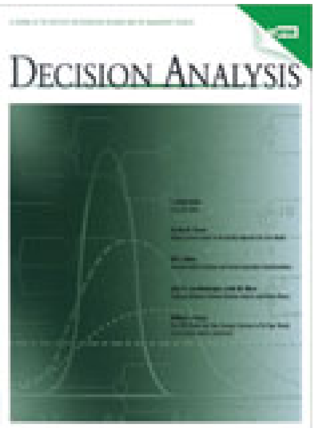

Decision Analysis

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Search Before Trade-offs Are Known

Otso Massala, Ilia Tsetlin

Decision Analysis
Vol. 12, No. 3, September 2015, pp. 105–121
ISSN 1545-8490 (print) | ISSN 1545-8504 (online)

http://dx.doi.org/10.1287/deca.2015.0313

Search Before Trade-offs Are Known

Otso Massala, Ilia Tsetlin
INSEAD, Singapore 138676
{otso.massala@insead.edu, ilia.tsetlin@insead.edu}

Search, broadly defined, is a critical managerial activity. Our contribution is a model of search for multiattribute alternatives, and our focus is on parallel search, where the decision is about the number of alternatives to explore. Most of the search literature considers univariate alternatives, and it can be applied to a multiattribute setting provided that the trade-offs to be used at the final selection stage were known at the search stage. However, uncertainty about trade-offs is likely to occur, especially in settings that involve parallel search (e.g., vendor selection, new product development, innovation tournaments). We show that incorporating uncertainty about trade-offs into a model changes its search strategy recommendations. Failing to account for such uncertainty, which is likely in practice, leads to suboptimal search and potentially large losses. For parallel search and a multivariate elliptical (e.g., normal) distribution of the alternatives, the solution is equivalent to univariate search with appropriately adjusted standard deviation. We prove that, in this setting, the optimal number of alternatives to explore increases if uncertainty about trade-offs increases, and we discuss the value of information about uncertain trade-offs.

Keywords: simultaneous search; parallel search; unknown trade-offs; multiattribute search
History: Received on July 24, 2014. Accepted by Editor-in-Chief Rakesh K. Sarin on January 28, 2015, after 2 revisions. Published online in *Articles in Advance* May 5, 2015.

1. Introduction

After searching for five months and considering more than a hundred chief executive officer (CEO) contenders, Microsoft directors chose internal candidate Satya Nadella. The search process was exhausting; Microsoft's board of directors discussed not only potential CEO candidates but also the pros and cons of an outsider versus an insider, future governance of the $300 billion company, and Bill Gates assuming a new board role as technology advisor, which in turn affected the CEO appointment (Microsoft 2014, Ovide et al. 2014). In the language of decision analysis, Microsoft's board of directors searched for multiattribute alternatives, and the trade-offs among attributes were not precisely known at the start.

The extensive search (i.e., considering more than 100 alternatives) conducted by Microsoft is consistent with insight from the search literature: a bigger difference of payoffs from different alternatives justifies the higher search effort and cost. We focus on other aspects of this example—that the search is for multiattribute alternatives and that trade-offs among different attributes (e.g., relative importance of the ability to lead change versus an understanding of Microsoft's complex business) are not entirely known at the beginning yet are clarified later.

Search, broadly defined, is ubiquitous. Research and development activities, new product development, idea generation, innovation tournaments, and finding a supplier or business partner are just some examples. In sequential search, the decision at each step is whether to explore one more alternative or to stop. In parallel search, which is the main focus of our paper, there is only one decision: how many alternatives to explore. In practically all situations, the alternatives are characterized by multiple attributes. Within the expected utility framework, the value of each multiattribute alternative is given by its utility, which depends on trade-offs among the attributes. This implies that search for multiattribute alternatives boils down to univariate search, subject to one caveat: the trade-offs to be used at the final selection stage are known in advance.

The search process can be viewed as consisting of two stages: the search stage, during which alternatives are explored; and the selection stage, at which one of the explored alternatives is chosen. At the selection stage, the discovered alternatives are ranked by explicit comparison of their utilities or, more simply, via ordinal ranking. In either case, this ranking is based on (implicit or explicit) trade-offs among the attributes. If the search process takes time, then it is reasonable to assume that, during that time, new information about the trade-offs (independent of the search process itself) will come to light; therefore, the trade-offs to be used at the selection stage might not be known at the search stage. As noted by March (1978, p. 597),

> We expect change in our preferences. As we contemplate making choices that have consequences in the future, we know that our attitudes about possible outcomes will change in ways that are substantial but not entirely predictable. The subjective probability distribution over possible future preferences (like the subjective probability distribution over possible future consequences) increases its variance as the horizon is stretched. As a result, we have a tendency to want to take actions now that maintain future options for acting when future preferences are clearer.

Uncertain trade-off might occur even in the case of a single-attribute utility function. Bell (1988) discusses the setting where the mean–variance trade-off is affected by the resolution of a side bet. That setting would match ours if the side bet is resolved between the search and selection stages, and we focus on how this anticipated resolution affects the search strategy.

Our methodological contribution is a model of search for multiattribute alternatives. In §4 we consider parallel search when the trade-offs between attributes are unknown at the search stage and the decision maker is risk neutral. If the alternatives are drawn from a multivariate elliptical distribution (defined in §4.1), then the solution to multivariate search is equivalent to the solution to univariate search with appropriately adjusted standard deviation, and it is optimal to search more as the uncertainty about trade-offs increases. The family of elliptical distributions is a large one and includes the multivariate normal, t, logistic, exponential, and other distributions. As a multivariate elliptical (e.g., normal) distribution is likely to occur in practice, our results provide managerial guidelines for search decisions, as illustrated in §4.3. In §5 we consider parallel search with outside (already available) alternatives and the implications for sequential search.

A review of the literature is in §2, where we give special attention to settings where parallel search is used (e.g., vendor selection) and to the multiattribute nature of the alternatives in such settings. Parallel search is used when exploring an alternative takes so much time that sequential search is precluded. Then the search and selection decisions are separated by substantial time, during which some information about trade-offs is likely to arrive. It is for these cases that our model is most relevant. The univariate parallel search is reviewed in §3.

2. Background and Literature Review

Search models have been extensively studied and applied to a variety of settings (for a recent review, see Rogerson et al. 2005). In sequential search, the decision is whether to explore one more alternative or to stop (e.g., Lippman and McCall 1976); in parallel search (Nelson 1961, Stigler 1961), the decision concerns the number of alternatives to explore, after which the alternative with the highest payoff is selected. A more general model combines these two decisions, so that in each period the searcher decides the number of alternatives to explore (Morgan 1983, Morgan and Manning 1985).

The search literature focuses almost exclusively on univariate alternatives. A few papers (e.g., Bearden and Connolly 2007, Lim et al. 2006; see also the references therein) consider sequential search for multiattribute options along with different costs for discovering a new option and for learning the attributes of a particular option. In that research, the focus is on balancing search over options and search within options. In spirit, this is similar to the case of univariate options, whose uncertain values can be learned at extra cost (Lippman and McCardle 1991).

2.1. When Parallel Search Is Used

When exploring an alternative takes a lot of time, the decision maker seldom has the luxury of searching for more than one period. In such cases, parallel search is the only feasible option. Parallel search is also favored by the very design of tournaments and procurements.

Indeed, the model of parallel search has been applied in many contexts that include procurement, new product development, and innovation tournaments. In a procurement setting, the time-consuming assessment of suppliers and the demand for transparency both argue for a parallel search approach (Costantino et al. 2012, Heijboer and Telgen 2002). With new product development, parallel search is used to reduce the time required to develop and test new solutions with the aim of benefitting from swifter responses to market opportunities (Loch et al. 2001). Dahan and Mendelson (2001) and Srinivasan et al. (1997) apply a parallel search model to new product development. In innovation tournaments, both the tournament design and the low marginal sampling cost favor a parallel search process (Terwiesch and Ulrich 2009, Terwiesch and Xu 2008). Kornish and Ulrich (2011) show empirically that redundancy (one potential weakness of parallel search for new ideas) is quite small. The various streams of the literature have developed special terms for the optimal number of alternatives to explore, such as "economic tender quantity" (de Boer et al. 2000) and "optimal batch size" for companies developing new products (Loch et al. 2001).

2.2. Multiattribute Alternatives and Uncertain Trade-offs

Most of the time, alternatives are characterized by multiple attributes. In the context of vendor selection, Ho et al. (2010, p. 21) conclude, "The traditional single criterion approach based on lowest cost bidding is no longer supportive and robust enough in contemporary supply management." An alternative criterion for evaluating bids is the "most economically advantageous tender" (see Costantino et al. 2012, p. 190). A survey of procurement practitioners concludes that incorporating multiple criteria (in addition to price) into the decision support tool used to determine the optimal number of bidders would increase considerably the number of situations to which the model is applicable (Heijboer and de Boer 2001). Different attributes are typically aggregated by using a weighted sum (Lorentziadis 2010), which corresponds to additive utility in the decision analysis literature. The value of an alternative with attributes $y_1, \ldots, y_M$ is given by $k_1 y_1 + \cdots + k_M y_M$; here, $k_1, \ldots, k_M$ are the attribute weights, often called "importance weights" (Lorentziadis 2010). Because weight k_j is also a trade-off between money (search cost) and attribute j, $j = 1, \ldots, M$, throughout the paper we will refer to these weights as *trade-offs*.

If parallel search is used, then there is a significant time delay between the search stage (when the number of alternatives to explore is chosen) and the selection stage (when the best of the explored alternatives is selected). Trade-offs $k_1, \ldots, k_M$ are numbers to be used at the selection stage, and given the significant time that elapses between search and selection decisions, they are likely to be unknown at the search stage. Lorentziadis (2010) discusses just such a case in the context of supplier selection.

3. Univariate Model of Parallel Search

In the classical model of parallel search (Nelson 1961, Stigler 1961), the decision variable is the number n of independent identical draws of random variable X. The searcher knows the distribution of X. Each draw costs $c > 0$. After collecting $n \geq 1$ draws $X_1, \ldots, X_n$, the decision maker chooses the best of them. Thus, the payoff equals $\max_{i=1,\ldots,n}(X_i)$, the highest of these n draws, minus the search cost nc. For n draws, the expected payoff[1] is given by $\pi(c, n; X) = -nc + E_{X_1,\ldots,X_n}[\max_{i=1,\ldots,n}(X_i)]$.

Theorem 1. *Let $E[|X|]$ exist. Then the expected payoff $\pi(c, n; X) = -nc + E[\max_{i=1,\ldots,n}(X_i)]$ is concave in n and $\lim_{n\to\infty}(\pi(c, n+1; X) - \pi(c, n; X)) = -c$. So for $c > 0$, the expected payoff is maximized at*

$$n^*(c; X) = \arg\max_n(\pi(c, n; X)) = \min\{n: \pi(c, n+1; X) - \pi(c, n; X) < 0\}.$$

Proof. Denote by F_X the cumulative distribution function (cdf) of X. Then the cdf of $\max_{i=1,\ldots,n}(X_i)$ is F_X^n. Since $E[|X|]$ exists, it follows that $E[|\max_{i=1,\ldots,n}(X_i)|]$ also exists and that, by David and Nagaraja (2003, p. 34, Footnote 1, and Equation (3.1.10′)),

$$E_{X_1,\ldots,X_n}\Big[\max_{i=1,\ldots,n}(X_i)\Big] = \int_0^\infty (1 - F_X^n(x) - F_X^n(-x))\,dx.$$

Therefore,

$$\pi(c, n+1; X) - \pi(c, n; X) = \int_0^\infty \big(F_X^n(x)(1 - F_X(x)) + F_X^n(-x)(1 - F_X(-x))\big)\,dx - c.$$

[1] This expected payoff is shown to be concave in Nelson (1961) for bounded X and in Benhabib and Bull (1983), de Boer et al. (2000), and Srinivasan et al. (1997) for X bounded from below. For completeness, we formally state and prove Theorem 1, because we could find no proof in the literature for the case where X is unbounded from below and from above.

