

Paradoxes in Learning and the Marginal Value of Information

Peter I. Frazier, Warren B. Powell

April 14, 2010

Abstract

We consider the Bayesian ranking and selection problem, in which one wishes to allocate an information collection budget as efficiently as possible to choose the best among several alternatives. In this problem, the marginal value of information is not concave, leading to algorithmic difficulties and apparent paradoxes. Among these paradoxes is that when there are many identical alternatives, it is often better to ignore some completely and focus on a smaller number than it is to spread the measurement budget equally across all the alternatives. We analyze the consequences of this non-concavity in several classes of ranking and selection problems, showing that the value of information is “eventually concave,” i.e., concave when the number of measurements of each alternative is large enough. We also present a new fully sequential measurement strategy that addresses the challenge that non-concavity it presents.

1 Introduction

Often in business we must choose from a set of alternatives whose values are uncertain. We typically have the ability to learn more about each alternative, but are constrained by time and financial budgets. A natural desire is to allocate our learning budget across the potential decisions to most improve our chances of making a good final decision.

Some examples of these decisions include:

- We would like to find the best supplier for a component part. We know the price of the component, but we do not know the reliability of the service or the quality of the product. We can collect information on service and product quality by placing small orders.

- We need to identify the best set of features to include in a new laptop we are manufacturing. We can estimate market response by running market tests, but these are time consuming and delay the product launch.
- We need to identify the best advertising message to maximize the number of clicks on a website. We can run market tests to evaluate the effectiveness of different messages.
- A charter jet operator has developed a simulator to predict the performance of a particular mix of aircraft types. The simulator takes a day to run, and we need to find the best mix of aircraft types.
- A political candidate must decide in which states to invest his remaining time for campaigning. He decides which states would benefit the most through telephone polls, but has to allocate a fixed budget for polling. How should he allocate his polling budget?

These are each instances of the ranking and selection (R&S) problem, which can be phrased more generally as follows. We start with a distribution of belief about each alternative from which we must choose. We then collect information about particular alternatives to improve our knowledge of them. As we collect more information, our knowledge of an alternative improves, sharpening our ability to make the best choice.

Now imagine, as is typically the case, that we have a budget for collecting information. To illustrate, consider the polling example where we must decide in which of 20 states to focus our attention. Assume we have enough money to make 10,000 phone calls (it is a low budget campaign). If we spread this around all 20 states, we will be able to make only 500 phone calls per state. If our candidate is polling near 50 percent, our estimates of how many people like our candidate will have sampling errors of around plus or minus 9 percent, which will be of little value for making decisions. But if we focus our polling on six states, we can drop our sampling error to plus or minus 5 percent, which offers sufficient precision to actually guide campaigning, but only in six states.

The problem of too many choices has received recognition. In “The Paradox of Choice: Why More is Less” (Schwartz (2004)), numerous examples of consumers facing too many

choices are provided, but without any formal analysis of the implications. The allocation of limited resources to collect information depends on the nature of competing alternatives, our confidence in our initial estimates, and the quality of the information that we can collect. In this paper, we investigate the factors that impact the marginal value of information, and show that since the value of information can be non-concave, we may need to restrict the number of alternatives we consider when trying to identify the best.

The value of information has a long history spanning the literatures of several disciplines. Within the economics literature, the influential work Stigler (1961) considers the value of information in the problem faced by buyers searching for the best price. The seminal work Howard (1966) formulates the value of information in a decision-theoretic context, and spawned a great deal of work in this area both within the decision-theoretic and operations research communities and beyond, including medical decision-making (see Yokota and Thompson (2004) for a recent review) and computer science (see, e.g., Kaelbling et al. (1998) for its use in reinforcement learning). The classic text Raiffa and Schlaifer (1968) poses the Bayesian R&S problem and defines the associated value of information, which marked the beginning of a number of research articles treating the value of information within Bayesian R&S, including Guttman and Tiao (1964); Tiao and Afonja (1976); Berger and Deely (1988). This work continues in the simulation optimization community, and also within computer science where it is known as the budgeted learning problem (Lizotte et al., 2003; Kapoor and Greiner, 2005; Guha and Munagala, 2007; Goel et al., 2008).

Within simulation-optimization, the R&S problem appears when allocating a computing budget to determine the set of parameters that produces the best results in a simulation. Chen (1995) introduces Optimal Computing Budget Allocation (OCBA) as a solution to this problem, which is studied in a series of articles (Chen et al. (1996, 1997, 2000a) and Chen et al. (2000b)). Other algorithmic solutions include LL(S) (Chick and Inoue (2001)), LL(1) (Chick et al., 2007, 2009) and the knowledge gradient (Frazier et al., 2008). Although this line of research proposes algorithms for the collection of information, some of which are based upon the marginal value of information, it does not formally investigate the underlying structure of this marginal value.

Intuition suggests that the value of information should be concave in the amount of information collected – as we collect more information it is natural to expect that the marginal value of this information should decrease. In fact, this is not true, and there are many documented cases

The structure of the value of information, and in particular its lack of concavity, has been explored in many information collection problems other than R&S. Howard (1966) describes an auction problem in which the value of perfect information is not concave. Radner and Stiglitz (1984) show general conditions under which the value of sampling is not concave. Chade and Schlee (2002) and De Lara and Gilotte (2007) extend these results to more general settings. Moscarini and Smith (2002) give asymptotic conditions under which the value of information is concave in a finite-state setting with a single information source. In a modified version of the R&S problem, Weibull et al. (2007) demonstrate that non-concavities in the value of information can lead to paradoxical discontinuities in optimal behavior, and Mattsson et al. (2004) describe in an unpublished working paper a case where it may be better to select an alternative that appears to be worse. Bickel and Smith (2006) demonstrate, using the setting of oil exploration (as a form of information collection) that information collection can exhibit increasing marginal returns, but do not address the more general problem of allocating resources for information collection.

We use the R&S framework with a single-stage of normal samples of known variance to study the structure of the marginal value of information and to characterize its concavity. While there are cases in which the value of information is not concave, we show that there exist thresholds such that the value of information is concave when the number of measurements of each alternative exceeds the corresponding threshold. We call this property “eventual concavity,” and it may be understood as implying that the value of information is concave as long as the number of measurements is large enough. In the special case of measuring a single alternative, we give an explicit expression for the region on which the value of information is concave.

While this eventual concavity result demonstrates that lack of concavity occurs in a restricted set of situations, it is nevertheless important in some cases. We illustrate some para-

doxes that result from the non-concavity, including the possible non-optimality of spreading the measurement budget equally among the alternatives when the prior is symmetric. It may be better to completely ignore some alternatives (chosen completely at random) and to allocate the remaining budget among those that remain. While paradoxes resulting from the lack of concavity in the value of information in other problems have been described in the literature, the strange behaviors resulting from this lack of concavity in R&S have not been fully explored elsewhere.

We also consider sequential measurement policies, where the non-concavity of information can cause difficulties for algorithms that allocate one or a small number of measurements at a time based on value of information calculations. We describe a policy, called $KG^{(*)}$, for overcoming these complications introduced by non-concavity in the value of information. This policy was introduced in Chick and Frazier (2009a,b) as part of a larger numerical study, where it was described only very briefly. We provide a more complete description.

We start by describing the Bayesian R&S problem formally in Section 2. We also give preliminary results to be used later. We then discuss geometric properties of the value of information in a number of special cases. Section 3 considers the case in which we measure only a single alternative. Section 4 considers the case of a homogeneous prior. Section 5 considers the case of only two alternatives. In Section 6 we return to the general problem, present an example in which the non-concavity of the value of information causes non-intuitive behavior in the optimal allocation, and give a theoretical result that shows that the value of information is “eventually concave,” in the sense that there exists a region of the measurement space that expands out to infinity in all directions on which the value of information is concave. Section 7 considers the consequences of non-concavity for sequential measurement policies and describes the $KG^{(*)}$ policy for overcoming this non-concavity. Except where provided in the text, proofs may be found in the appendix.

2 The Ranking and Selection Problem

We begin by providing a formal definition of the R&S problem. This definition follows the standard definition in the literature — see, e.g., Raiffa and Schlaifer (1968). Suppose that we have a collection of M alternatives, and associated with each of these alternatives is a sampling distribution. The sampling distribution for alternative $i \in \{1, \dots, M\}$ is normal with mean θ_i and variance λ_i , and samples are independent of each other when conditioned on $\theta_1, \dots, \theta_M$. The sampling variances $\lambda_1, \dots, \lambda_M$ are known, but the sampling means $\theta_1, \dots, \theta_M$ are unknown. We begin with a normally distributed Bayesian prior belief on the sampling means that is independent across alternatives, $\theta_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$. We refer to the vector $[\theta_1, \dots, \theta_M]$ as θ . The vectors μ and σ^2 are defined similarly.

In this problem we are interested in determining through sampling which of the alternatives has the largest sampling mean. Consider any sampling allocation $n \in \mathbb{Z}_+^M$, by which we mean the sampling strategy that draws n_i samples from each alternative i , where i ranges over $1, \dots, M$. Here, $\mathbb{Z}_+ = \{0, 1, \dots\}$. Call the resulting set of observations Y . It can then be shown (see, e.g., DeGroot (1970)) that the posterior distribution on θ is normal with independent components. The posterior variance of θ_i is $\text{Var}[\theta_i | Y, n] = \left(\frac{1}{\sigma_i^2} + \frac{n_i}{\lambda_i}\right)^{-1}$. Note that it does not depend on the observations Y . The posterior mean $\mathbf{E}[\theta_i | Y, n]$, on the other hand, does depend on the observations. Its predictive distribution is normal, and the mean of this predictive distribution is given by the tower property of conditional expectation, $\mathbf{E}[\mathbf{E}[\theta_i | Y, n]] = \mathbf{E}[\theta_i] = \mu_i$. We denote by $\tilde{\sigma}_i^2(n_i)$ the variance of the predictive distribution, and it can be computed by the conditional variance formula as

$$\tilde{\sigma}_i^2(n_i) = \text{Var}[\mathbf{E}[\theta_i | Y, n]] = \text{Var}[\theta_i] - \mathbf{E}[\text{Var}[\theta_i | Y]] = \sigma_i^2 - \left(\frac{1}{\sigma_i^2} + \frac{n_i}{\lambda_i}\right)^{-1} = \frac{\sigma_i^2 n_i}{(\lambda_i/\sigma_i^2) + n_i}. \quad (1)$$

We write $\tilde{\sigma}(n)$ to indicate the vector $[\tilde{\sigma}_1(n_1), \dots, \tilde{\sigma}_M(n_M)]$. Expressions for the first and second derivatives of $\tilde{\sigma}_i$ are needed in proofs, and are provided in the appendix in Lemma 2.

