staples 2005). Another specialist firm called Threadless uses a similar approach in the domain of apparel and creative content design (Lakhani and Kanji 2010). These are good examples of extremely open-ended problems, where there is not even a specific problem to solve, all innovative ideas are accepted, and the "best" ones rewarded. The domain of innovation tournaments is rapidly expanding. While many of these applications and extensions are interesting in their own right, examination of these issues awaits future work.

Acknowledgments
The authors acknowledge the feedback on this paper from Sreekumar Bhaskaran, Stelios Kavadias, and seminar participants at the University of Utah, Harvard Business School, and Queens University. Financial support to the second author from the Sheryl and Harvey White Endowment at the Rady School of Management, University of California, San Diego is gratefully acknowledged.

References
Anderson, E., A. Davis-Blake, S. Erzurumlu, N. Joglekar, G. Parker. 2008. The effects of outsourcing, offshoring, and distributed product development organizations on coordinating the NPD process. C. Loch, S. Kavadias, eds. *Handbook on New Product Development*. Elsevier, Boston, 165–197.

Auer, P., N. Cesa-Bianchi, P. Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.* **47**(2–3) 235–256.

Cameron, K. S. 1986. Effectiveness as a paradox: Consensus and conflicts in conceptations of organizational effectiveness. *Management Sci.* **32**(5) 539–553.

Che, Y.-K., I. Gale. 2003. Optimal design of research contests. *Amer. Econom. Rev.* **93**(3) 646–671.

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

Eppinger, S. D., A. R. Chitkara. 2006. The new practice of global product development. *MIT Sloan Management Rev.* **47** 22–30.

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

Gaba, A., A. Kalra. 1999. Risk behavior in response to quotas and contests. *Marketing Sci.* **18**(3) 417–434.

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

Huston, L., N. Sakkab. 2006. Connect and develop: Inside Procter & Gamble's new model for innovation. *Harvard Bus. Rev.* **84**(3) 58–66.

Kauffman, S., S. Levin. 1987. Towards a general theory of adaptive walks on rugged landscapes. *J. Theoret. Biol.* **128**(1) 11–45.

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

Lakhani, K., P. A. L. Lars Bo Jeppesen, J. A. Panetta. 2007. The value of openness in scientific problem solving. Working Paper 07-050, Harvard Business School, Boston.

Lakhani, K. R., Z. Kanji. 2010. Threadless: The business of community. HBS Case 608-707, Harvard Business School, Boston.

Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Sci.* **43**(7) 934–950.

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

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

Minsky, M. 1961. Steps toward artificial intelligence. *Proc. Inst. Radio Engineers* **49**(1) 8–29.

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

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

Simon, H. A. 1969. *The Sciences of the Artificial*, 1st ed. MIT Press, Cambridge, MA.

Skipso. 2009. The Rwanda Rural Electrification Challenge. Accessed August 25, 2011, http://skipso.wordpress.com/2009/08/08/the-rwanda-rural-electrification-challenge.

Smith, G. F. 1988. Towards a heuristic theory of problem structuring. *Management Sci.* **34**(12) 1489–1506.

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

Staples. 2005. America wants the SpindleRack! Staples® Invention Quest™ unveils the next great office product. Accessed August 25, 2011, http://multivu.prnewswire.com/mnr/staples/22668/.

Szymanskia, S., T. M. Valletti. 2005. Incentive effects of second prizes. *Eur. J. Political Econom.* **21**(2) 467–481.

Taylor, C. R. 1995. Digging for golden carrots: An analysis of research tournaments. *Amer. Econom. Rev.* **85**(4) 872–890.

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

Terwiesch, C., K. Ulrich. 2009. *Innovation Tournaments: Creating and Selecting Exceptional Opportunities*. Harvard Business School Press, Cambridge, MA.

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

Toubia, O. 2006. Idea generation, creativity, and incentives. *Marketing Sci.* **25**(5) 411–425.

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

Weitzman, M. L. 1980. Efficient incentive contracts. *Quart. J. Econom.* **94**(4) 719–730.