This difference is decreasing in n (as the expression inside the integral is decreasing in n) and approaches $-c$ as n goes to infinity. Expected payoff $\pi(c, n; X)$ is concave in n because $\pi(c, n+1; X) - \pi(c, n; X)$ decreases with n. □

For a random variable X, let

$$\omega(n; X) = E\Big[\max_{i=1,\dots,n}(X_i)\Big],$$
$$n^*(c; X) = \min\big(n: \omega(n+1; X) - \omega(n; X) < c\big). \quad (1)$$

Note that the optimal number of draws is adjusted in a simple manner if the distribution of X is modified via a linear transformation (i.e., if the shape of the distribution is preserved): for $X = \mu + \sigma Z$ with $\sigma > 0$, we have $\omega(n; X) = \mu + \sigma\omega(n; Z)$ and $n^*(c; X) = n^*(c/\sigma; Z)$; for any Z, this can be computed numerically.

4. Parallel Search Before Trade-offs Are Known

We now extend the parallel search model from §3 to multiattribute alternatives with uncertain trade-offs. Each draw yields a multiattribute alternative $\mathbf{Y}$, and the uncertain trade-offs are denoted by $\mathbf{K}$. (We use boldface capital letters to indicate random vectors.) The vector $\mathbf{Y}$ captures uncertainty among attribute values of alternatives that might be available because of search, and the vector $\mathbf{K}$ captures uncertainty about relative merits of different components of $\mathbf{Y}$. Components of $\mathbf{Y}$ can be dependent and components of $\mathbf{K}$ can be dependent, but components of $\mathbf{Y}$ are independent of components of $\mathbf{K}$. Independence of $\mathbf{K}$ and $\mathbf{Y}$ is realistic (the former captures uncertainty about future preference among the attributes, and the latter captures uncertainty about different alternatives from the search process) and also makes our problem tractable.

For a particular $\mathbf{k}$, the payoff from alternative $\mathbf{y}$ is given by the valuation function $v(\mathbf{y}, \mathbf{k})$. If $v(\mathbf{y}, \mathbf{k}) = k_1 y_1 + \cdots + k_M y_M$, then each k_j is a trade-off between attribute y_j and money (search cost); we will focus on this case and refer to $\mathbf{k}$ as trade-offs. However, some of our theoretical results (in particular, Theorems 2 and 4) are applicable to arbitrary $v(\mathbf{y}, \mathbf{k})$; then $\mathbf{k}$ corresponds to uncertain parameters of the valuation function.

As in §3, the decision maker collects n independent identical draws from a multivariate distribution that is known to this decision maker, and each draw costs c. If alternative $\mathbf{y}$ is selected, then the payoff is $-nc + v(\mathbf{y}, \mathbf{k})$, so the payoff depends both on the selected alternative $\mathbf{y}$ and on the realized trade-offs $\mathbf{k}$. A risk-neutral decision maker maximizes the expected payoff.

The decision maker faces two decisions: choosing the number of draws n at the search stage and, after collecting n alternatives, choosing one of them at the selection stage. The second decision is fairly straightforward. The first decision—about the number of draws—is more challenging and depends on when the uncertainty about $\mathbf{K}$ is resolved.

There are three different scenarios that correspond to the different times at which $\mathbf{K}$ may become known (see Figure 1). First, suppose $\mathbf{k}$ is known at the search stage—that is, before the decision about the number of draws. Then, after observing n draws with realizations $\mathbf{y}_1, \dots, \mathbf{y}_n$, the decision maker will choose the alternative that maximizes $v(\mathbf{y}_i, \mathbf{k})$. The expected payoff for n draws is

$$\pi(c, n; v(\mathbf{Y}, \mathbf{k})) = -nc + E_{\mathbf{Y}_1,\dots,\mathbf{Y}_n}\Big[\max_{i=1,\dots,n} v(\mathbf{Y}_i, \mathbf{k})\Big],$$

and by Theorem 1, the optimal number of draws is $n^*(c; X)$ with $X = v(\mathbf{Y}, \mathbf{k})$.

Second, suppose $\mathbf{K}$ is not known at the search stage and remains unknown at the selection stage (i.e., when the decision maker is choosing one of the n alternatives). In this scenario, after observing n draws, the decision maker will choose the alternative that maximizes $E_{\mathbf{K}}[v(\mathbf{y}_i, \mathbf{K})]$. The expected payoff for n draws is $\pi(c, n; E_{\mathbf{K}}[v(\mathbf{Y}, \mathbf{K})]) = -nc + E_{\mathbf{Y}_1,\dots,\mathbf{Y}_n}(\max_{i=1,\dots,n} E_{\mathbf{K}}[v(\mathbf{Y}_i, \mathbf{K})])$ and—again by Theorem 1—the optimal number of draws is $n^*(c; X)$ with $X = E[v(\mathbf{Y}, \mathbf{K}) \mid \mathbf{Y}] = E_{\mathbf{K}}[v(\mathbf{Y}, \mathbf{K})]$. In sum, if trade-offs either are known before the search stage or remain unknown at the selection stage, then multiattribute search is equivalent to the classical univariate search problem in which each alternative is characterized by a single attribute (payoff). To determine the optimal number of draws in such cases, it is enough to estimate the univariate distribution of $v(\mathbf{Y}, \mathbf{k})$ in the first scenario and of $E_{\mathbf{K}}[v(\mathbf{Y}, \mathbf{K})]$ in the second.

The focus of our paper is on the third scenario, where trade-offs become known before selecting one

Figure 1 **Resolving Uncertainty About Trade-offs: Three Scenarios**

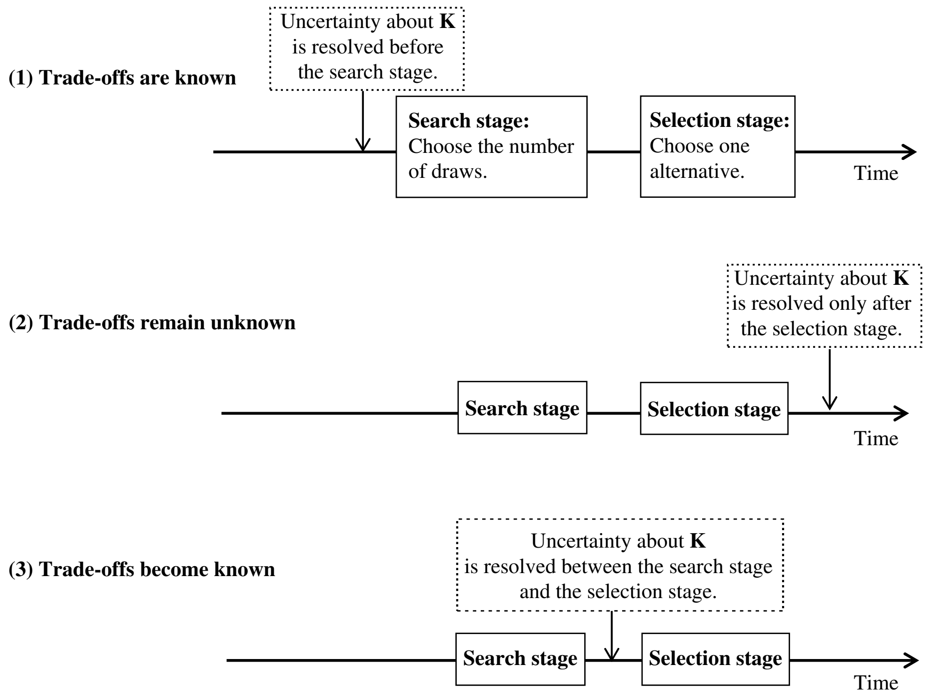

of the discovered alternatives but only after the decision about the number of draws has been made. (We refer to this scenario as "trade-offs become known.") After n alternatives are drawn, the choice among $\mathbf{y}_1, \ldots, \mathbf{y}_n$ might depend on the realized value of $\mathbf{k}$. In turn, the optimal number of draws is affected by knowing that the uncertainty about $\mathbf{K}$ will be resolved before the selection stage. The expected payoff for n draws is given by

$$\pi(c,n;\mathbf{Y},\mathbf{K}) = -nc + E\Big[\max_{i=1,\ldots,n} v(\mathbf{Y}_i,\mathbf{K})\Big], \qquad (2)$$

where the expectation is over $\mathbf{K}$ and $\mathbf{Y}_1, \ldots, \mathbf{Y}_n$.

Theorem 2. *Let $E[|v(\mathbf{Y},\mathbf{K})|]$ exist. Then the expected payoff $\pi(c,n;\mathbf{Y},\mathbf{K})$ given by (2) is concave in n and $\lim_{n\to\infty}(\pi(c,n+1;\mathbf{Y},\mathbf{K}) - \pi(c,n;\mathbf{Y},\mathbf{K})) = -c$. So for $c > 0$, the expected payoff is maximized at $n^*(c;\mathbf{Y},\mathbf{K}) = \arg\max_n(\pi(c,n;\mathbf{Y},\mathbf{K})) = \min\{n\colon\ \pi(c,n+1;\mathbf{Y},\mathbf{K}) - \pi(c,n;\mathbf{Y},\mathbf{K}) < 0\}$.*

Proof. Using independence of $\mathbf{Y}$ and $\mathbf{K}$, the expected payoff (2) can be written as

$$\pi(c,n;\mathbf{Y},\mathbf{K}) = E_{\mathbf{K}}\Big[-nc + E_{\mathbf{Y}_1,\ldots,\mathbf{Y}_n}\Big[\max_{i=1,\ldots,n} v(\mathbf{Y}_i,\mathbf{K})\Big]\Big]$$
$$= E_{\mathbf{K}}\big[\pi(c,n;v(\mathbf{Y},\mathbf{K}))\big].$$

It is the expectation (with respect to $\mathbf{K}$) of a function that, by Theorem 1, is concave in n and has marginal change approaching $-c$ as n goes to infinity. Therefore, the optimal number of draws $n^*(c;\mathbf{Y},\mathbf{K})$ is the smallest n for which the change in the expected payoff becomes negative. □

By Theorem 2, the expected payoff when trade-offs become known is a well-behaved function of n, and maximizing it numerically (given the distributions of $\mathbf{Y}$ and $\mathbf{K}$) is feasible in each particular case. Because it is beneficial to know trade-offs before choosing one of the alternatives, the expected payoff (2) is not less than $\pi(c,n;E_{\mathbf{K}}[v(\mathbf{Y},\mathbf{K})]) = -nc + E_{\mathbf{Y}_1,\ldots,\mathbf{Y}_n}[\max_{i=1,\ldots,n} E_{\mathbf{K}} v(\mathbf{Y}_i,\mathbf{K})]$, which is the expected payoff when trade-offs remain unknown (and the two payoffs are equal only if the optimal alternative is, almost surely, the same for any $\mathbf{k}$).