It is important to keep in mind that there are four distinct but related normal distributions involved with this problem. First, there is the sampling distribution $\mathcal{N}(\theta_i, \lambda_i)$ that governs the samples observed from alternative i . Then, there is the prior $\mathcal{N}(\mu_i, \sigma_i^2)$ on θ_i

as well as the posterior $\mathcal{N}(\mathbf{E}[\theta_i | Y, n], \text{Var}[\theta_i | Y, n])$. Finally, since $\mathbf{E}[\theta_i | Y, n]$ is a conditional expectation whose value is unknown until the observations Y are seen, it also has a distribution which is $\mathcal{N}(\mu_i, \tilde{\sigma}_i^2(n_i))$.

After sampling according to the sampling allocation n , which directs us to take n_i samples from each alternative i , we choose one alternative and receive a reward equal to its sampling mean. Assuming that we are most interested in maximizing the expected value of this reward, and have no risk aversion, the best choice is the alternative with the largest posterior mean $\mathbf{E}[\theta_i | Y, n]$, and the expected value of choosing this alternative is $\max_i \mathbf{E}[\theta_i | Y, n]$.

We now calculate the value $v(n)$ of the information obtained from the sampling allocation n . This value of information $v(n)$ is defined to be the incremental improvement over the best expected value that can be obtained without measurement, which is $\max_i \mu_i$.

$$v(n) = \mathbf{E} \left[\max_i \mathbf{E}[\theta_i | Y, n] \mid n \right] - \max_i \mu_i. \quad (2)$$

Given a sampling budget $N \in \mathbf{Z}_+$, the R&S problem is to find the sampling allocation satisfying the budget constraint $\sum_i n_i \leq N$ that maximizes the value of the information obtained. That is,

$$\max_{n \in \mathbf{Z}_+^M : \sum_i n_i \leq N} v(n). \quad (3)$$

If the prior parameters μ and σ^2 come directly from a practitioner's initial belief, this problem is called the single-stage Bayesian R&S problem. If, instead, μ and σ^2 are the parameters of a posterior resulting from the combination of some initial belief and some previous set of measurements, this problem is the second stage of what is called a two-stage Bayesian R&S problem. In this article, we study the structure of the function v , both as interesting on its own, and as a foundation for solving one-stage, two-stage, and fully sequential (see Section 7) Bayesian R&S problems.

We may obtain an expression for $v(n)$ that is more explicit than (2) using the predictive distribution $\mathbf{E}[\theta_i | Y, n] \sim \mathcal{N}(\mu_i, \tilde{\sigma}_i^2(n_i))$. Defining independent standard normals Z_1, \dots, Z_N by $Z_i = (\mathbf{E}[\theta_i | Y, n] - \mu_i) / \tilde{\sigma}_i(n_i)$, we have

$$v(n) = \mathbf{E} \left[\max_i \mu_i + \tilde{\sigma}_i(n_i) Z_i \right] - \max_i \mu_i. \quad (4)$$

This expression offers a convenient platform for analysis. When suitably transformed, it can also be computed efficiently as a 1-dimensional integral of a function of normal cumulative distribution functions.

Although sampling allocations n are generally discrete in nature, we may extend the function v continuously onto \mathbf{R}_+^M using the definition (4). By dropping the integrality constraint in (3) and solving the resulting relaxed problem, we obtain an upper bound on the value of the optimal integral solution. Also, by rounding the solution to the relaxed problem we obtain a feasible solution to (3) and a lower bound on the value of the optimal integral solution. If the gap between upper and lower bounds is small, we may take the rounded solution as a good suboptimal allocation.

Although the primary motivation for the continuous extension of v is as an analytically convenient tool for obtaining approximate solutions to the integer-constrained R&S problem (3), one may also consider problems in which observation occurs continuously through time, and the measurement variance λ_i is the variance per unit time of the observation process. In this case, the value n_i is the amount of time to observe alternative i , and need not satisfy an integrality constraint. Such problems are considered elsewhere in the literature. See, e.g., Chick and Gans (2009).

It is sometimes useful to work with v in terms of a function $g : \mathbf{R}_+^M \mapsto \mathbf{R}_+$ defined by

$$g(s) = \mathbf{E} \left[\max_i \mu_i + s_i Z_i \right] - \max_i \mu_i. \quad (5)$$

This definition is chosen so that $v(n) = g(\tilde{\sigma}(n))$. Expressions for the first and second partial derivatives of g are needed in proofs, and are provided in the appendix in Lemmas 3 and 4.

We conclude by noting the following remark, which is useful in computation. Here and throughout, Φ is the standard normal cumulative distribution function and φ is the standard normal probability density function.

Remark 1. For $c \in \mathbf{R}_+$, and Z a standard normal random variable,

$$\mathbf{E} [Z \mathbf{1}_{\{Z \geq c\}}] = \frac{1}{\sqrt{2\pi}} \int_{[c, \infty)} z \exp(-z^2/2) dz = \frac{1}{\sqrt{2\pi}} \int_{[c^2/2, \infty)} \exp(-u) du = \varphi(c),$$

which follows from the substitution $u = z^2/2$.

3 Measuring a Single Alternative

We begin by studying the structure of the function $n \mapsto v(ne_i)$ for a *fixed* alternative i . This function gives the value of measuring that alternative n times without measuring any other alternatives.

This function and its derivative may be computed analytically, as in the following proposition. We provide a proof in the appendix for completeness, but its development is similar to proofs that may be found elsewhere, including Frazier et al. (2008) and Gupta and Miescke (1996).

Proposition 1. *Define a constant $\Delta_i = |\mu_i - \max_{j \neq i} \mu_j|$ and a function $f : \mathbb{R} \mapsto \mathbb{R}$ by $f(z) = z\Phi(z) + \varphi(z)$. Then,*

$$v(ne_i) = \begin{cases} \tilde{\sigma}_i(n_i) f(-\Delta_i/\tilde{\sigma}_i(n_i)), & \text{if } n_i > 0, \\ 0, & \text{if } n_i = 0, \end{cases}$$

$$\frac{\partial}{\partial n_i} v(n_i e_i) = \tilde{\sigma}'_i(n_i) \varphi\left(\frac{\Delta_i}{\tilde{\sigma}_i(n_i)}\right), \quad \text{for } n_i > 0.$$

This proposition shows that the value of measurement increases with the number of measurements. This monotonicity is natural, in light of the fact that more information allows us to make better decisions. The value of measurement is plotted in Figure 1 for $\Delta_i = 10$. The prior variance σ_i^2 and the noise variance λ_i are both fixed to 1.

Notice that the value of measurement is not concave in general. It is concave when there are many measurements, but when there are few the value of measurement may be initially convex. This is especially true when the measurement noise is large relative to the differences between alternatives. This lack of concavity may be understood by realizing that when measurements are few and/or noisy, it is very unlikely that they will provide enough evidence to cause us to change our mind about which alternative is best. Even if the alternative being measured is actually better than the other alternatives, the strength of our prior and the weakness of the evidence from measurement will cause us to ignore the evidence collected. Instead, we will choose as best the same alternative that we thought was




Figure 1: The value of measuring a single alternative, $v(n_i e_i)$, as a function of n_i , how many times it is measured. The left-hand plot shows the value, and the right-hand plot shows the (natural) log of its value. In this example, $\sigma_i = 1$, $\lambda_i = 1$, and $\Delta_i = 10$.

best before the measurements. Only when measurements are sufficiently precise or numerous will they have the ability to affect our ultimate decision.

We may think of $\tilde{\sigma}_i(n)$ as one measure of the amount of information collected. Indeed, $\tilde{\sigma}_i^2(n)$ is the variance of the change in our belief due to measurement. As noted previously, $\tilde{\sigma}(n)$ is concave in n . But $v(n e_i) = g(\tilde{\sigma}_i(n) e_i)$, and g is a convex function. We may think of g as the function which translates a more well-behaved measure of the amount of information collected into the value of that information. Thus, even if the “amount” of information collected is concave in the number of measurements, the ability of this information to affect the final decision can be non-concave, and so is its overall value. In some sense, it is the discreteness of the ultimate implementation decision that causes this non-concavity in the value of measurement.

In the case in which we only measure a single alternative, the following theorem completely characterizes convexity and concavity of the value of measurement.