It might be tempting to conjecture that if the uncertainty about $\mathbf{K}$ is resolved before the selection stage, the optimal number of draws should be no smaller than if $\mathbf{K}$ remains unknown, because in the former case it is more beneficial to create a wider choice of alternatives ("maintain future options for acting when future preferences are clearer"; see

March 1978, p. 597). In §4.1 we show that this is indeed the case if $v(\mathbf{y}, \mathbf{k})$ is linear in $\mathbf{y}$ and if $\mathbf{Y}$ has a multivariate elliptical distribution. In §4.2 we present counterexamples that underscore the difficulty of generalizing this result to nonelliptical distributions.

4.1. Search from a Multivariate Elliptical Distribution

Let the random vector $\mathbf{Y}$ have M-variate elliptical distribution with mean vector $\boldsymbol{\mu}_\mathbf{Y} = (\mu_1, \ldots, \mu_M)^T$ and covariance matrix $\boldsymbol{\Sigma}$. (Superscript T denotes transposition.) The family of elliptical distributions is large and includes multivariate normal, nonnormal variance mixtures of multinormal, t, exponential, logistic, and other distributions (Fang et al. 1990, Table 3.1; see also Owen and Rabinovitch 1983). There are several definitions of elliptical distributions.

Definition 1. Random vector $\mathbf{Y}$ has elliptical distribution if and only if, for all M-component scalar vectors $\mathbf{a}$, all the univariate random variables $\mathbf{a}^T\mathbf{Y}$ such that $\mathrm{Var}(\mathbf{a}^T\mathbf{Y})$ is constant follow the same distribution (Owen and Rabinovitch 1983, Definition (b)). We denote the corresponding standardized random variable by $Z_\mathbf{Y}$, so that $\mathbf{a}^T\mathbf{Y}$ has the same distribution as $E(\mathbf{a}^T\mathbf{Y}) + \mathrm{Var}(\mathbf{a}^T\mathbf{Y})Z_\mathbf{Y} = \mathbf{a}^T\boldsymbol{\mu}_\mathbf{Y} + (\mathbf{a}^T\boldsymbol{\Sigma}\mathbf{a})Z_\mathbf{Y}$.

There are alternative definitions,[2] and the differences are important if first or second moments do not exist. To simplify the exposition, we assume that these moments do exist for $\mathbf{Y}$. Note that $Z_\mathbf{Y}$ from Definition 1 can also be defined as $Z_\mathbf{Y} = (Y_j - \mu_j)/\sqrt{\boldsymbol{\Sigma}_{jj}}$ for any $j = 1, \ldots, M$, and it corresponds to the standardized marginal distribution of all components of $\mathbf{Y}$. For example, if $\mathbf{Y}$ is M-variate normal, then $Z_\mathbf{Y}$ is standard normal.

In this section we assume also that $v(\mathbf{y}, \mathbf{k}) = \sum_{j=1}^{M} y_j k_j = \mathbf{k}^T\mathbf{y}$. As before, $\mathbf{K}$ is independent of $\mathbf{Y}$, and we assume that $E(|K_j|)$ exists for $j = 1, \ldots, M$. Let $\boldsymbol{\mu}_\mathbf{K} = E(\mathbf{K})$. Lemma 1 summarizes the solution for optimal parallel search in this setting and shows that it is equivalent to the solution for univariate search with appropriately adjusted standard deviation.

[2] For example, another definition (Owen and Rabinovitch 1983, Definition (c)) is that *random vector* $\mathbf{Y}$ *has elliptical distribution if its density function is only a function of the quadratic form* $(\mathbf{Y} - \boldsymbol{\mu}_\mathbf{Y})^T\boldsymbol{\Sigma}^{-1}(\mathbf{Y} - \boldsymbol{\mu}_\mathbf{Y})$.

Lemma 1. *Let* $\mathbf{Y}$ *be multivariate elliptical with mean vector* $\boldsymbol{\mu}_\mathbf{Y}$ *and covariance matrix* $\boldsymbol{\Sigma}$, *and let* $v(\mathbf{y}, \mathbf{k}) = \mathbf{k}^T\mathbf{y}$. *Then the expected payoff* (2) *is* $\pi(c, n; \mathbf{Y}, \mathbf{K}) = -nc + \boldsymbol{\mu}_\mathbf{K}^T\boldsymbol{\mu}_\mathbf{Y} + E[\sigma(\mathbf{K})]\omega(n; Z_\mathbf{Y})$. *Here,* $\sigma(\mathbf{k}) = \sqrt{\mathbf{k}^T\boldsymbol{\Sigma}\mathbf{k}}$, $Z_\mathbf{Y} = (Y_1 - \mu_1)/\sqrt{\boldsymbol{\Sigma}_{11}}$ *(by Definition* 1*), and* $\omega(n; Z_\mathbf{Y})$, *as defined in* (1), *is the expected value of the nth-order statistic of* $Z_\mathbf{Y}$. *For* $c > 0$, *the optimal number of draws is* $n^*(c/E[\sigma(\mathbf{K})]; Z_\mathbf{Y})$ *with* n^* *defined in* (1).

Proof. For $v(\mathbf{y}, \mathbf{k}) = \mathbf{k}^T\mathbf{y}$, Equation (2) becomes $\pi(c, n; \mathbf{Y}, \mathbf{K}) = -nc + E_\mathbf{K}[E_{\mathbf{Y}_1, \ldots, \mathbf{Y}_n}[\max_{i=1,\ldots,n}(\mathbf{K}^T\mathbf{Y}_i)]]$. For given $\mathbf{k}$, the random variable $\mathbf{k}^T\mathbf{Y}$ has mean $\mathbf{k}^T\boldsymbol{\mu}_\mathbf{Y}$ and standard deviation $\sigma(\mathbf{k}) = \sqrt{\mathbf{k}^T\boldsymbol{\Sigma}\mathbf{k}}$, and (for elliptical $\mathbf{Y}$, by Definition 1) its shape corresponds to $Z_\mathbf{Y}$—that is, $\mathbf{k}^T\mathbf{Y} = \mathbf{k}^T\boldsymbol{\mu}_\mathbf{Y} + \sigma(\mathbf{k})Z_\mathbf{Y}$. Then, by (1), we have $E_{\mathbf{Y}_1, \ldots, \mathbf{Y}_n}[\max_{i=1,\ldots,n}(\mathbf{k}^T\mathbf{Y}_i)] = \mathbf{k}^T\boldsymbol{\mu}_\mathbf{Y} + \sigma(\mathbf{k})\omega(n; Z_\mathbf{Y})$; taking the expectation over $\mathbf{K}$ yields $\pi(c, n; \mathbf{Y}, \mathbf{K}) = -nc + \boldsymbol{\mu}_\mathbf{K}^T\boldsymbol{\mu}_\mathbf{Y} + E[\sigma(\mathbf{K})]\omega(n; Z_\mathbf{Y})$. By Theorem 1, the optimal number of draws is $n^*(c/E[\sigma(\mathbf{K})]; Z_\mathbf{Y})$. □

As Lemma 1 shows, multiattribute search (in the "trade-offs become known" scenario) from a multivariate elliptical distribution $\mathbf{Y}$ is equivalent to search from a $Z_\mathbf{Y}$-shaped univariate distribution with mean $\boldsymbol{\mu}_\mathbf{K}^T\boldsymbol{\mu}_\mathbf{Y}$ and standard deviation $E[\sigma(\mathbf{K})]$. In the "trade-offs remain unknown" scenario, the expected payoff from n draws is equal to $-nc + \boldsymbol{\mu}_\mathbf{K}^T\boldsymbol{\mu}_\mathbf{Y} + \sigma(\boldsymbol{\mu}_\mathbf{K})\omega(n; Z_\mathbf{Y})$, which is equivalent to the expected payoff for the search from a $Z_\mathbf{Y}$-shaped univariate distribution with mean $\boldsymbol{\mu}_\mathbf{K}^T\boldsymbol{\mu}_\mathbf{Y}$ and standard deviation $\sigma(\boldsymbol{\mu}_\mathbf{K})$. Theorem 3 compares these two scenarios.

Theorem 3. *Let* $\mathbf{Y}$ *be multivariate elliptical, and let* $v(\mathbf{y}, \mathbf{k}) = \mathbf{k}^T\mathbf{y}$. *The optimal number of draws if trade-offs become known is never less than the optimal number of draws if trade-offs remain unknown (or are known to equal* $\boldsymbol{\mu}_\mathbf{K}$*).*

Proof. First, we prove that $\sigma(\mathbf{k})$ is convex; that is, for any $\mathbf{k}_1$, $\mathbf{k}_2$ and $\gamma \in (0, 1)$,

$$\sqrt{(\gamma\mathbf{k}_1 + (1-\gamma)\mathbf{k}_2)^T\boldsymbol{\Sigma}(\gamma\mathbf{k}_1 + (1-\gamma)\mathbf{k}_2)} \le \gamma\sqrt{\mathbf{k}_1^T\boldsymbol{\Sigma}\mathbf{k}_1} + (1-\gamma)\sqrt{\mathbf{k}_2^T\boldsymbol{\Sigma}\mathbf{k}_2}.$$

Let $X_1 = \gamma\mathbf{k}_1^T\mathbf{Y}$ and $X_2 = (1-\gamma)\mathbf{k}_2^T\mathbf{Y}$. Then the inequality's left-hand side is the standard deviation of $X_1 + X_2$, and its right-hand side is the sum of the standard deviations of X_1 and X_2 (since $\boldsymbol{\Sigma} =$

$E(\mathbf{Y}\mathbf{Y}^T) - \boldsymbol{\mu}_\mathbf{Y}\boldsymbol{\mu}_\mathbf{Y}^T$). The inequality is strict unless X_1 and X_2 are perfectly positively correlated. Therefore, $E[\sigma(\mathbf{K})] \geq \sigma(\boldsymbol{\mu}_\mathbf{K})$.

Second, by Lemma 1, if trade-offs become known, then the optimal number of draws is $n^*(c/E[\sigma(\mathbf{K})]; Z_\mathbf{Y})$, whereas if trade-offs remain unknown, then the optimal number of draws is $n^*(c/\sigma(\boldsymbol{\mu}_\mathbf{K}); Z_\mathbf{Y})$. Since $n^*(c; Z_\mathbf{Y})$ is decreasing in c, $E[\sigma(\mathbf{K})] \geq \sigma(\boldsymbol{\mu}_\mathbf{K})$ implies $n^*(c/E[\sigma(\mathbf{K})]; Z_\mathbf{Y}) \geq n^*(c/\sigma(\boldsymbol{\mu}_\mathbf{K}); Z_\mathbf{Y})$. □

Theorem 3 exploits two facts. First, the standard deviation $\sigma(\mathbf{k})$ of a random variable $v(\mathbf{Y}, \mathbf{k}) = \mathbf{k}^T\mathbf{Y}$ is convex in $\mathbf{k}$, which is true for any distribution of $\mathbf{Y}$. Second, $\mathbf{k}^T\mathbf{Y}$ is a linear transformation of $Z_\mathbf{Y}$ because $\mathbf{Y}$ is elliptical. Then $\pi(c, n+1; \mathbf{k}^T\mathbf{Y}) - \pi(c, n; \mathbf{k}^T\mathbf{Y}) = -c + \sigma(\mathbf{k})(\omega(n+1; Z_\mathbf{Y}) - \omega(n; Z_\mathbf{Y}))$ is also convex in $\mathbf{k}$, because $\omega(n+1; Z_\mathbf{Y}) > \omega(n; Z_\mathbf{Y})$. For nonelliptical distribution of $\mathbf{Y}$, the shape of the distribution of $\mathbf{k}^T\mathbf{Y}$ depends on $\mathbf{k}$. Because of that, as shown in §4.2, it is difficult to extend Theorem 3 to nonelliptical distributions.

We now discuss how the optimal number of alternatives to explore in the "trade-offs become known" scenario depends on the distributions of $\mathbf{Y}$ and $\mathbf{K}$. Table 1 summarizes the results, which are discussed later in this section.

What is the impact of variability of $\mathbf{Y}$ (i.e., of $\boldsymbol{\Sigma}$) on the optimal number of draws? Let $\boldsymbol{\Sigma}_{ij} = \sigma_i\sigma_j\rho_{ij}$. From

$$\sigma(\mathbf{k}) = \sqrt{\mathbf{k}^T\boldsymbol{\Sigma}\mathbf{k}} = \sqrt{\sum_{i=1}^{M}\sum_{j=1}^{M} k_i k_j \sigma_i \sigma_j \rho_{ij}},$$

we can see that for k_ik_j greater (less) than zero, $\sigma(\mathbf{k})$ increases (decreases) with correlations ρ_{ij}. Hence for $\mathbf{K} \geq \mathbf{0}$, increasing (decreasing) correlations will increase (decrease) $E[\sigma(\mathbf{K})]$ and thus the optimal number of draws. Increasing standard deviation σ_j of component Y_j increases $\sigma(\mathbf{k})$ if $k_ik_j\rho_{ij} \geq 0$ for $i = 1, \ldots, M$, but otherwise, the effect of changing σ_j without changing ρ_{ij} is ambiguous: the optimal number of draws could go either way as σ_j increases. A similar effect of increasing standard deviations of multivariate normal distributions appears in target-oriented situations (Tsetlin and Winkler 2007), in the case of "mixex utility" (Tsetlin and Winkler 2009), and in comparing distributions via multivariate infinite-degree stochastic dominance (Denuit et al. 2013).

Table 1 The Impact of the Distributions of Y and K on the Optimal Number of Draws

	The optimal number of draws	Result in
K becomes riskier	Increases	Proof of Theorem 3
Correlation between K_1 and K_2 increases	Increases for $M = 2$ and $\mathbf{K} \geq \mathbf{0}$	Corollary 1
Standard deviation of Y_j increases	Increases for $\mathbf{K} \geq \mathbf{0}$ and $\rho_{ij} \geq 0, i = 1, \ldots, M$	Lemma 1, discussion below
Correlation between Y_i and Y_j (ρ_{ij}) increases	Increases for $\mathbf{K} \geq \mathbf{0}$	Lemma 1, discussion below

We now consider the impact of the distribution of trade-offs $\mathbf{K}$ on the optimal number of draws. Because $\sigma(\mathbf{k})$ is convex (by the proof of Theorem 3), the optimal number of draws increases if the distribution of trade-offs becomes riskier (i.e., if multivariate zero-mean noise is added to $\mathbf{K}$; see Müller and Stoyan 2002, Theorem 3.4.2.a). Furthermore, Corollary 1 shows that, for $M = 2$ and $\mathbf{K} \geq \mathbf{0}$, one should search less if $\mathbf{K}$ changes in the direction of larger quadrant dependence (Yanagimoto and Okamoto 1969, Definition 7.1).[3]

Corollary 1. *For $M = 2$, let $\mathbf{K} \geq \mathbf{0}$ and $\mathbf{K}_0 \geq \mathbf{0}$ have the same marginal distributions and* $\Pr(K_1 \leq k_1,\ K_2 \leq k_2) \geq \Pr(K_{01} \leq k_1, K_{02} \leq k_2)$ *for all (k_1, k_2). Then one should search more if trade-offs are distributed according to $\mathbf{K}_0$—that is, $n^*(c/E[\sigma(\mathbf{K}_0)]; Z_\mathbf{Y}) \geq n^*(c/E[\sigma(\mathbf{K})]; Z_\mathbf{Y})$. This result does not extend to $M > 2$.*

Proof. For $M = 2$, observe that

$$\frac{\partial^2}{\partial k_1 \partial k_2}\sigma(\mathbf{k}) = \frac{\partial^2}{\partial k_1 \partial k_2}\sqrt{\boldsymbol{\Sigma}_{11}k_1^2 + 2\boldsymbol{\Sigma}_{12}k_1k_2 + k_2^2\boldsymbol{\Sigma}_{22}}$$
$$= -\frac{k_1k_2(\boldsymbol{\Sigma}_{11}\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{12}^2)}{(\sigma(\mathbf{k}))^3} \leq 0.$$

Therefore, $\sigma(\mathbf{k})$ corresponds to correlation-averse utility ("multivariate risk averse utility" in Richard 1975), and from Epstein and Tanny (1980, Theorem 6) it follows that $E[\sigma(\mathbf{K}_0)] \geq E[\sigma(\mathbf{K})]$. This result does not

[3] If $\mathbf{K}$ and $\mathbf{K}_0$ are discrete, then $\mathbf{K}$ has larger quadrant dependence than $\mathbf{K}_0$ if and only if $\mathbf{K}$ can be obtained from $\mathbf{K}_0$ by a finite sequence of elementary correlation-increasing transformations (Epstein and Tanny 1980, Theorem 1; their Theorem 2 extends this result to the case of continuous distributions). By definition, an elementary correlation-increasing transformation is one that shifts some probability mass from (x_1, y_1) to (x_1, y_2) and from (x_2, y_2) to (x_2, y_1), for $x_1 < x_2$, $y_1 < y_2$ (Epstein and Tanny 1980, Definition 1).

extend to $M > 2$ because in that case the sign of $(\partial^2/\partial k_1 \partial k_2)\sigma(\mathbf{k})$ can be either positive or negative. Indeed, for $M = 3$,

$$\frac{\partial^2}{\partial k_1 \partial k_2}\sigma(\mathbf{k}) = -\frac{k_1 k_2(\boldsymbol{\Sigma}_{11}\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{12}^2)}{(\sigma(\mathbf{k}))^3} + \frac{k_3^2 \boldsymbol{\Sigma}_{12}\boldsymbol{\Sigma}_{33}}{(\sigma(\mathbf{k}))^3},$$

which is negative for $\boldsymbol{\Sigma}_{12} \le 0$ but is positive for $\boldsymbol{\Sigma}_{12} > 0$ and large enough k_3. □

In Corollary 1, the correlation between components of $\mathbf{K}$ is greater than the correlation between components of $\mathbf{K}_0$. For $M = 2$, that leads to a lower value of $E[\sigma(\mathbf{K})]$ and then to a lower number of draws. This is not necessarily the case for $M > 2$, because the sign of $(\partial^2/\partial k_1 \partial k_2)\sigma(\mathbf{k})$ depends on k_3.

To gain more intuition about Theorem 3, it is useful to consider a special case in which $\boldsymbol{\Sigma}$ is the identity matrix. Then $\sigma(\mathbf{k}) = \sqrt{\mathbf{k}^T\mathbf{k}}$ is the length of $\mathbf{k}$, which is convex because the length of the sum is less than the sum of the lengths. Uncertainty about the trade-off vector $\mathbf{K}$ can then be decomposed into uncertainty about length (i.e., the value of $\sigma(\mathbf{k})$) and uncertainty about direction (i.e., the unit vector $\mathbf{k}/\sigma(\mathbf{k})$). For example, consider $M = 2$; let $\mathbf{K}^{*T}$ be either $(0.5, 0.5)$ or $(1.5, 1.5)$ with equal probability, and let $\mathbf{K}_0^T$ be either $(1.5, 0.5)$ or $(0.5, 1.5)$ with equal probability. Then there is no uncertainty about direction if uncertain trade-offs are given by $\mathbf{K}^*$ and no uncertainty about length if uncertain trade-offs are given by $\mathbf{K}_0$. According to Corollary 2 below, it is worth resolving uncertainty about trade-offs before the selection stage (i.e., being in the "trade-offs become known" scenario is preferable to being in the "trade-offs remain unknown" scenario) if trade-offs are given by $\mathbf{K}_0$ but not if they are given by $\mathbf{K}^*$. By contrast, it is advantageous to resolve uncertainty about trade-offs before the search stage (i.e., to be in the "trade-offs are known" scenario rather than the "trade-offs become known" scenario) if trade-offs are given by $\mathbf{K}^*$ but not if they are given by $\mathbf{K}_0$. Indeed, if there is no uncertainty about direction, then the ranking of the discovered alternatives does not depend on $\mathbf{k}$. Hence the action at the selection stage is independent of $\mathbf{k}$—in which case the information about $\mathbf{K}$ has no value. If there is no uncertainty about length (i.e., $\sigma(\mathbf{K})$ is the same for all realizations of $\mathbf{K}$), then information about $\mathbf{K}$ does not change the decision at the search stage as the optimal number of draws is given by $n^*(c/\sigma(\mathbf{k}); Z_{\mathbf{Y}})$. Note also that, in line with Corollary 1, the optimal number of draws when trade-offs become known is greater for $\mathbf{K}_0$ than for $\mathbf{K}^*$; the reason is that $\mathbf{K}_0$ is obtained from $\mathbf{K}^*$ by a correlation-decreasing transformation.

No generality is lost by assuming that $\boldsymbol{\Sigma}$ is the identity matrix. Indeed, the $M \times M$ covariance matrix $\boldsymbol{\Sigma}$ can be written as $\boldsymbol{\Sigma} = \mathbf{A}^T\mathbf{A}$, where $\mathbf{A}$ is a $p \times M$ matrix with $\text{rank}(\boldsymbol{\Sigma}) = p$. Then there exists a p-variate spherically distributed $\mathbf{Y}'$ with $E(Y_j') = 0$ and $\text{Var}(Y_j') = 1$, $j = 1, \ldots, p$, such that $\mathbf{Y} = \boldsymbol{\mu}_{\mathbf{Y}} + \mathbf{A}^T\mathbf{Y}'$ (Fang et al. 1990, Definition 2.2).[4] For example, if $\mathbf{Y}$ is multivariate normal, then $\mathbf{Y}'$ consists of independent standard normal variables. Note that $Z_{\mathbf{Y}}$ (defined in Definition 1 and Lemma 1) has the same distribution as Y_j', $j = 1, \ldots, p$, and therefore $\omega(n; Z_{\mathbf{Y}}) = \omega(n; Y_1')$. Let $\mathbf{K}' = \mathbf{A}\mathbf{K}$. Then $\sigma(\mathbf{k}) = \sqrt{\mathbf{k}^T\boldsymbol{\Sigma}\mathbf{k}} = \sqrt{\mathbf{k}'^T\mathbf{k}'}$. In other words, the parallel search setting in which alternatives are drawn from an M-variate elliptical distribution with covariance matrix $\boldsymbol{\Sigma} = \mathbf{A}^T\mathbf{A}$ is equivalent to a setting where alternatives are drawn from a p-variate spherical distribution ($p \le M$) and where the distribution of trade-offs is adjusted as $\mathbf{K}' = \mathbf{A}\mathbf{K}$. Returning to the original trade-offs $\mathbf{k}$, we can say that uncertainty about length is uncertainty about $\sigma(\mathbf{k}) = \sqrt{\mathbf{k}^T\boldsymbol{\Sigma}\mathbf{k}}$ and that uncertainty about direction is uncertainty about $\mathbf{A}\mathbf{k}/\sigma(\mathbf{k})$.