Theorem 1. $n_i \mapsto v(n_i e_i)$ is convex on $(0, n_i^*]$ and concave on (n_i^*, ∞) , where

$$n_i^* = \frac{\lambda_i}{8\sigma_i^4} \left[\Delta_i^2 - \sigma_i^2 + \sqrt{\Delta_i^4 + 14\sigma_i^2 \Delta_i^2 + \sigma_i^4} \right]. \quad (6)$$

Proof. Let $n_i > 0$. By Proposition 1, $\frac{\partial^2}{\partial n_i^2} v(n_i e_i) = \frac{\partial}{\partial n_i} \tilde{\sigma}_i'(n_i) \varphi\left(\frac{-\Delta_i}{\tilde{\sigma}_i(n_i)}\right)$, which can be rewrit-

ten as

$$\frac{\partial^2}{\partial n_i^2} v(n_i e_i) = \varphi \left(\frac{-\Delta_i}{\tilde{\sigma}_i(n_i)} \right) \frac{1}{\tilde{\sigma}_i(n_i)} \left[\tilde{\sigma}_i''(n_i) \tilde{\sigma}_i(n_i) + \left(\frac{\Delta_i \tilde{\sigma}_i'(n_i)}{\tilde{\sigma}_i(n_i)} \right)^2 \right]. \quad (7)$$

Letting $r = \lambda_i / \sigma_i^2$, the expressions for the first and second derivatives of $\tilde{\sigma}_i$ from Lemma 2 in the appendix can be rewritten as,

$$\begin{aligned} \tilde{\sigma}_i''(n_i) \tilde{\sigma}_i(n_i) &= -\frac{\sigma_i^2 r}{4} n_i^{-1} (r + n_i)^{-3} (r + 4n_i), \\ \tilde{\sigma}_i'(n_i) / \tilde{\sigma}_i(n_i) &= \frac{r}{2} n_i^{-1} (r + n_i)^{-1}. \end{aligned}$$

This allows us to evaluate the term in the square brackets in (7) as

$$\tilde{\sigma}_i''(n_i) \tilde{\sigma}_i(n_i) + \left(\frac{\Delta_i \tilde{\sigma}_i'(n_i)}{\tilde{\sigma}_i(n_i)} \right)^2 = \frac{r}{4} n_i^{-2} (r + n_i)^{-3} [-4\sigma_i^2 n_i^2 + r(\Delta_i - \sigma_i^2)n_i + \Delta_i^2 r^2].$$

The quadratic expression $-4\sigma_i^2 n_i^2 + r(\Delta_i^2 - \sigma_i^2)n_i + \Delta_i^2 r^2$ has two roots,

$$\frac{r}{8\sigma_i^2} \left[\Delta_i - \sigma_i^2 \pm \sqrt{(\Delta_i^2 - \sigma_i^2)^2 + 16\sigma_i^2 \Delta_i^2} \right] = \frac{r}{8\sigma_i^2} \left[\Delta_i - \sigma_i^2 \pm \sqrt{\Delta_i^4 + 14\sigma_i^2 \Delta_i^2 + \sigma_i^4} \right],$$

call them n_+ and n_- , where n_+ is the larger of the two. Note that $n_- < 0 \leq n_+$, where the strict inequality $n_- < 0$ is implied by $\sigma_i^2 > 0$. Then writing the quadratic as $(n_i - n_+)(n_i - n_-)$, we have

$$\frac{\partial^2}{\partial n_i^2} v(n_i e_i) = \varphi \left(\frac{-\Delta_i}{\tilde{\sigma}_i(n_i)} \right) \frac{1}{\tilde{\sigma}_i(n_i)} \frac{r}{4} n_i^{-2} (r + n_i)^{-3} (n_i - n_+)(n_i - n_-).$$

Since all the terms except $(n_i - n_+)$ are strictly positive for non-negative n_i , $n_i \mapsto v(n_i e_i)$ is convex on $(0, n_+]$ and concave on $[n_+, \infty)$. The theorem follows from setting n_* equal to n_+ . \square

This theorem tells us that, when Δ_i or λ_i is larger, the region of non-concavity is larger. This makes sense in light of our previous understanding, since larger values for Δ_i or λ_i imply that we need a more substantial body of evidence (in the form of more measurements) to affect our decision-making.

Theorem 1 has as a consequence the following corollary, which gives two conditions under which the value of measuring an alternative is concave everywhere. When an alternative is tied for the best ($\Delta_i = 0$), even a single measurement with very large measurement

noise is enough to cause us to change our decision – our prior belief is ambivalent about which implementation decision is better, and so there is no bias imposed by the prior that measurements must overcome. When measurements are perfect ($\lambda_i = 0$), any non-zero amount of measurement is enough to overcome any bias imposed by the prior.

Corollary 1. *If $\Delta_i = 0$ or $\lambda_i = 0$, then $n_i \mapsto v(n_i e_i)$ is concave on \mathbf{R}_{++} .*

Proof. $\Delta_i = 0$ or $\lambda_i = 0$ imply $n_i^* = 0$, and concavity on \mathbf{R}_{++} follows from Theorem 1. \square

We also have another type of structure in the single-measurement case which is given in the following proposition. This proposition shows log-concavity of the value of measurement, but only for a measurement of a single alternative.

Proposition 2. *For each alternative i , the function $n \mapsto v(n e_i)$ with domain \mathbf{R}_+ is log-concave.*

4 The Homogeneous Case

Consider the measurement problem that arises when our prior parameters and measurement variances are constant across alternatives, i.e., $\mu_i = \mu$, $\sigma_i^2 = \sigma^2$ and $\lambda_i = \lambda$ for all i and some μ , σ^2 , and λ . We call such measurement problems *homogeneous*.

Homogeneous problems are quite common in practice. They arise when we cannot distinguish between alternatives before measuring them, but we are willing to make a priori judgments of the population *as a whole*. As an example, consider evaluating job applicants. Fairness may require that we do not favor a priori one applicant over another (i.e., that μ_i and σ_i^2 are constant), but our past experience gives us a good idea of the distribution of quality within the applicant pool as a whole (guiding our choice of μ and σ^2). In other situations, for example when evaluating different designs for a very complex and poorly understood system using simulation, the a priori indistinguishability between alternatives may arise from a lack of strong prior knowledge or a desire for objectivity. In still other situations, for example when purchasing a piece of capital equipment from among several of




Figure 2: The value of information as a function of n_1 for a homogenous problem with $M = 100$ alternatives, $n_i = 1$ for $i > 1$, $\mu = 0$, $\sigma^2 = 10$, and $\lambda = 10^6$. Two different ranges of n_1 are pictured.

the same model produced by the same company, it may arise because alternatives genuinely lack distinguishing characteristics.

As noted previously, the value of information is not concave in general, and this lack of concavity may be seen in the homogeneous case. Figure 2 shows the value of information for a homogeneous problem with $M = 100$ alternatives, $\mu = 0$, $\sigma^2 = 10$ and $\lambda = 10^6$. The value of information is pictured as a function of the number of measurements of the first alternative, where exactly 1 measurement is allocated to each of the other alternatives. The figure clearly shows that the value of information is not concave in n_1 . This lack of concavity occurs despite the fact that constant μ_i implies $\Delta_i = 0$, and so Corollary 1 shows that the value of measuring a *single alternative* is concave. This seeming disparity may be understood by realizing that, after measuring all but the first alternative, alternative 1 must overcome a gap of $(\max_{i \neq 1} \mathbf{E}[\theta_i | Y, n]) - \mu_1$ to be implemented.

In the homogeneous case, intuition and symmetry suggest that our best course of action would be to spread our measurement budget equally across the alternatives. Indeed, existing general purpose allocation techniques for this single-stage R&S problem with linear loss, LL(B) (Chick and Inoue, 2001) and OCBA for linear loss (He et al., 2007), both recommend the equal allocation in the homogeneous case. (Although they recommend equal allocation in the homogeneous case, they make other recommendations in other cases). Surprisingly, however, we can often obtain more value by focusing our sampling on a strict subset of the

alternatives and ignoring the rest.

To formalize this class of policies that focus effort on less than the full set of alternatives, let $m \leq M$ be an integer. The allocation given by m is the one that chooses m alternatives uniformly at random from the full set of M alternatives. It then allocates the measurement budget equally across these m alternatives, so each one receives N/m measurements. This class of policies includes the equal allocation policy as $m = M$.

Figure 3(a) shows the value obtained by such policies. Each line in the plot shows, for a different value of the measurement noise λ , the value obtained as a function of m . The problems considered have $M = 50$ alternatives, a measurement budget of $N = 50$, and a homogeneous prior $\mu = 0$, $\sigma^2 = 1$. The values of λ generating each line, from top to bottom, are $1/2, 1, 2, 4, 8, 16$. Circles show the maximum attained by each line.

From the right-most edge of the figure, where $m = M$ and the policy pursued is the equal allocation policy, we see that the equal allocation policy is often suboptimal. The only case pictured for which equal allocation is optimal is the top line, $\lambda = 1/2$. As we move down the figure and the noise variance λ increases, the equal allocation becomes more suboptimal. Thus, when the measurement noise is large, one may improve significantly upon the naive equal allocation policy by ignoring some alternatives completely and focusing measurement effort on those that remain. Equal allocation over the full set of alternatives dilutes the measurement budget, while ignoring some alternatives and allocating the measurement budget over those that remain concentrates it, and allows enough evidence to accumulate to affect our decisions in a significant way. As the noise variance increases, we need to further concentrate the budget, and we must ignore more alternatives.