Decomposing uncertainty about trade-offs into uncertainty about length and uncertainty about direction is useful for considering the value of information (i.e., how much the decision maker would benefit if trade-offs were resolved earlier). As Corollary 2 shows, it is beneficial to know the trade-offs before (rather than after) the selection stage if there is uncertainty about direction, and it is beneficial to know the trade-offs before (rather than after) the search stage if there is uncertainty about length.

Corollary 2. (i) *If there is no uncertainty about direction, then the expected payoff and the optimal number of draws are the same in the "trade-offs remain unknown" and "trade-offs become known" scenarios.*

(ii) *If there is no uncertainty about length, then the expected payoff and the optimal number of draws are the same in the "trade-offs become known" and "trade-offs are known" scenarios.*

[4] A spherical distribution is a symmetric elliptical distribution in which all components are uncorrelated.

PROOF. Consider the setting in Lemma 1, and let $\mathbf{\Sigma} = \mathbf{A}^T\mathbf{A}$.

(i) If there is no uncertainty about direction, then there exists a unit vector $\mathbf{s}$ such that $\mathbf{AK} = \mathbf{s}\sigma(\mathbf{K})$, and so $\sigma(\boldsymbol{\mu}_\mathbf{K}) = E[\sigma(\mathbf{K})]$. By Lemma 1, the payoffs in the "trade-offs remain unknown" and "trade-offs become known" scenarios are identical.

(ii) If there is no uncertainty about length, then there exists a value b such that $\sigma(\mathbf{K})$ equals b almost surely. If trade-offs are known, then the optimal number of draws for a particular realization of $\mathbf{k}$ is given by $n^*(c/\sigma(\mathbf{k}); Z_\mathbf{Y}) = n^*(c/b; Z_\mathbf{Y})$. Note that $E[\sigma(\mathbf{K})] = b$, and therefore the optimal decision at the search stage is the same in both scenarios. The decision at the selection stage is also the same, given that in each scenario the trade-offs are known by the selection stage. □

4.2. Examples Where It Is Optimal to Search Less If Trade-offs Become Known

Here, we present two examples that illustrate the difficulty of generalizing §4.1's results to nonelliptical distributions. Example 1 is more straightforward and uses a discrete three-point distribution of $\mathbf{Y}$ to demonstrate that the resolution of uncertainty about $\mathbf{K}$ before the selection stage may either decrease or increase the optimal number of draws as compared with the "trade-offs remain unknown" scenario. Example 2 demonstrates the same with $\mathbf{Y}$ being independently uniform.

EXAMPLE 1. Let

$$\mathbf{Y}^T = (Y_1, Y_2) = \begin{cases} (1,0) & \text{with probability } p, \\ (20,-10) & \text{with probability } (1-p)/2, \\ (-20,10) & \text{with probability } (1-p)/2; \end{cases}$$

$$\mathbf{K}^T = (K_1, K_2) = \begin{cases} (1,1) & \text{with probability } 0.5, \\ (1,3) & \text{with probability } 0.5; \end{cases}$$

$$v(\mathbf{y}, \mathbf{k}) = y_1k_1 + y_2k_2; \quad c = 0.01.$$

Set $p = 0.05$. If $\mathbf{K}$ remains unknown, then the optimal number of draws is 31, but if $\mathbf{K}$ becomes known, then the optimal number of draws is 10.

Set $p = 0.5$. Now if $\mathbf{K}$ remains unknown, then the optimal number of draws is 5, but if $\mathbf{K}$ becomes known, then the optimal number of draws is 18. (See the appendix for details.) The intuition behind this result is that if $\mathbf{K}$ remains unknown, then the alternative $(1,0)$ is preferable because it gives the highest expected payoff. If $\mathbf{K}$ becomes known, then one wants to get two alternatives, $(20,-10)$ and $(-20,10)$, that yield the highest respective payoffs for $k_2 = 1$ and $k_2 = 3$. If p is small (e.g., $p = 0.05$), then when $\mathbf{K}$ remains unknown, many draws are required (at a small cost of $c = 0.01$ each) to discover one alternative of $(1,0)$. However, if $\mathbf{K}$ becomes known, then even a relatively small number of draws would practically guarantee that there will be one $(20,-10)$ and one $(-20,10)$ alternative. If p is not small (e.g., $p = 0.5$), then the opposite holds: there is no need to take many draws when $\mathbf{K}$ remains unknown, since the chances of getting $(1,0)$ are quite high; but one must search more to discover one $(20,-10)$ and one $(-20,10)$ alternative. □

Example 1 shows that (a) the resolution of uncertainty about $\mathbf{K}$ at the selection stage does affect the optimal number of draws at the search stage and (b) if $\mathbf{Y}$ is nonelliptical then that number can be either greater or less than the optimal number of draws when $\mathbf{K}$ remains unknown. However, Example 1 involves a discrete distribution of $\mathbf{Y}$, and its logic is based on obtaining specific alternatives if trade-offs either remain unknown or become known. In our next example, $\mathbf{Y}$ is continuous.

EXAMPLE 2. Consider $M = 2$, and let Y_1 and Y_2 be independent and uniformly distributed on $[0,1]$. Let

$$\mathbf{K}^T = (K_1, K_2) = \begin{cases} (1,0) & \text{with probability } 0.5, \\ (0,1) & \text{with probability } 0.5; \end{cases}$$

and $v(\mathbf{y}, \mathbf{k}) = y_1k_1 + y_2k_2$.

If trade-offs become known, then

$$E\Big[\max_{i=1,\dots,n} v(\mathbf{Y}_i, \mathbf{K})\Big] = 0.5E\Big(\max_{i=1,\dots,n} Y_{1i}\Big) + 0.5E\Big(\max_{i=1,\dots,n} Y_{2i}\Big) = \omega(n; Z_U),$$

where Z_U is uniformly distributed on $[0,1]$. Therefore, search with trade-offs becoming known is equivalent to univariate search from a uniform distribution on $[0,1]$. The optimal number of draws is $n^*(c; Z_U)$, as defined in (1), with $\omega(n; Z_U) = n/(n+1)$.

If trade-offs remain unknown, then the distribution of $Z_T = E_{\mathbf{K}} v(\mathbf{Y}, \mathbf{K}) = 0.5(Y_1 + Y_2)$ is symmetric triangular on $[0, 1]$. Then the optimal number of draws is $n^*(c; Z_T)$ with $\omega(n; Z_T) = \int_0^{0.5} n2^{1-n}x^{2n}\,dx + \int_{0.5}^1 nx(2-x)(4x-x^2-2)^{n-1}2^{1-n}\,dx$.

If it were always optimal to search more in the "trade-offs become known" scenario, then $n^*(c; Z_U) \geq n^*(c; Z_T)$ for all c. By (1), this is equivalent to $\omega(n+1; Z_U) - \omega(n; Z_U) \geq \omega(n+1; Z_T) - \omega(n; Z_T)$ for all n. Calculation reveals that this inequality holds for $n \leq 7$ but not for $n > 7$. As a result, $n^*(c; Z_U) \geq n^*(c; Z_T)$ for $c \geq 0.012$ and $n^*(c; Z_U) \leq n^*(c; Z_T)$ for $c < 0.012$, with strict inequalities for some c.

In this example, although the standard deviation of $v(\mathbf{Y}, \mathbf{k}) = k_1 Y_1 + k_2 Y_2$ is convex in $\mathbf{k}$, the shape of this distribution depends on $\mathbf{k}$. It is uniform for $k_1 = 1$ and $k_2 = 0$ but is triangular for $k_1 = 0.5$ and $k_2 = 0.5$ (which corresponds to the scenario where trade-offs remain unknown). For a uniform distribution, the largest-order statistic approaches the upper bound fairly quickly, so that additional gains of increasing n are small. Reaching the upper bound for a triangular distribution is more difficult, and this leads to violation of Theorem 3 for small c. □

In Example 2, $\mathbf{Y}$ is uniform over a square. We remark that if $\mathbf{Y}$ were uniform over a circle (or an ellipse) then it would be spherically distributed, and by Theorem 3, the optimal number of draws would be greater when trade-offs become known, because in this case the shape of the distribution of $\mathbf{k}^T\mathbf{Y}$ would not depend on $\mathbf{k}$.

Overall, for Theorem 3 to hold, the distribution of $\mathbf{Y}$ should not be too different from elliptical (alternatively, the shape of $\mathbf{k}^T\mathbf{Y}$ should not change too much with $\mathbf{k}$). For example, consider $M = 2$, $\mathbf{Y}$ independent standard normal, and $v(\mathbf{y}, \mathbf{k}) = k_1(1-e^{-ry_1})/r + k_2(1-e^{-ry_2})/r$ for $r > 0$. Here, Theorem 3 holds if the joint distribution of $((1-e^{-ry_1})/r, (1-e^{-ry_2})/r)$ is close to elliptical, which is the case if r is small enough. As r increases, the distribution of $(1-e^{-ry_1})/r$ moves further from normal—leading to violation of Theorem 3. As numerical calculation shows, if $r = 1$, $c = 0.045$, and $\mathbf{K}^T$ is either $(1, 0)$ or $(0, 1)$ with equal probability (as in Example 2), then the optimal number of draws is five if trade-offs become known and six if trade-offs remain unknown. The intuition behind this result is similar to that behind Example 2: if trade-offs become known, then the shape of $v(\mathbf{Y}, \mathbf{k})$ is negative lognormal, but if trade-offs remain unknown, then the shape of $E_{\mathbf{K}} v(\mathbf{Y}, \mathbf{K})$ is the sum of two negative lognormals, which has a fatter right tail. If r is small, then there is little difference between lognormal and normal distributions, and the conclusion of Theorem 3 holds for virtually all values of c.

4.3. Illustration: Bivariate Normal Distribution of Alternatives

Let $\mathbf{Y} = (Y_1, Y_2)^T$ be bivariate normal with respective standard deviations σ_1, σ_2 and correlation ρ. We consider the case of one uncertain trade-off $\mathbf{K}^T = (1, K_2)$ and assume without loss of generality that $E(K_2) = 1$. Then $v(\mathbf{y}, \mathbf{k}) = y_1 + k_2 y_2$.

In a setting with $K_1 \equiv 1$, the first attribute is in the same units as search cost. In the case of monetary search cost, y_1 can be price, cost, or profit. The second attribute is nonmonetary—for example, capacity, reliability, quality, or supplier sustainability. Then k_2 is the weight of the second attribute, or the trade-off between y_2 and y_1 (e.g., between quality and price), and it might be not known at the search stage.

If the trade-off remains unknown at the selection stage, then the optimal number of draws is $n^*(c/\sigma(\boldsymbol{\mu}_{\mathbf{K}}); Z_N)$ with $\sigma(\boldsymbol{\mu}_{\mathbf{K}}) = \sqrt{\sigma_1^2 + 2\sigma_1\sigma_2\rho + \sigma_2^2}$. If the trade-off becomes known by the selection stage, then, by Lemma 1, the optimal number of draws is $n^*(c/E[\sigma(\mathbf{K})]; Z_N)$ with

$$E[\sigma(\mathbf{K})] = E\left(\sqrt{\sigma_1^2 + 2\sigma_1\sigma_2\rho K_2 + \sigma_2^2 K_2^2}\right).$$

Silver (1987, Table 1) presents the values of $n^*(c; Z_N)$ and uses them to estimate the number of bids. An interesting observation from Silver is that, numerically, $n^*(c; Z_N)$ is close to $0.5/c$—so if the values of the alternatives are normally distributed with standard deviation σ, then one should spend about 0.5σ on search.[5]

Figure 2 plots the optimal number of draws in these two scenarios as a function of ρ, where $K_2 = 1 \pm 0.8$

[5] As the search cost c decreases, the total optimal search cost also decreases, albeit rather slowly. For example, if $c = 0.002$, then $n^*(0.002; Z_N) = 169$, and the total optimal search cost is 0.338. This value is not that far from 0.5 when one considers the extremely small search cost of $c = 0.002$.

Figure 2 **Optimal Number of Draws as a Function of Correlation ρ for $K_2 = 1.0 \pm 0.8$ (with Equal Probability)**

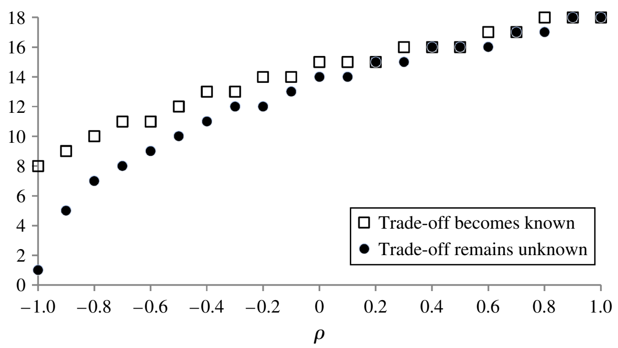

Note. The optimal number of draws is given along the vertical axis; $c = 0.05$ and $\sigma_1 = \sigma_2 = 1$.

(with equal probability), $c = 0.05$, and $\sigma_1 = \sigma_2 = 1$. Confirming Theorem 3, the graph shows that one should search more if the trade-offs become known. This effect is stronger for negative correlations (e.g., for $\rho = -0.9$, the optimal number of draws is five when the trade-offs remain unknown but is nine when the trade-offs become known) and vanishes for positive correlations. This dynamic is consistent with Corollary 2: uncertainty about direction decreases with correlation and goes to zero at $\rho = 1$, in which case $E[\sigma(\mathbf{K})] = \sigma(\boldsymbol{\mu}_{\mathbf{K}})$.

Figure 3 plots the difference between the expected payoff with the optimal number of draws and the expected payoff with a single draw for the three different scenarios (of when the trade-offs become known) described in Figure 1. (This difference does not depend on $\boldsymbol{\mu}_{\mathbf{Y}}$.) As expected, the earlier the trade-off is known, the better. We can also see that, at $\rho = -1$, the expected payoffs are the same for the "trade-offs are known" and "trade-offs become known" scenarios; in other words, there is no benefit in resolving uncertainty about the trade-off before the search stage as opposed to resolving it before the selection stage. This result is consistent with Corollary 2 because $\sigma(\mathbf{K}) \equiv 0.8$ at $\rho = -1$, eliminating all uncertainty about length. At $\rho = 1$ the situation is different. Now the payoffs are the same for the "trade-offs become known" and "trade-offs remain unknown" scenarios, because in this case there is no uncertainty about direction.

Figure 3 **Difference Between Expected Payoff with the Optimal Number of Draws and Expected Payoff with a Single Draw as a Function of Correlation ρ, for Three Different Information Scenarios**

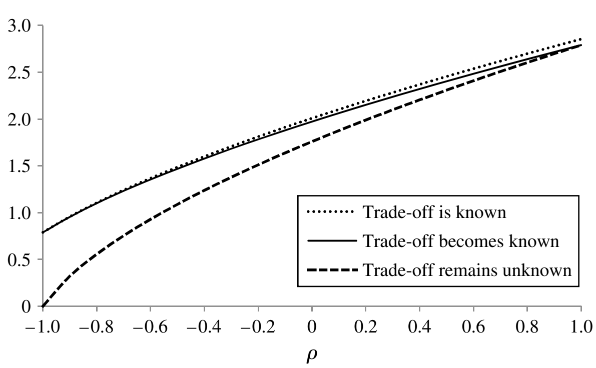

Notes. The difference in expected payoffs is given along the vertical axis; $K_2 = 1.0 \pm 0.8$ (with equal probability), $c = 0.05$, and $\sigma_1 = \sigma_2 = 1$.

Alternatives are typically characterized by multiple attributes, and in §2 we cite works that feature multiple attributes in the vendor selection context. For purposes of illustration we have focused on the $M = 2$ case with a single uncertain trade-off. We now discuss a realistic situation in which the foregoing analysis could prove useful.

A corporation needs a new floor space and must decide how many vendors it should invite to bid on the construction tender. According to observed practice, the corporation will evaluate the bids in terms of two attributes: the price to construct the building and the speed (in number of months) required for completion (Palaneeswaran and Kumaraswamy 2000).

Meanwhile, the corporation will rent temporary space until the permanent building is constructed. The monthly rental cost determines the price–speed trade-off that will be used for selecting the winner of the tender. At the time of deciding how many bidders to invite, the rental cost of the temporary space is unknown. Yet the tender process for construction lasts several months (Elfving et al. 2005), so the corporation will learn the rental cost before selecting a construction vendor.

The corporation may follow one of two different approaches when deciding how many vendors should be asked to bid. If the corporation accounts for the resolution of rental cost uncertainty, then its decision making follows the "trade-offs become known" scenario, and the decision is based on the probability distribution of the rental cost; if the corporation does not account for the resolution of rental cost uncertainty (or entirely ignores this uncertainty), then its decision making follows the "trade-offs remain unknown" scenario, and the decision is based on the expected rental cost. (The "trade-offs are known" scenario would occur if the corporation had an option to invite the bidders after the rental cost becomes known.)

For illustrative purposes, assume that the cost c of including each vendor is \$50,000, that the monthly rental cost k_2 is either \$0.2 or \$1.8 million (with equal probability), and that the attribute values (Y_1, Y_2) for price and speed are jointly normally distributed with $\sigma_1 = \$1$ million and $\sigma_2 = 1$ month.[6] These numbers correspond to Figures 2 and 3, where the latter's vertical axis is denominated in millions of dollars.

Figure 2 shows that the corporation can safely ignore the uncertainty about the price–speed trade-off if it expects ρ to be positive. That would be the case if different vendors use different technologies, some fitting better to the corporation's requirements and some fitting worse. Then both price and speed are either good or bad. However, we would rather expect a negative correlation, as when a vendor offering rapid construction demands a higher price for speediness. This scenario corresponds to negative ρ in Figure 2, in which case the corporation should approach more vendors than would be recommended by a model that uses only the expected rental cost (and thus ignores uncertainty about that cost at the search stage). As mentioned previously, for $\rho = -0.9$, the optimal number of vendors is five if the trade-off (rental cost) remains unknown and nine if the trade-off becomes known.

Figure 3 plots the benefit of soliciting bids from the optimal number of vendors relative to approaching just a single vendor. When $\rho = -0.9$, the payoff in the "trade-offs become known" scenario is \$0.63 million higher than that in the "trade-offs remain unknown" scenario. In this case, failing to account for resolution of the rental cost uncertainty leads not only to underinvesting in search but also to underestimating the benefits of search. The latter is costly because it may lead the corporation to forgo the tender altogether.

[6] We assume a bivariate normal distribution for illustrative purposes. As follows from §4.1, all results would remain qualitatively the same for any other bivariate elliptical distribution.

5. Search with Outside Alternatives and Sequential Search

So far we have assumed that the decision maker must make at least one draw (i.e., $n \geq 1$), that there are no previously discovered (and still available) alternatives, and that there is no option to abandon. We now consider the case where outside alternatives are present. This setting also allows us to discuss when one should stop searching and the implications for sequential search.

Consider the setting of Theorem 2, and assume that the decision maker has already discovered $m \geq 1$ alternatives $\mathbf{y}_1^*, \ldots, \mathbf{y}_m^*$. We call them "outside alternatives." (An option to abandon would be an alternative that yields zero payoff regardless of the value of $\mathbf{K}$.) Set $\Theta = \{\mathbf{y}_1^*, \ldots, \mathbf{y}_m^*\}$. Then, for $n \geq 0$ draws, the expected payoff is given by

$$\pi(c, n; \mathbf{Y}, \mathbf{K} \mid \Theta) = -nc + E\Big[\max_{i=1,\ldots,n,\, l=1,\ldots,m} \{v(\mathbf{Y}_i, \mathbf{K}), v(\mathbf{y}_l^*, \mathbf{K})\}\Big]. \quad (3)$$

We can adapt the proof of Theorem 2 to show that the expected payoff (3) is concave in n. However, computing this payoff becomes more tedious because the optimal n could depend on all available alternatives $\mathbf{y}_1^*, \ldots, \mathbf{y}_m^*$. That situation arises whenever, for each l ($l = 1, \ldots, m$), there is a positive probability that $\max\{v(\mathbf{y}_1^*, \mathbf{K}), \ldots, v(\mathbf{y}_m^*, \mathbf{K})\} = v(\mathbf{y}_l^*, \mathbf{K})$—in other words, if none of the alternatives $\mathbf{y}_1^*, \ldots, \mathbf{y}_m^*$ is dominated by others. However, if trade-offs remain unknown at the time of selection, then the decision maker needs to retain only one of the m alternatives, the one that maximizes $E_{\mathbf{K}}[v(\mathbf{y}_l^*, \mathbf{K})]$.

Example 3. Consider the case $M = 2$, and suppose that $\mathbf{Y}$ is spherically distributed with uniform density over the circle centered at $(0, 0)$ with radius 1. Let $\mathbf{K}^T$ be either $(1, 0)$ or $(0, 1)$ with equal probability, $v(\mathbf{y}, \mathbf{k}) = y_1 k_1 + y_2 k_2$, and let $m = 1$, $\mathbf{y}_1^* = (1, 0)^T$. If trade-offs become known, then

$$E\Big[\max_{i=1,\ldots,n,\, l=1,\ldots,m} \{v(\mathbf{Y}_i, \mathbf{K}), v(\mathbf{y}_l^*, \mathbf{K})\}\Big] = 0.5 + 0.5E\Big(\max_{i=1,\ldots,n} \max(Y_{2i}, 0)\Big).$$

If trade-offs remain unknown, then the distribution of $E_{\mathbf{K}} v(\mathbf{Y}, \mathbf{K}) = 0.5(Y_1 + Y_2)$ is the same as of $(1/\sqrt{2})Y_2$, and

$$E\Big[\max_{i=1,\dots,n,\, l=1,\dots,m}\{E_{\mathbf{K}} v(\mathbf{Y}_i, \mathbf{K}), E_{\mathbf{K}} v(\mathbf{y}_l^*, \mathbf{K})\}\Big] = \frac{1}{\sqrt{2}} E\bigg(\max_{i=1,\dots,n} \max\Big(Y_{2i}, \frac{1}{\sqrt{2}}\Big)\bigg).$$

Calculation shows that, for $c = 0.0025$, the optimal number of draws is 22 in both scenarios; for $c > 0.0025$, it is greater or equal if trade-offs become known, and for $c < 0.0025$, it is greater or equal if trade-offs remain unknown. For example, for $c = 0.0005$, the optimal number of draws is 57 if trade-offs become known and 70 if trade-offs remain unknown. □

As this example shows, Theorem 3 does not generalize to the case where there are outside alternatives. As in Example 2, the shape of the distribution of $[\max_{i=1,\dots,n,\, l=1,\dots,m}\{v(\mathbf{Y}_i, \mathbf{k}), v(\mathbf{y}_l^*, \mathbf{k})\}]$ depends on $\mathbf{k}$, and that leads to violation of Theorem 3.