We now provide an expression for allocation value as a function of m that is useful for finding the best value of m . The value of the allocation generated by m is $v\left(\frac{N}{m} \sum_{i=1}^m e_i\right)$, which implicitly depends upon M , λ , and σ^2 in addition to depending explicitly on N and

m . Momentarily fix m and write the value of the corresponding allocation,

$$\begin{aligned} v \left(\frac{N}{m} \sum_{i=1}^m e_i \right) &= \mathbf{E} \left[\max \mu_i + \tilde{\sigma}_i \left(\frac{N}{m} \mathbf{1}_{\{i \leq m\}} \right) Z_i \right] - \max_i \mu_i \\ &= \begin{cases} \mathbf{E} \left[(\max_{i \leq m} \tilde{\sigma}_i (N/m) Z_i)^+ \right] & \text{if } m < M \\ \mathbf{E} \left[\max_{i \leq m} \tilde{\sigma}_i (N/m) Z_i \right] & \text{if } m = M. \end{cases} \end{aligned}$$

Define $r = \frac{\lambda}{N\sigma^2}$, so that $\tilde{\sigma}_i(N/m) = \sigma/\sqrt{1+rm}$. Then,

$$v \left(\frac{N}{m} \sum_{i=1}^m e_i \right) = \begin{cases} \sigma \mathbf{E} \left[(\max_{i \leq m} (1+rm)^{-1/2} Z_i)^+ \right] & \text{if } m < M, \\ \sigma \mathbf{E} \left[\max_{i \leq m} (1+rm)^{-1/2} Z_i \right] & \text{if } m = M. \end{cases}$$

These expressions may be evaluated using one-dimensional numerical integration. We provide details for the case $m < M$, and the case $m = M$ may be evaluated similarly. Let $w(m, r)$ be the value obtained given r and $m < M$.

$$w(m, r) = \mathbf{E} \left[\max_{i \leq m} (1+rm)^{-1/2} Z_i \right] = \int_0^\infty \mathbb{P} \left\{ \max_{i \leq m} (1+rm)^{-1/2} Z_i > u \right\} du.$$

We rewrite the inner probability as

$$\begin{aligned} \mathbb{P} \left\{ \max_{i \leq m} Z_i > u\sqrt{1+rm} \right\} &= 1 - \mathbb{P} \left\{ \max_{i \leq m} Z_i \leq u\sqrt{1+rm} \right\} \\ &= 1 - \mathbb{P} \left\{ Z_i \leq u\sqrt{1+rm}, \forall i \leq m \right\} = 1 - \mathbb{P} \left\{ Z_1 \leq u\sqrt{1+rm} \right\}^m = 1 - \Phi \left(u\sqrt{1+rm} \right)^m. \end{aligned}$$

Thus, the value obtained for $m < M$ is

$$w(m, r) = \int_0^\infty 1 - \Phi \left(u\sqrt{1+rm} \right)^m du.$$

This expression for $w(m, r)$ extends continuously to non-integral values of m .

We may find the best m simply by evaluating $v \left(\frac{N}{m} \sum_{i=1}^m e_i \right) / \sigma$ over the set $m \in \{1, \dots, M\}$. Equivalently, we may divide this expression by the constant σ and optimize that expression instead. Observe that $v \left(\frac{N}{m} \sum_{i=1}^m e_i \right) / \sigma$ depends only upon r for $m < M$.

This leads us to recommend the following method for finding an approximately optimal value of m . We temporarily drop the integrality constraint and upper and lower bounds on m , allowing it to take values in $(0, \infty)$, and solve the resulting continuous optimization problem. Let $m^*(r)$ be the solution, $m^*(r) = \arg \max_{m \in [0, \infty)} w(m, r)$. The solution depends only upon r and is plotted in Figure 3(b).




Figure 3: (a) shows the value of allocations that spread the measurement budget equally across $m \leq M$ alternatives in the homogeneous case. $\mu = 0$, $\sigma^2 = 1$, and $M = N = 50$. From top to bottom, the lines use $\lambda = 1/2, 1, 2, 4, 8, 16$. Circles show the maximal point on each line. (b) shows $m^*(r)$ as a function of r .

In Figure 3(b), as $1/r$ decreases the optimal $m^*(r)$ decreases. Recalling the definition $1/r = N\sigma^2/\lambda$, we see that decreasing either N/λ or the prior variance σ^2 causes $m^*(r)$ to decrease as well. The quantity N/λ is an “effective” measurement budget, and gives the total measurement precision that may be partitioned among the alternatives. When it is smaller, discerning the best alternative is more difficult. When σ^2 is smaller, the value of different alternatives are typically closer together, and it is again harder to discern the best alternative. Thus, decreasing either quantity forces us to concentrate our measurement budget onto a smaller number of alternatives. This generalizes the observation from Figure 3(a) that as the measurement precision $1/\lambda$ decreases, so too does $m^*(r)$.

5 The Two-Alternative Case

We now briefly consider the case when the number of alternatives M is exactly 2, and the concavity of the value of information in this case. Because of a certain symmetry, it is much easier to find the optimal allocation in this case than when $M > 2$. This optimal allocation is described in Raiffa and Schlaifer (1968). We briefly review this optimal allocation, and its impact on the concavity of the value of information.

First, we have the following easy-to-use expression for the value of information. This

expression is well-known in the literature (see, for example, Raiffa and Schlaifer (1968)). Although the proof is straightforward, we provide one in the appendix for completeness.

Lemma 1. *Suppose $M = 2$, and let $s(n) = \sqrt{\tilde{\sigma}_1^2(n_1) + \tilde{\sigma}_2^2(n_2)}$. Then $v(n) = s(n)f(-|\mu_1 - \mu_2|/s(n))$, where we recall that $f(z) = z\Phi(z) + \varphi(z)$.*

Although $s(n) = \sqrt{\tilde{\sigma}_1^2(n_1) + \tilde{\sigma}_2^2(n_2)}$ is a concave function, neither f nor v is generally concave. Despite this lack of concavity, the optimal allocation is easily obtained by transforming the problem. The function $s \mapsto sf(-|\mu_1 - \mu_2|/s)$ is strictly increasing (because f is strictly increasing), and so maximizing the non-concave function $v(n)$ is equivalent to maximizing $s(n)$. Furthermore, because $s \mapsto \sqrt{s}$ is strictly monotonic on \mathbf{R}_+ , maximizing $s(n)$ (and hence also $v(n)$) is the same as maximizing $s(n)^2 = \tilde{\sigma}_1^2(n_1) + \tilde{\sigma}_2^2(n_2)$.

When measurement variances are equal, this maximization is accomplished by allocating one sample at a time to the alternative with the largest posterior variance. Note that the posterior variance is a deterministic function of the number of samples allocated, and so the resulting allocation is deterministic. If the prior variances are equal across alternatives and the sample size is even, then this results in an equal allocation. This deterministic allocation is also optimal among all *sequential* allocations (Frazier et al., 2008).

The value of information is not concave in general in the two-alternative case, as is easily seen by noting that the non-concavity of the value of measuring a single alternative occurs in the two-alternative case. However, we do have concavity in a special case noted in the following remark.

Remark 2. *The value of information is concave when $M = 2$ and $\mu_1 = \mu_2$. This is because Lemma 1 implies $v(n) = s(n)f(0) = \varphi(0) \sqrt{\tilde{\sigma}_1^2(n_1) + \tilde{\sigma}_2^2(n_2)}$, which is concave.*

6 The General Case

We now consider the general case, with no special assumptions on μ , σ , M , or the number of alternatives measured. We begin with an example in which the lack of concavity in the value of measurement has a dramatic and counterintuitive effect on the optimal allocation. We then

present a theoretical result that shows, despite the existence of such oddly behaved examples, the value of information is concave over a space of sufficiently large allocations. This may allow finding the optimal allocation, or a reasonably tight lower bound, by optimizing over this space using efficient convex programming algorithms.

Consider an example with $M = 3$ alternatives and prior $\mu = [1, 1, 0]$ and $\sigma^2 = [0, .005, .5]$. Under this prior, we know the value of the first alternative perfectly, we have a strong belief that the second alternative's value close to 1, and we think the third alternative is smaller (close to 0) but we are uncertain about this.

We consider different measurement budgets N , and for each one we find the optimal allocation. Because this illustrative example has a very small number of alternatives, we can find the optimal allocation by enumerating the set of possible allocations, calculating the value of each, and choosing the one with the largest value. The optimal policy allocates no measurements to the first alternative because its value is already known perfectly, so the set of allocations given N is $\{x \in \mathbb{Z}_+^3 : x_1 = 0, x_2 + x_3 = N\}$, and any allocation in this set is completely determined by the choice of, say, x_2 .

Figure 4 shows the value of various allocations in this example, as well as the behavior of the optimal allocation. Each line in Figure 4(a) shows the value obtained for a single value of N as a function of the proportion x_2/N . The optimal proportion for a given value of N is the maximal point on the corresponding curve. As N increases, the value increases and the optimal proportion changes. Figure 4(b) shows the optimal proportion as a function of N . The circles are optimal among those restricted to integral numbers of measurements, while the dashed line is optimal among all allocations.

When we have a single measurement to make ($N = 1$), the optimal allocation measures alternative 2, which is the one whose mean is largest but whose variance is extremely small. When our measurement budget is just slightly larger ($1 < N \leq 4$), the optimal allocation switches abruptly and spends its entire sampling budget on alternative 3, which is the one with smaller mean but larger variance. This dramatic shift in behavior is due to the lack of concavity in the value of sampling, and parallels similar behavior observed in the homogeneous case.

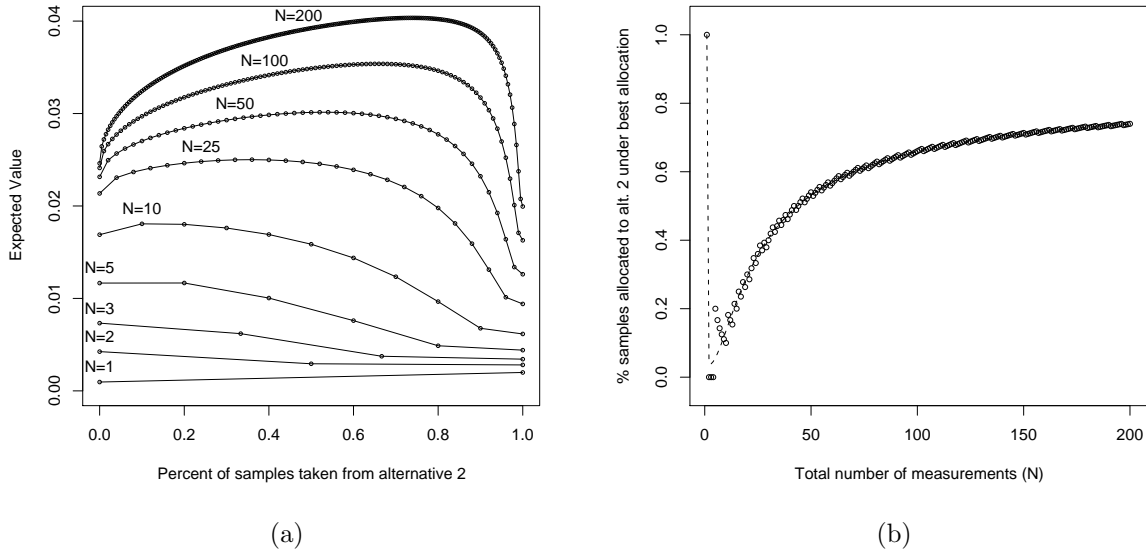


Figure 4: Value of information in the example from the text with $\mu = [1, 1, 0]$, $\sigma^2 = [0, .005, .5]$, and $\lambda = 1$. Panel (a) shows the value of the full range of allocations possible for various total measurement budgets N . Panel (b) shows the optimal allocation as a percentage of samples taken from the second alternative both with integrality constraints (circles) and without (dashed line).