We now consider the decision to stop searching if m alternatives are already available. Taking zero draws is equivalent to stopping the search; for that stopping decision, the decision maker must compare the expected payoff with no additional draws ($n = 0$) and the expected payoff with at least one additional draw ($n \geq 1$). By the concavity of (3) in n, one should stop searching if $\pi(c, 0; \mathbf{Y}, \mathbf{K} \mid \Theta) > \pi(c, 1; \mathbf{Y}, \mathbf{K} \mid \Theta)$; that is, one need only compare the expected payoffs with $n = 0$ and $n = 1$. Let us compare the stopping decisions in two scenarios—one where trade-offs become known and one where they remain unknown—for the same set of outside alternatives $\mathbf{y}_1^*, \dots, \mathbf{y}_m^*$.

Theorem 4. *Suppose that exactly one outside alternative is available (i.e., $m = 1$). If it is optimal to continue searching when trade-offs remain unknown, then it is optimal to continue searching when trade-offs become known.*

Proof. If one more alternative is explored, then the expected payoff is

$$-c + E_{\mathbf{Y}}[\max\{E_{\mathbf{K}}[v(\mathbf{Y}, \mathbf{K})], E_{\mathbf{K}}[v(\mathbf{y}_1^*, \mathbf{K})]\}] \quad \text{if trade-offs remain unknown}$$

or

$$-c + E_{\mathbf{K}}[E_{\mathbf{Y}}(\max\{v(\mathbf{Y}, \mathbf{K}), v(\mathbf{y}_1^*, \mathbf{K})\} \mid \mathbf{K})] \quad \text{if trade-offs become known.}$$

If the search is stopped, then the expected payoff in either scenario is $E_{\mathbf{K}}[v(\mathbf{y}_1^*, \mathbf{K})]$. When trade-offs remain unknown, it is optimal to continue searching if and only if

$$-c + E_{\mathbf{Y}}\big[\max\{E_{\mathbf{K}}[v(\mathbf{Y}, \mathbf{K})], E_{\mathbf{K}}[v(\mathbf{y}_1^*, \mathbf{K})]\}\big] > E_{\mathbf{K}}\big[v(\mathbf{y}_1^*, \mathbf{K})\big];$$

when trade-offs become known, it is optimal to continue searching if and only if

$$-c + E_{\mathbf{K}}\big[E_{\mathbf{Y}}\big(\max\{v(\mathbf{Y}, \mathbf{K}), v(\mathbf{y}_1^*, \mathbf{K})\} \mid \mathbf{K}\big)\big] > E_{\mathbf{K}}\big[v(\mathbf{y}_1^*, \mathbf{K})\big].$$

Theorem 4 now follows from the inequality

$$E_{\mathbf{Y}}\big[\max\{E_{\mathbf{K}}[v(\mathbf{Y}, \mathbf{K})], E_{\mathbf{K}}[v(\mathbf{y}_1^*, \mathbf{K})]\}\big] \leq E_{\mathbf{K}}\big[E_{\mathbf{Y}}\big(\max\{v(\mathbf{Y}, \mathbf{K}), v(\mathbf{y}_1^*, \mathbf{K})\} \mid \mathbf{K}\big)\big]. \quad \square$$

Theorem 4 exploits the fact that if one stops searching (i.e., $n = 0$) then, for $m = 1$, the expected payoff is the same in both scenarios; however, if one more alternative is drawn, then the expected payoff is greater when trade-offs become known. For $m \geq 2$, the expected payoff with $n = 0$ is also greater in the scenario where trade-offs become known. As a result, Theorem 4 does not extend to $m \geq 2$; it could be that (for some choice of $\mathbf{y}_1^*, \dots, \mathbf{y}_m^*$) the searcher will stop if trade-offs become known at the selection stage yet will continue searching if those trade-offs remain unknown, as the following example illustrates.

Example 4. Consider the same setting as in Example 3, and let $m = 2$, $\mathbf{y}_1^* = (1, 0)^T$, and $\mathbf{y}_2^* = (0, 1)^T$. If trade-offs become known, then the searcher will stop because the highest possible payoff has already been achieved. If trade-offs remain unknown, then the best possible alternative is $(1/\sqrt{2}, 1/\sqrt{2})^T$. So for a low enough search cost, it is optimal to continue searching to get closer to $(1/\sqrt{2}, 1/\sqrt{2})^T$. □

We now discuss sequential search before trade-offs are known. The simplest setting is the same as in Figure 1 except for change to the search stage: instead of choosing the number of draws (the only decision variable in parallel search), the decision maker performs a sequential search, drawing alternatives one by one and paying cost c for each draw. Once the searcher judges that expected improvement from an additional draw is less than cost c, the searcher stops, and the search stage is completed. Search with recall corresponds to a situation where, at the selection stage, the decision maker can choose any of the discovered alternatives. Search without recall corresponds to a situation where, once an alternative is drawn, it must be either accepted (in which case the search stage ends) or abandoned forever.[7]

In a standard sequential search model (e.g., Lippman and McCall 1976) with independent payoffs drawn from the same distribution, the stopping rule is characterized by a reservation price. Then infinite-horizon searches with and without recall are equivalent because the decision maker never needs to keep an alternative that is below the reservation price and stops searching if the discovered alternative is above the reservation price. The same holds for the sequential multivariate search if trade-offs remain unknown: the decision maker will stop when an alternative $\mathbf{y}$ with $E_{\mathbf{K}}[v(\mathbf{y}, \mathbf{K})]$ above the reservation price is discovered. However, in the "trade-offs become known" scenario, the decision maker would like to keep multiple discovered alternatives (all that are not dominated), and the decision to stop depends, potentially, on all previously discovered alternatives. Also note that, for the search without recall in the "trade-offs become known" scenario, the decision maker accepts the first alternative with $E_{\mathbf{K}}[v(\mathbf{y}, \mathbf{K})]$ above the reservation price, and thus the stopping decision is equivalent to the one where trade-offs remain unknown. To conclude, the stopping decision in the "trade-offs become known" scenario is different from the one in a univariate search, and the settings with and without recall are different.

Another difference concerns draws from different distributions. Weitzman (1979) considers a setting in which the searcher has several closed boxes (i.e., different univariate distributions) to explore. Each box is characterized by a reservation price, and the search strategy is to open the box with the highest reservation price. If trade-offs become known, then the search strategy is not that simple and elegant. Each box does not have a single reservation value because the attractiveness of opening a particular box depends on the combination of alternatives discovered previously. For example, suppose the searcher can take any number of draws (sequentially, with a fixed cost per draw) from two different distributions. In a univariate setting, one would never use both distributions: all draws will be taken from the distribution with the highest reservation price. If trade-offs become known, then the searcher might switch between the distributions, depending on previously discovered alternatives.

[7] Such a sequential search model is applicable in a situation where all search activity must be completed before some externally set deadline (e.g., a board meeting at which trade-offs will become known and one of the discovered alternatives will be chosen), and it assumes that no information about trade-offs is revealed during the search stage. As mentioned at the end of §2, our model is more relevant to the parallel search setting.

6. Summary and Discussion

Almost all existing models of search consider single-attribute alternatives, even though real-world alternatives are usually characterized by multiple attributes. The payoff from a multiattribute alternative is determined by the trade-offs among—or weights attached to—its attributes. A decision maker explores the alternatives at the search stage and then, at the selection stage, chooses one of them.

If the trade-offs determining final payoffs are known at the search stage, then a search for multiattribute alternatives is equivalent to a univariate search in which each alternative is characterized by its payoff. If trade-offs are not known at the search stage and remain unknown at the selection stage, then the multiattribute setting is again equivalent to a univariate search; in this scenario, each alternative is characterized by its expected payoff, and we can assume that each uncertain trade-off is equal to its expected value. In both cases, trade-offs that will be used at the selection stage are known at the search stage. However, when the uncertainty about trade-offs is resolved after the search stage yet before the selection stage, the multiattribute search setting differs from a univariate one.

We consider parallel search in §4. If trade-offs become known, then, in general, the optimal number of draws (i.e., the optimal decision at the search stage) could be either greater or less than the optimal number of draws if trade-offs remain unknown (Examples 1 and 2). In a setting where alternatives are drawn from a multivariate elliptical distribution, the optimal number of draws when the trade-offs become known cannot be less than the optimal number of draws when trade-offs remain unknown (Theorem 3). The family of elliptical distributions, which is a large one, includes the multivariate normal distribution. The emerging prescription is that one should search more if some information about trade-offs is likely to arrive before the selection stage. A realistic application is discussed in §4.3.

If the uncertainty about trade-offs is resolved by the selection stage, then losses from nonoptimal search—as when one wrongly assumes that trade-offs remain unknown at the selection stage (which is equivalent to ignoring the uncertainty about trade-offs)—can be substantial. As illustrated in §4.3 for a two-attribute example, these losses are higher when attributes are negatively correlated (which could reasonably be expected in many situations) and also when trade-offs are negatively correlated (Corollary 1).

Naturally, it is better to resolve uncertainty about trade-offs sooner rather than later; thus, the expected payoff in the scenario where trade-offs remain unknown is smaller than the one where trade-offs become known, which in turn is smaller than the payoff where trade-offs are known at the outset. However, these benefits depend on the nature of the uncertainty (Corollary 2). Uncertainty about trade-offs can be decomposed into uncertainty about direction and uncertainty about length of the adjusted trade-off vector. Resolving uncertainty about trade-offs at the selection stage matters to the extent that it might change the best alternative, which is captured by uncertainty about direction (an uncertainty that, as Theorem 3 shows, in turn affects the search stage decision). Resolving uncertainty about trade-offs at the search stage matters to the extent that it might change the optimal number of draws, which is captured by uncertainty about length.

In §5 we extend our setting to the case in which some alternatives (e.g., an option to abandon) are present before search starts, and we also consider the decision of when to stop searching. An important difference here from univariate search (and thus from the setting where trade-offs remain unknown) is that the decision maker must retain multiple alternatives because, depending on information about trade-offs at the selection stage, different alternatives might turn out to be optimal. When only one outside alternative is available, it is better to search more if trade-offs become known at the selection stage (Theorem 4). We would think that the case of a single outside alternative does occur often in practice. After all, a decision maker who perceives the alternatives to be univariate has no reason to retain multiple alternatives—not even if some had been discovered previously. Another difference from univariate search is that a searcher who can draw from different distributions might prefer switching among them in the "trade-offs become known" scenario, in response to previously discovered alternatives. The reason is that, unlike the case of univariate search, here each distribution cannot be characterized by a single reservation price.

Our conclusion is that both parallel (§4) and sequential (§5) search with multiattribute alternatives—and with trade-offs becoming known at the selection stage—are quite different from univariate search. We expect that this conclusion and our model will be relevant also in richer settings. In Morgan and Manning (1985), the searcher decides how many alternatives to draw in each period. Chade and Smith (2006) consider the portfolio choice problem in which the decision maker determines how many draws to take simultaneously and also which distributions to explore (as in the literature on directed search, where participants in the labor market identify the jobs for which they should apply). Smith and Ulu (2012) consider technology adoption with uncertain costs and qualities. In all these settings it would be reasonable to view each alternative as having multiple attributes, and it would also be reasonable to expect that some additional information about trade-offs will become available over time (i.e., before the selection decision is made).