As N increases further ($N > 4$), we suffer diminishing marginal returns from measurements of the low-mean high-variance alternative, and the optimal allocation allocates some of the growing measurement budget to the low-variance alternative. The proportion of samples allocated to the low-variance alternative is small initially, but then grows toward an asymptotic level. This asymptotic level does not depend on the prior variances — only upon the prior means.

The asymptotic proportion toward which the optimal policy is growing is the proportion that is optimal for large values of N . In this example, the asymptotic proportion is approximately 0.84. (This was established by calculating the optimal proportion for finite but very large values of N , and observing that the optimal proportions were all between 0.8383 and 0.8385 for values tested from 10^8 up to the largest value tested, 10^{10} .) This asymptotic proportion is often far from the optimal proportion for the given value of N . For example, Figure 4(b) shows that the best proportion when $N = 25$ is 0.36 which is quite far from 0.84. The predominant methods for finding good single-stage allocations, LL(B)

(Chick and Inoue, 2001) and the OCBA (Chen (1996), He et al. (2007)), approximate this optimal asymptotic proportion and then suggest that we allocate our available (finite) sampling budget according to this proportion. The difference between asymptotic and finite-horizon optimal proportions and its effect on these policies may be an interesting topic for future research.

Figure 5 shows contours of the value of information for this example, as well as circles for the optimal (integrality-constrained) allocations for values of N up to 11. It also shows a dark solid line dividing the region on which the value of information is not concave (to the left of the line) from the region on which it is concave (right of the line). Although the value of information is not concave in general, it is concave on most of the domain. Furthermore, the optimal policy’s switch from putting all its measurements on alternative 2 to putting all on alternative 3 occurs as we transition from the non-concave to the concave regions. Afterward, in the concave region, the optimal allocation behaves more reasonably. Figure 6 shows similar partitions of the allocation space into regions on which the value of information is concave and non-concave for this prior and two others. In all three priors, we see that, as long as the number of measurements allocated to each alternative is large enough, the value of information is concave.

The following theorem shows that such a region of “ n large enough” on which the value of information is concave always exists. This suggests that one way to improve the quality of an allocation found through existing methods, e.g., via LL(B) or OCBA, is to take the allocation given by either of these existing methods, and then perform gradient ascent or a second order convex optimization method. If both the starting allocation and the global maximum lie within this concave region, as seems likely when the measurement budget is large, then this technique will find the globally optimal allocation. If not, it may still lead to better allocations.

Theorem 2. *There exists an $N^* \in \mathbf{R}_+^M$ such that v is concave on the region $\{n \in \mathbf{R}_+^M : n_i \geq N_i^* \forall i\}$.*




Figure 5: Value of information contours for $\mu = [1, 1, 0]$, $\sigma^2 = [0, .005, .5]$, $\lambda = 1$. Circles show optimal integrality-constrained allocations. The dark vertical line separates the region on which the value of information is not concave (left of the line) from the region where it is.




Figure 6: Regions of concave value of information for three different priors. For a given prior (A, B or C), the value of information is concave in the region above and to the right of the corresponding line. Prior A is $\mu = [1, 1, 0]$ $\sigma^2 = [0, .005, .5]$; Prior B is $\mu = [1, 1.3, 1.3]$ $\sigma^2 = [0, .1, .1]$; and Prior C is $\mu = [1, 1.3, 1]$ $\sigma^2 = [0, .1, .1]$. In all cases $\lambda = 1$.

7 The Sequential Case

Previously we considered the non-concavity of the value of information for single-stage sampling. Now we consider the implications of this non-concavity for fully sequential policies, and in particular for a myopic fully sequential sampling policy known as the $(R1, \dots, R1)$ policy (Gupta and Miescke, 1996), and analyzed further in Frazier et al. (2008), where it was called the knowledge-gradient (KG) policy.

A fully sequential policy is one in which the allocation of each measurement is decided separately. Moreover, each decision is based upon the results of all previous measurements. Thus, a sequential policy may use the information collected so far to hone later measurement decisions, and so we may measure more efficiently with a sequential policy than we can with the single-stage policies discussed so far. In particular, many good sequential policies use early measurements to roughly identify which alternatives are likely to be the best, and then focus later effort on these alternatives in order to discern the best alternative from those that are merely good.

In the sequential setting, the value of a measurement depends upon the measurements

that will be taken in the future. Such values may be calculated, at least in theory, through dynamic programming, although the computational challenge of computing them for all but the smallest problems currently seem insurmountable. The details of such a dynamic programming formulation may be found in Frazier et al. (2008).

Due to the difficulty of computing the value of information in the sequential setting, a number of sequential policies make decisions based upon the value of information for single-stage allocations. These policies include the KG policy and the LL(S) policy of Chick and Inoue (2001). The single-stage value of an allocation n is the $v(n)$ defined earlier, where the prior is replaced by the current posterior belief. Of course, the value of information for single stage allocations is only an approximation to the value of information for sequential sampling, and so policies that maximize the single-stage value of information with each allocation are not optimal overall.

In this section, to match the usage of sequential policies, the value $v(n)$ of an allocation $n \in \mathbb{Z}_+^M$ will be computed under the current posterior, rather than under the prior. That is,

$$v(n) = \mathbf{E}_t \left[\max_i \mathbf{E}_t[\theta_i | Y, n] | n \right] - \max_i \mathbf{E}_t[\theta_i], \quad (8)$$

where t is the number of measurements observed so far, and \mathbf{E}_t is the conditional expectation with respect to these observations. Thus, $v(n)$ is implicitly a function of the information collected so far, but we suppress this dependence in the notation.

7.1 KG and KG(*) Policies

As described above, several sequential policies use the single-stage value of information v to make allocation decisions, including the KG policy. The KG policy is defined as the policy that measures, at each point in time, the alternative $\arg \max_x v(e_x)$. The quantity $v(e_x)$ is called the KG factor, and the motivation for the KG policy is that a measurement's KG factor, which is its single-stage value, approximates its sequential value.

The KG policy is optimal when only a single measurement remains in the budget ($N = 1$). The previous discussion, however, establishes that the optimal single-stage allocation may

vary greatly depending on the budget, and the best alternatives to measure when the budget is small may be far from best when the budget is large. Thus, it is natural to be concerned about problems induced in the KG policy by this effect. From a convergence result in Frazier et al. (2008), we can be sure that the effect is not severe enough to prevent convergence to the best alternative given large enough N , but we may be concerned about performance at intermediate values of N .

Observe that $v(me_x)/m$ is the average benefit per measurement obtained by measuring alternative x a total of m times. One fundamental problem encountered by the KG policy is that the KG factor, which is $v(e_x)$, may be many orders of magnitude smaller than the average benefit of some larger number of measurements. This can mislead the policy into making poor decisions, particularly when some alternatives have this scaling problem and others do not.

With these considerations in mind, we propose a policy based upon the KG policy that considers the value of multiple measurements of the same alternative. We call this policy KG(*). The value of measuring an alternative is approximated by the best average value that can be obtained from measuring it a fixed number of times.

$$x^n \in \arg \max_x \max_{1 \leq m \leq N-n} v(me_x)/m. \tag{9}$$

We call the quantity $\max_{1 \leq m \leq N-n} v(me_x)/m$ the KG(*) factor. This maximum may either be taken over the integers $\{1, \dots, N - n\}$, or over the entire continuous interval $[1, N - n]$. The KG(*) policy first appeared in Chick and Frazier (2009a,b), where it was described briefly and used in a numerical study. The current article describes and justifies this policy in a more complete way.

The following theorem shows that the KG(*) policy retains the theoretical guarantees of the KG policy of optimality at $N = 1$ and $N = \infty$ shown in Frazier et al. (2008), and in Section 7.3 we see that it improves upon the KG policy's performance at other values of N .

Theorem 3. *The sequential value of the KG(*) policy is equal to that of the optimal sequential policy when $N = 1$, and also in the limit as $N \rightarrow \infty$.*

The statement that the KG(*) policy and the optimal sequential policy have asymptoti-

cally equal values in the limit as $N \rightarrow \infty$ is equivalent to the statement that, as the number of measurements grows large, we discover and implement the best alternative with probability under the prior growing to 1. In Frazier and Powell (2009) this is known as “convergence to a global optimum.”

The effect of considering this best average value measuring a single alternative is to smooth the non-concavity in the value of measurement. It allows the policy to value an alternative according to our ability to measure it within the budget, rather than myopically assuming this budget is 1. This valuation is still only an approximation to the true sequential value of information computed using the fact that the policy may behave in a fully sequential way. Nevertheless it is a better approximation than the KG policy, avoiding some problems caused by the non-concavity of the value of information.

Figure 7 illustrates the KG(*) policy. The situation pictured is one with three alternatives, where the first alternative has $\mu_1 = -1$, $\sigma_1 = 1.1$, the second alternative has $\mu_2 = 0$, $\sigma_2 = 0.003$ and the third alternative has a known value of 0. A sampling variance of $\lambda = 10$ is common across the alternatives. The solid line without circles (henceforth called the “solid” line) shows the single-stage value of sampling from the first alternative as a function of the number of samples taken, while the solid line with circles (the “circled” line) shows this value for sampling from the second alternative. The dotted line starts at the origin and is tangent to the solid line. The right panel is a magnification of the left panel.

The decision of the KG policy is given by the value of a single sample allocated to a single alternative. This quantity is seen in the right panel as the solid and circled lines at $n_x = 1$. Because the circled line (corresponding to alternative 2) is higher than the solid line at $n_x = 1$, the KG policy measures the second alternative. It does so despite the fact that the left panel shows that much more value can be obtained from repeated measurements of alternative 1 than can from alternative 2. Indeed, with a miniscule variance of 0.003, measuring this alternative is worth very little — 25 measurements allocated in a single stage provide an expected value of less than .002, while the left panel shows that 25 measurements of the other unknown alternative provide a value that is 30 times larger.

The KG(*) factor may be computed from the dotted line, which begins at the origin and

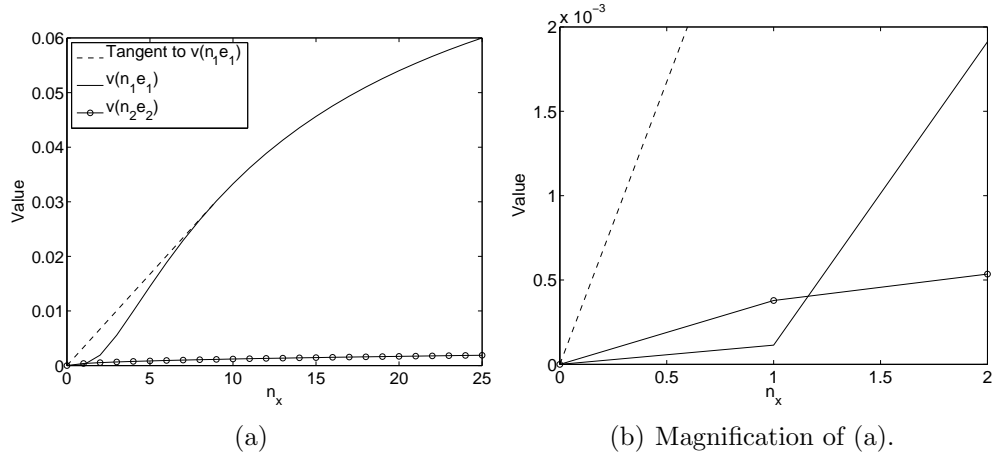


Figure 7: The value of information for single-stage sampling, as a function of the number of measurements. There are three alternatives: $\mu_1 = -1$, $\sigma_1 = 1.1$; $\mu_2 = 0$, $\sigma_2 = 0.003$; $\mu_3 = 0$, $\sigma_0 = 0$. The sampling variance is $\lambda = 10$. Solid lines without and with circles show the value of sampling from alternatives 1 and 2 respectively. The dotted line is tangent to the solid line without circles, and illustrates the KG^* value of sampling from alternative 1.

lies tangent to the $v(n_x)$ curve. Observe that the slope of a line that begins at the origin and intersects $v(n_x)$ at a point $n_x = m$ is $v(m)/m$. Since the dotted line is tangent to $v(n_x)$, its slope is highest among all such lines intersecting $v(n_x)$, and so is equal to $\max_{m \geq 0} v(m)/m$. If the point of intersection is between 1 and N , as it is for the dotted line in Figure 7, then its slope is the KG^* factor evaluated over a continuous rather than discrete range. Furthermore, the value of the dotted line at $n_x = 1$ is its slope, and so the KG^* factor for alternative 2 is given by the dotted line at $n_x = 1$. This parallels the KG factor for alternative 2, which is equal to the solid line at $n_x = 1$.

In contrast to the solid line, the circled line is concave (as implied by Corollary 1 and $\mu_2 = \max_x \mu_x$), and the KG^* factor for this alternative is the same as the KG factor. From this, we see that the KG^* policy measures the first alternative, which is clearly a better decision than the KG policy's decision to measure the second. We revisit a similar example later in Section 7.3.

Observe that the curve resulting from following the dotted line up to the point of intersection, and the solid line afterward, is concave. Furthermore, it is the concave envelope of the value of information curve. With this understanding, we see that the KG^* policy chooses

measurements to maximize a “concavified” version of the value of measurement. As in the case with the solid line, if a single measurement has very little value on its own, but repeated measurement will have significant value, then the KG^* factor captures this fact. It does so by dividing the eventually significant value that may be obtained among the number of measurements necessary to obtain it, crediting each measurement with the portion that it contributes to the whole. On the other hand, if a single measurement already puts us in the region where the value of measurement is concave, then the KG^* factor is the same as the KG factor. Also observe that if the remaining budget is so small that it prevents us from reaching what would otherwise be the tangent point (the number of measurements giving the largest average value per measurement), then the KG^* factor gives only the best average value per measurement obtainable with the remaining budget. In particular, when $N = 1$, the KG^* policy is the same as the KG policy.

Although the KG^* policy addresses the non-concavity of the value of information and the problems it causes for the KG policy, the KG^* factors it uses remain only approximations to the value of information for sequential sampling. They ignore that one may sample more than one alternative in the future, and also that future decisions of how many times to sample an alternative may be made sequentially rather than in a single stage. Nevertheless, Section 7.3 shows that it significantly outperforms the KG policy in some cases, which is significant because the KG policy performed well when compared to other fully sequential policies in an empirical study (Frazier et al., 2008).

7.2 Computing the KG^* Policy

To compute the KG^* policy, we need to find the maximum in (9). This is facilitated by the following result. The lower bound \underline{m} from this result was derived as an approximation, but not shown to be an upper bound, in Chick and Frazier (2009b).

Proposition 3. *If $\Delta_x = 0$ then $m \mapsto v(me_x)/m$ is strictly decreasing on \mathbf{R}_{++} . If $\Delta_x \neq 0$, then $m \mapsto v(me_x)/m$ is strictly unimodal over \mathbf{R}_{++} and its unique maximum m^* satisfies*

$\underline{m} \leq m^* \leq \bar{m}$, where

$$\underline{m} = \frac{\lambda_x}{4\sigma_x^2} \left(-1 + r + \sqrt{1 + 6r + r^2} \right),$$

$$\bar{m} = \frac{\lambda_x}{4\sigma_x^2} \left(1 + r + \sqrt{1 + 10r + r^2} \right),$$

and where $r = \Delta_x^2 / \sigma_x^2$.

This proposition has the following immediate consequences for solutions to the constrained optimization problems $\arg \max_{m=1, \dots, N-n} v(me_x)/m$ and $\arg \max_{m \in [1, N-n]} v(me_x)/m$.

- If $N - n \leq \underline{m}$ then

$$\arg \max_{m=1, \dots, N-n} v(me_x)/m = \arg \max_{m \in [1, N-n]} v(me_x)/m = N - n.$$

- If $\Delta_x = 0$ or $1 \geq \bar{m}$ then

$$\arg \max_{m=1, \dots, N-n} v(me_x)/m = \arg \max_{m \in [1, N-n]} v(me_x)/m = 1.$$

- Otherwise,

$$\arg \max_{m=1, \dots, N-n} v(me_x)/m \in [\max(1, \underline{m}), \min(N - n, \bar{m})],$$

$$\arg \max_{m \in [1, N-n]} v(me_x)/m \in \{\max(1, \text{floor}(\underline{m})), \dots, \min(N - n, \text{ceil}(\bar{m}))\}.$$

In the last case, we have several options. In the integer-constrained optimization problem, we may simply evaluate the objective function at each integer in the range. Or, in both the integer-constrained and continuous problems, we may find m^* numerically using a line-search on the strictly decreasing function h defined in the proof of Proposition 3. This function h has its only root at m^* . The value of m^* then immediately provides the solution to either constrained optimization problem. Or, one may simply approximate the solution, for example, by taking the average of the upper and lower bounds \underline{m} and \bar{m} . In practice, such an approximation is computationally convenient and seems to cause little degradation in performance.

7.3 Computational Experiments

It is instructive to compare the decisions made by KG and KG(*) policies on a simple problem. Consider three alternatives, where one is known perfectly, a great deal is known about another, and less is known about the last alternative. Furthermore, suppose that the two well-known alternatives are equal in value, while we estimate the last alternative to be worse. This example may be written as $\sigma_1 = 0$, $\sigma_2 \approx 0$, $\sigma_3 > 0$, $\mu_1 = 0$, $\mu_2 = 0$, $\mu_3 < 0$.

In this problem, because alternative 1 is perfectly known, there is no value in measuring it. The question we then face is whether to measure alternative 2 or 3. The value of measuring alternative 2 is quite low because it has a very low variance – even though measuring it may reveal it to be better than alternative 1, the amount by which it may improve upon alternative 1 is very likely to be small. In contrast, if σ_3 is large enough, and if μ_3 is not too small, then alternative 3 may be significantly better than both alternatives 1 and 2 and there may be significant value in measuring it. Furthermore, for larger values of either μ_3 or σ_3 , there is a greater potential for discovering substantial value in alternative 3, and measuring it should have greater value.

In Figure 8(a) we fix $\sigma_2 = 10^{-2}$ and $\mu_2 = 0$, and examine the decisions of the KG and KG(*) policies as a function of μ_3 and σ_3 . For each policy a curve is plotted, and the region above and to the right contains those priors for which the policy measures alternative 3 first. Below and to the left of each curve are those priors for which alternative 3 is less attractive, and the policy instead measures alternative 2.

The figure shows that KG(*) is more willing than KG to measure alternative 3, the alternative with large variance. Indeed, KG’s tendency to measure alternative 2 seems extreme and counterproductive. Even if we were to learn its value perfectly, we would only obtain a reward of

$$\mathbf{E}[\max(\theta_1, \theta_2)] = \mathbf{E}[\max(0, \theta_2)] = \sigma_2 f(\mu_2/\sigma_2) = \sigma_2 f(0) = 10^{-2} \frac{1}{\sqrt{2\pi}} \approx 0.004.$$

In reality, the reward resulting from measuring alternative 2 is even smaller because of measurement noise, and because later measurements of alternative 3 may reveal it to be better than either alternative 2 or alternative 1.