Our approach and results could be useful in modeling innovation tournaments and delegated search (i.e., search by agents on behalf of the principal). Chao and Erat (2012) discuss a setting in which the contestants are engaged in search activity but the principal will (with some probability) evaluate the quality of their solution as zero. Such a setting can be viewed as an extreme case of uncertainty about trade-offs, and more detailed modeling of this uncertainty could well yield extra insights. Chao et al. (2014) consider a stage-gate process whereby the agent searches and the principal is responsible for deciding whether (and when) the project should be stopped. A natural extension of that model would be to consider multiattribute projects, for which one could reasonably assume that the agent has some uncertainty about the trade-offs to be used by the principal. The model could also incorporate some gradual learning about the trade-offs.

It might be often overlooked that uncertainty about trade-offs will be reduced at the selection stage. This oversight could be due to the commonly observed underappreciation of uncertainty: a decision maker who does not perceive trade-offs to be uncertain at the search stage will, of course, be unaware that this uncertainty may be reduced at the selection stage. Furthermore, decision makers often omit important objectives (Bond et al. 2008) and make typical mistakes when assessing trade-offs (Keeney 2002); similar mistakes are likely to occur when assessing uncertainty about those trade-offs. Accounting both for uncertain trade-offs and for multiattribute alternatives places more of a cognitive burden on the decision maker, who must estimate not only the joint multivariate distribution of alternatives' attributes (as opposed to the univariate distribution of alternatives' payoffs) but also the reduction in trade-off uncertainty at the selection stage. Even so, the resulting solution need not be overly complicated—for instance, it is fairly tractable in the case of a multivariate elliptical distribution of attributes. More importantly, such analysis leads to a better search decision (i.e., one with a higher expected payoff) and also enables us to assess the value of information about uncertain trade-offs.

Acknowledgments

The authors are grateful to Yakov Bart, Neil Bearden, Tian Heong Chan, Stephen Chick, Sameer Hasija, Nitish Jain, Kevin McCardle, Claudio Mezzetti, Serguei Netessine, Dana Popescu, Ling Yue, and Sasa Zorc for helpful comments and suggestions. The associate editor and the referees provided very useful comments. The authors are especially grateful to Jim Smith for the reference to elliptical distributions.

Appendix. Derivation of Example 1

Let $\mathbf{y}_1 = (1, 0)^T$, $\mathbf{y}_2 = (20, -10)^T$, and $\mathbf{y}_3 = (-20, 10)^T$.

First consider the case where $\mathbf{K}$ remains unknown. Then $E_{\mathbf{K}}[v(\mathbf{y}_1, \mathbf{K})] = 1$, $E_{\mathbf{K}}[v(\mathbf{y}_2, \mathbf{K})] = 0$, $E_{\mathbf{K}}[v(\mathbf{y}_3, \mathbf{K})] = 0$,

$$X = E_{\mathbf{K}}[v(\mathbf{Y}, \mathbf{K})] = \begin{cases} 1 & \text{with probability } p, \\ 0 & \text{with probability } 1-p, \end{cases}$$

and the expected payoff after n draws is given by $\pi(c, n; X) = -nc + 1 - (1-p)^n$. In accordance with Theorem 1, this value is maximized at $n^*(c; X) = \max(1, \lfloor \ln(c/p)/\ln(1-p) \rfloor)$; here, $\lfloor \cdot \rfloor$ signifies rounding down to the nearest integer.

Now consider the case where $\mathbf{K}$ becomes known before choosing one of the n alternatives. For both possible realizations of $\mathbf{K}$, we have

$$\max_{i=1,\ldots,n} v(\mathbf{Y}_i, \mathbf{K}) = \begin{cases} -10 & \text{with probability } ((1-p)/2)^n, \\ 1 & \text{with probability } (p + (1-p)/2)^n - ((1-p)/2)^n, \\ 10 & \text{with probability } 1 - (p + (1-p)/2)^n. \end{cases}$$

The expected payoff (2) is

$$\pi(c, n; \mathbf{Y}, \mathbf{K}) = -nc + 10\left(1 - \left(\frac{1+p}{2}\right)^n\right) + \left(\frac{1+p}{2}\right)^n - \left(\frac{1-p}{2}\right)^n - 10\left(\frac{1-p}{2}\right)^n.$$

Then, for any c and p, we can compute the optimal number of draws $n^*(c; \mathbf{Y}, \mathbf{K})$. □

References

Bearden JN, Connolly T (2007) Multi-attribute sequential search. *Organ. Behav. Human Decision Processes* 1031:147–158.

Bell DE (1988) The value of pre-decision side bets for utility maximizers. *Management Sci.* 34(6):797–800.

Benhabib J, Bull C (1983) Job search: The choice of intensity. *J. Political Econom.* 91(5):747–764.

Bond SD, Carlson KA, Keeney RL (2008) Generating objectives: Can decision makers articulate what they want? *Management Sci.* 54(1):56–70.

Chade H, Smith L (2006) Simultaneous search. *Econometrica* 74(5): 1239–1307.

Chao R, Erat S (2012) Competition and effort in contests with evaluation uncertainty. Working paper, Darden School of Business, Charlottesville, VA.

Chao R, Lichtendahl KC Jr, Grushka-Cockayne Y (2014) Incentives in a stage-gate process. *Production Oper. Management* 23(8): 1286–1298.

Costantino N, Dotoli M, Falagario M, Sciancalepore F (2012) Balancing the additional costs of purchasing and the vendor set dimension to reduce public procurement costs. *J. Purchasing Supply Management* 18(3):189–198.

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

David HA, Nagaraja HN (2003) *Order Statistics* (John Wiley & Sons, Hoboken, NJ).

de Boer L, Dijkhuizen G, Telgen J (2000) A basis for modeling the costs of supplier selection: The economic tender quantity. *J. Oper. Res. Soc.* 51(10):1128–1135.

Denuit M, Eeckhoudt L, Tsetlin I, Winkler R (2013) Multivariate concave and convex stochastic dominance. Biagini F, Richter A, Schlesinger H, eds. *Risk Measures and Attitudes*, European Actuarial Academy (EAA) Series (Springer, Berlin), 11–32.

Elfving JA, Tommelein ID, Ballard G (2005) Consequences of competitive bidding in project-based production. *J. Purchasing Supply Management* 11(4):173–181.

Epstein LG, Tanny SM (1980) Increasing general correlation: A definition and some economic consequences. *Canad. J. Econom.* 13(1):16–34.

Fang K, Kotz S, Ng K (1990) *Symmetric Multivariate and Related Distributions* (Chapman & Hall, New York).

Heijboer G, de Boer L (2001) Optimal number of tenders in practice. *The 10th Internat. Annual IPSERA Conf. 2001* (International Purchasing and Supply Education and Research Association, Jönköping, Sweden), 473–484.

Heijboer G, Telgen J (2002) Choosing the open or the restricted procedure: A big deal or a big deal? *J. Public Procurement* 2(2): 187–215.

Ho W, Xu X, Dey PK (2010) Multi-criteria decision making approaches for supplier evaluation and selection: A literature review. *Eur. J. Oper. Res.* 202(1):16–24.

Keeney RL (2002) Common mistakes in making value trade-offs. *Oper. Res.* 50(6):935–945.

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

Lim C, Bearden JN, Smith JC (2006) Sequential search with multiattribute options. *Decision Anal.* 3(1):3–15.

Lippman SA, McCall JJ (1976) Job search in a dynamic economy. *J. Econom. Theory* 12(3):365–390.

Lippman SA, McCardle KF (1991) Uncertain search: A model of search among technologies of uncertain values. *Management Sci.* 37(11):1474–1490.

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

Lorentziadis PL (2010) Post-objective determination of weights of the evaluation factors in public procurement tenders. *Eur. J. Oper. Res.* 200(1):261–267.

March JG (1978) Bounded rationality, ambiguity, and the engineering of choice. *Bell J. Econom.* 9(2):587–608.

Microsoft (2014) Microsoft Board names Satya Nadella as CEO. Press release (February 4), Microsoft, Redmond, WA. http://www.microsoft.com/en-us/news/press/2014/feb14/02-04newspr.aspx.

Morgan PB (1983) Search and optimal sample sizes. *Rev. Econom. Stud.* 50(4):659–675.

Morgan P, Manning R (1985) Optimal search. *Econometrica* 53(4): 923–944.

Müller A, Stoyan D (2002) *Comparison Methods for Stochastic Models and Risks* (John Wiley & Sons, Chichester, UK).

Nelson RR (1961) Uncertainty, learning, and the economics of parallel research and development efforts. *Rev. Econom. Statist.* 43(4):351–364.

Ovide S, Langley M, Lublin JS (2014) Inside the Microsoft CEO search: How Satya Nadella was chosen. *Wall Street Journal* (February 5), http://online.wsj.com/news/articles/SB10001424052702304450904579364693178656198.

Owen J, Rabinovitch R (1983) On the class of elliptical distributions and their applications to the theory of portfolio choice. *J. Finance* 38(3):745–752.

Palaneeswaran E, Kumaraswamy MM (2000) Contractor selection for design/build projects. *J. Construction Engrg. Management* 126(5):331–339.

Richard SF (1975) Multivariate risk aversion, utility independence and separable utility functions. *Management Sci.* 22(1):12–21.

Rogerson R, Shimer R, Wright R (2005) Search-theoretic models of the labor market: A survey. *J. Econom. Literature* 43(4):959–988.

Silver EA (1987) A procedure for estimating the number of competitive bids for procurement. *IIE Trans.* 19(4):468–471.

Smith JE, Ulu C (2012) Technology adoption with uncertain future costs and quality. *Oper. Res.* 60(2):262–274.

Srinivasan V, Lovejoy WS, Beach D (1997) Integrated product design for marketability and manufacturing. *J. Marketing Res.* 34(1):154–163.

Stigler GJ (1961) The economics of information. *J. Political Econom.* 69(3):213–225.

Terwiesch C, Ulrich KT (2009) *Innovation Tournaments: Creating and Selecting Exceptional Opportunities* (Harvard Business School Publishing, Boston).

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

Tsetlin I, Winkler RL (2007) Decision making with multiattribute performance targets: The impact of changes in performance and target distributions. *Oper. Res.* 55(2):226–233.

Tsetlin I, Winkler RL (2009) Multiattribute utility specifying a preference for combining good with bad. *Management Sci.* 55(12): 1942–1952.

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

Yanagimoto T, Okamoto M (1969) Partial orderings of permutations and monotonicity of a rank correlation statistic. *Ann. Inst. Statist. Math.* 21(3):489–506.

Otso Massala is a Ph.D. candidate at INSEAD. His research focuses on decision and risk analysis. He has particular interests in modeling search, innovation tournaments, project portfolio management, and delegation of search. Other areas of interest include intertemporal choice, resource allocation processes, strategy cascading, incentivization of creativity, and motivating organizational search.

Ilia Tsetlin is an associate professor of decision sciences at INSEAD. His research interests are in modeling decisions under uncertainty, with particular focus on decision making with multiple attributes and stochastic dominance. Other research streams are related to negotiation, auction theory, and collective choice.