Contrasting the low value of measuring alternative 2, the ultimate value of measuring alternative 3 can be quite large. For example, consider the case $\mu_3 = -1$ and $\sigma_3 = \sqrt{2}$. In this case, learning the value of alternative 3 perfectly (without learning about alternative 2), has value


$$\mathbb{E}[\max(0, \theta_3)] = \sigma_3 f(\mu_3/\sigma_3) = \sqrt{2} f(-1/\sqrt{2}) \approx 0.2.$$

This value is nearly two orders of magnitude larger than the value of learning alternative 2 perfectly. Yet, the KG policy chooses to measure alternative 2.

Of course, one does not learn the value of an alternative perfectly when one measures it. In the calculation behind the KG policy, we assume that only one measurement will be allowed, and under this assumption measuring alternative 2 is best. This is because of the non-concavity induced by the large measurement variance ($\lambda = 100$). Contrasting the myopically optimal action, if the budget is sufficiently large, it is prudent to first measure alternative 3 several times to determine whether its value is significantly better than 0 (the true value of alternative 1, and the initial estimate of the value of alternative 2), and then later, if sufficient budget remains, go back and learn more about alternative 2. It is in learning about alternative 3 that the most significant benefits are to be found.

Figure 8(b) shows the performance of the KG and KG(*) policies when $\mu_3 = -1$ and $\sigma_3 = \sqrt{2}$. The average opportunity cost over 1000 independent simulations of the two policies is plotted, where opportunity cost at a time n is the difference in value between the best decision that can be made with perfect information, and the best decision that can be made with the available information. This opportunity cost is written $\max_i \theta_i - \theta_{\arg \max_j \mu_j^n}$.

The KG policy measures alternative 2 through every iteration pictured, from 1 to 1000. Thus, it reduces its opportunity cost by a small amount when $\theta_2 > \theta_1$, but it misses the large benefit of learning about alternative 3. Eventually, under the KG policy, the variance of our belief about alternative 2 shrinks enough that we measure alternative 3, but in this example, the measurement budget ends before this eventuality is reached. In contrast, the KG(*) policy measures alternative 3 immediately, and sees a large reduction in opportunity cost as it learns whether θ_3 or θ_1 is better.



(a) Measurement decisions of KG and KG(*) for $\mu = [0, 0, \mu_3]$, $\sigma = [0, 10^{-2}, \sigma_3]$, $\lambda = 100$. The region above each algorithm's curve contains those values of μ_3 and σ_3 for which the algorithm measures alternative 3. The region below contains those values for which the algorithm measures alternative 2.

(b) Opportunity cost vs. number of iterations for KG and KG(*) for $\mu = [0, 0, -1]$, $\sigma = [0, 10^{-2}, \sqrt{2}]$, $\lambda = 100$. Plotted is the average opportunity cost over 1000 independent runs. Opportunity cost at time n is $\max_i \theta_i - \theta_{\arg \max_j \mu_j^n}$.

Figure 8: Comparison of KG and KG(*) for a problem with three measurement choices.

8 Conclusion

Managers realize intuitively that it is typically the case that they cannot evaluate all alternatives when faced with a decision. Often referred to as “analysis paralysis” they realize that they have to do a good job with a reasonable set of choices. We revisit the nonconcavity property in the value of information, itself a surprising behavior for many, and investigate its implications in terms of determining how to allocate limited resources for collecting information. In particular, we show that there is an optimal number of choices which depends on the measurement budget and the level of uncertainty in a measurement.

An effective policy for determining how to collect information sequentially is one which maximizes the expected value of a single measurement, sometimes referred to as the knowledge gradient policy. We show, however, that such a policy can perform very poorly in the presence of nonconcavity in the value of information. As an alternative, we introduce the KG(*) policy which chooses to measure the alternative which offers the highest average value when we are allowed to observe an alternative multiple times. This policy overcomes a major limitation of the original knowledge gradient policy for problems where a single observation is noisy and yields limited information. The KG(*) policy is effectively allowing us to pay

for a more precise measurement.

A Proofs

A.1 Derivative Expressions

The following lemmas give explicit expressions for the first and second derivatives of $\tilde{\sigma}_i$ and g needed in other proofs.

Lemma 2. *The first and second derivatives of $\tilde{\sigma}_i$ are given by*

$$\begin{aligned}\tilde{\sigma}'_i(n_i) &= \frac{\lambda_i}{2\sigma_i} n_i^{-1/2} \left(\frac{\lambda_i}{\sigma_i^2} + n_i \right)^{-3/2}, \\ \tilde{\sigma}''_i(n_i) &= -\frac{\lambda_i}{4\sigma_i} n_i^{-3/2} \left(\frac{\lambda_i}{\sigma_i^2} + n_i \right)^{-5/2} \left(\frac{\lambda_i}{\sigma_i^2} + 4n_i \right).\end{aligned}$$

The second derivative is negative, and so $\tilde{\sigma}_i$ is concave.

Proof. Define $r = \lambda_i/\sigma_i^2$. Then $\tilde{\sigma}_i(n_i) = \sqrt{\sigma_i^2 n_i / (r + n_i)}$. We compute the first derivative as

$$\tilde{\sigma}'_i(n_i) = \sigma_i \frac{1}{2} \left[\frac{n_i}{r + n_i} \right]^{-1/2} \frac{r}{(r + n_i)^2} = \frac{\sigma_i r}{2} (n_i(r + n_i)^3)^{-1/2}.$$

We compute the second derivative from the first derivative as

$$\tilde{\sigma}''_i(n_i) = \frac{\sigma_i r}{2} \left[-\frac{1}{2} (n_i(r + n_i)^3)^{-3/2} \right] [(r + n_i)^3 + 3n_i(r + n_i)^2] = -\frac{\sigma_i r}{4} \frac{r + 4n_i}{n_i^{3/2} (r + n_i)^{5/2}}.$$

□

Lemma 3.

$$\frac{\partial g}{\partial s_i}(s) = \begin{cases} \mathbf{E} \left[\varphi \left(\frac{W - \mu_i}{s_i} \right) \right], & \text{if } s_i > 0, \\ \varphi(0), & \text{if } s = 0 \text{ and } \mu_i = \max_{j \neq i} \mu_j, \\ 0, & \text{otherwise.} \end{cases}$$

where $W = \max_{j \neq i} \mu_j + s_j Z_j$.

Proof. We first calculate the right partial derivative of g with respect to s_i . If the limit exists, this right partial derivative is

$$\frac{\partial^+ g}{\partial s_i}(s) = \lim_{\epsilon \rightarrow 0^+} \frac{g(s + \epsilon e_i) - g(s)}{\epsilon} = \lim_{\epsilon \rightarrow 0^+} \mathbf{E} \left[\frac{\max(\mu_i + (s_i + \epsilon)Z_i, W) - \max(\mu_i + s_i Z_i, W)}{\epsilon} \right].$$

Since $|\max(a, c) - \max(b, c)| \leq |a - b|$ for $a, b, c \in \mathbf{R}$, the term inside the expectation is dominated by $|(s_i + \epsilon)Z_i - s_i Z_i|/\epsilon = |Z_i|$, which is integrable. Thus, by the dominated convergence theorem, the limit and expectation may be interchanged, and

$$\begin{aligned} \frac{\partial^+ g}{\partial s_i}(s) &= \lim_{\epsilon \rightarrow 0^+} \frac{g(s + \epsilon e_i) - g(s)}{\epsilon} = \mathbf{E} \left[\frac{\partial^+}{\partial s_i} \max(\mu_i + s_i Z_i, W) \right] \\ &= \mathbf{E} [Z_i \mathbf{1}_{\{\mu_i + s_i Z_i > W\}} + Z_i^+ \mathbf{1}_{\{\mu_i + s_i Z_i = W\}}]. \end{aligned}$$

Thus $\frac{\partial^+ g}{\partial s_i}(s)$ exists and is given by this expression.

We now consider two cases, $s_i = 0$ and $s_i \neq 0$. First, consider the case $s_i = 0$. In this case, the right derivative $\frac{\partial^+ g}{\partial s_i}(s)$ is equal to the derivative $\frac{\partial g}{\partial s_i}(s)$ because s_i is at the left edge of the domain \mathbf{R}_+^M of g . Furthermore, $\mathbf{E} [Z_i \mathbf{1}_{\{\mu_i + s_i Z_i > W\}}] = \mathbf{E} [Z_i \mathbf{1}_{\{\mu_i > W\}}] = \mathbf{E} [Z_i] \mathbb{P} \{\mu_i > W\} = 0$, where the first equality follows by substituting $s_i = 0$, the second equality follows from the independence of Z_i and W , and the third equality follows from $\mathbf{E} [Z_i] = 0$. Thus,

$$\frac{\partial g}{\partial s_i}(s) = \frac{\partial^+ g}{\partial s_i}(s) = \mathbf{E} [Z_i^+ \mathbf{1}_{\{\mu_i + s_i Z_i = W\}}] = \begin{cases} 0, & \text{if } s_i = 0, s \neq 0, \\ \frac{1}{\sqrt{2\pi}} \mathbf{1}_{\{\mu_i = W\}}, & \text{if } s = 0. \end{cases}$$

where we have used $\mathbb{P} \{\mu_i + s_i Z_i = W\} = \mathbb{P} \{\mu_i = W\} = 0$ to show the expression for $s_i = 0, s \neq 0$, and the fact that $\mathbf{1}_{\{\mu_i + s_i Z_i = W\}} = \mathbf{1}_{\{\mu_i = \max_{j \neq i} \mu_j\}}$ is deterministic together with $\mathbf{E}[Z_i^+] = \varphi(0) = 1/\sqrt{2\pi}$ (from Remark 1) to show the expression for $s = 0$. This shows the lemma for the case $s_i = 0$.

We now consider the case $s_i > 0$. In this case, $\mathbb{P} \{\mu_i + s_i Z_i = W\} = 0$ and $\frac{\partial^+ g}{\partial s_i}(s) = \mathbf{E} [Z_i \mathbf{1}_{\{\mu_i + s_i Z_i > W\}}]$. Then an argument similar to the one above, in which integrability is shown and then the dominated convergence theorem applied, shows

$$\begin{aligned} \frac{\partial^- g}{\partial s_i}(s) &= \mathbf{E} \left[\frac{\partial^-}{\partial s_i} \max(\mu_i + s_i Z_i, W) \right] = \mathbf{E} [Z_i \mathbf{1}_{\{\mu_i + s_i Z_i > W\}} + Z_i^- \mathbf{1}_{\{\mu_i + s_i Z_i = W\}}] \\ &= \mathbf{E} [Z_i \mathbf{1}_{\{\mu_i + s_i Z_i > W\}}], \end{aligned}$$

where the last equality follows from $\mathbb{P} \{\mu_i + s_i Z_i = W\} = 0$. The equality of left and right derivatives implies their equality with $\frac{\partial g}{\partial s_i}(s)$.

Applying the tower property provides

$$\frac{\partial g}{\partial s_i}(s) = \mathbf{E} [\mathbf{E} [Z_i \mathbf{1}_{\{\mu_i + s_i Z_i > W\}} \mid W]] = \mathbf{E} \left[\mathbf{E} \left[Z_i \mathbf{1}_{\left\{Z_i > \frac{W - \mu_i}{s_i}\right\}} \mid W \right] \right] = \mathbf{E} \left[\varphi \left(\frac{W - \mu_i}{s_i} \right) \right],$$

where the last equality follows from Remark 1. This shows the lemma for the case $s_i > 0$. \square

Lemma 4. For $s \neq 0$ and $i \neq j$,

$$\begin{aligned}\frac{\partial^2 g}{\partial s_i^2} &= \mathbb{E} \left[\left(\frac{W - \mu_i}{s_i^3} \right) \varphi \left(\frac{W - \mu_i}{s_i} \right) \right], \\ \frac{\partial^2 g}{\partial s_i \partial s_j} &= \mathbb{E} \left[\left(\frac{-W + \mu_i}{s_i} \right) \varphi \left(\frac{W - \mu_i}{s_i} \right) Z_j \mathbf{1}_{\{\mu_j + s_j Z_j \geq W'\}} \right],\end{aligned}$$

where $W = \max_{k \neq i} \mu_k + s_k Z_k$ and $W' = \max_{k \neq i, j} \mu_k + s_k Z_k$. Furthermore, these derivatives are continuous in s for all $s \neq 0$. Thus, $\nabla g(s)$ and $\nabla^2 g(s)$ exist and are continuous at all $s \neq 0$.

Proof. Let $s \neq 0$. We first show the claimed expression for $\frac{\partial^2 g}{\partial s_i^2}$. By Lemma 3,

$$\frac{\partial^2 g}{\partial s_i^2} = \frac{\partial}{\partial s_i} \mathbb{E} \left[\varphi \left(\frac{W - \mu_i}{s_i} \right) \right] = \lim_{u \rightarrow s_i} \mathbb{E} \left[\frac{1}{u - s_i} \left(\varphi \left(\frac{W - \mu_i}{u} \right) - \varphi \left(\frac{W - \mu_i}{s_i} \right) \right) \right]. \quad (10)$$

Consider the open ball $B = (s_i/2, 3s_i/2)$ containing s_i . We now bound over $u \in B$ the expression inside the expectation in (10). For all $u \in B$,

$$\frac{1}{|u - s_i|} \left| \varphi \left(\frac{W - \mu_i}{u} \right) - \varphi \left(\frac{W - \mu_i}{s_i} \right) \right| \leq \sup_{u' \in B} \left| \frac{\partial}{\partial u'} \varphi \left(\frac{W - \mu_i}{u'} \right) \right| \leq \left| \frac{W - \mu_i}{(u')^3} \right| \varphi(0),$$

where we note that $\frac{\partial}{\partial u'} \varphi \left(\frac{W - \mu_i}{u'} \right) = \frac{W - \mu_i}{(u')^3} \varphi \left(\frac{W - \mu_i}{u'} \right)$. This bound is integrable, and so by the dominated convergence theorem we may exchange the limit and expectation in (10) to obtain

$$\frac{\partial^2 g}{\partial s_i^2} = \mathbb{E} \left[\frac{\partial}{\partial s_i} \varphi \left(\frac{W - \mu_i}{s_i} \right) \right] = \mathbb{E} \left[\frac{W - \mu_i}{s_i^3} \varphi \left(\frac{W - \mu_i}{s_i} \right) \right],$$

which is the claimed expression for $\frac{\partial^2 g}{\partial s_i^2}$.

We now show the claimed expression for $\frac{\partial^2 g}{\partial s_i \partial s_j}$. By Lemma 3,

$$\frac{\partial^+}{\partial s_j} \left(\frac{\partial g}{\partial s_i} \right) = \frac{\partial^+}{\partial s_j} \mathbb{E} \left[\varphi \left(\frac{W - \mu_i}{s_i} \right) \right] = \lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} \mathbb{E} \left[\varphi \left(\frac{W_\epsilon - \mu_i}{s_i} \right) - \varphi \left(\frac{W - \mu_i}{s_i} \right) \right] \quad (11)$$

where $W_\epsilon = \max \{ \max_{k \neq i, j} \mu_k + s_k Z_k, \mu_j + (s_j + \epsilon) Z_j \}$.

For any $a > 0$ and $\epsilon' \in (0, a)$,

$$\begin{aligned}\frac{1}{\epsilon'} \left| \varphi \left(\frac{W_{\epsilon'} - \mu_i}{s_i} \right) - \varphi \left(\frac{W - \mu_i}{s_i} \right) \right| &\leq \sup_{\epsilon \in (0, a)} \left| \frac{\partial^+}{\partial \epsilon} \varphi \left(\frac{W_\epsilon - \mu_i}{s_i} \right) \right| \\ &\leq \sup_{\epsilon \in (0, a)} \frac{|W_\epsilon - \mu_i|}{s_i} |Z_j| \varphi(0) = (|W - \mu_i| + a|Z_j|) |Z_j| \varphi(0) / s_i,\end{aligned}$$

where we have used

$$\frac{\partial^+}{\partial \epsilon} \varphi \left(\frac{W_\epsilon - \mu_i}{s_i} \right) = \left(\frac{-W_\epsilon + \mu_i}{s_i} \right) \varphi \left(\frac{W_\epsilon - \mu_i}{s_i} \right) Z_j \mathbf{1}_{\{\mu_j + (s_j + \epsilon) Z_j \geq W'\}}.$$

The bound $(|W - \mu_i| + a|Z_j|)|Z_j|\varphi(0)/s_i$ on the integrand of (11) is integrable, and so by the dominated convergence theorem we may exchange the limit and expectation to rewrite (11) as

$$\frac{\partial^+}{\partial s_j} \left(\frac{\partial g}{\partial s_i} \right) = \mathbb{E} \left[\frac{\partial^+}{\partial s_j} \varphi \left(\frac{W - \mu_i}{s_i} \right) \right] = \mathbb{E} \left[\left(\frac{-W + \mu_i}{s_i} \right) \varphi \left(\frac{W - \mu_i}{s_i} \right) Z_j \mathbf{1}_{\{\mu_j + s_j Z_j \geq W'\}} \right].$$

Using a similar argument, in which we use the dominated convergence theorem to exchange a limit with an integral, we may calculate the left derivative as

$$\frac{\partial^-}{\partial s_j} \left(\frac{\partial g}{\partial s_i} \right) = \mathbb{E} \left[\frac{\partial^+}{\partial s_j} \varphi \left(\frac{W - \mu_i}{s_i} \right) \right] = \mathbb{E} \left[\left(\frac{-W + \mu_i}{s_i} \right) \varphi \left(\frac{W - \mu_i}{s_i} \right) Z_j \mathbf{1}_{\{\mu_j + s_j Z_j > W'\}} \right].$$

The only difference is that the inequality in the indicator function is now strict, because the corresponding expression for $\frac{\partial^-}{\partial \epsilon} \varphi \left(\frac{W_\epsilon - \mu_i}{s_i} \right)$ has a strict inequality in the indicator function.

Then, $s \neq 0$ implies $\mathbb{P}\{\mu_j + s_j Z_j = W'\} = 0$, and

$$\frac{\partial^+}{\partial s_j} \left(\frac{\partial g}{\partial s_i} \right) = \frac{\partial^-}{\partial s_j} \left(\frac{\partial g}{\partial s_i} \right) = \mathbb{E} \left[\left(\frac{-W + \mu_i}{s_i} \right) \varphi \left(\frac{W - \mu_i}{s_i} \right) Z_j \mathbf{1}_{\{\mu_j + s_j Z_j \geq W'\}} \right]$$

The left and right derivatives agree, and so are equal to the derivative $\frac{\partial^2 g}{\partial s_i \partial s_j}$.

We now show continuity of $\frac{\partial^2 g}{\partial s_i^2}$ and $\frac{\partial^2 g}{\partial s_i \partial s_j}$. For $u \in \mathbb{R}^M$, let $W_u = \max_{k \neq i} \mu_k + u_k Z_k$ and $W'_u = \max_{k \neq i, j} \mu_k + u_k Z_k$.

First, consider $\frac{\partial^2 g}{\partial s_i^2}$. We have

$$\begin{aligned} \lim_{u \rightarrow s} \frac{\partial^2 g}{\partial s_i^2}(u) &= \lim_{u \rightarrow s} \mathbb{E} \left[\left(\frac{W_u - \mu_i}{u_i^3} \right) \varphi \left(\frac{W_u - \mu_i}{u_i} \right) \right] = \mathbb{E} \left[\lim_{u \rightarrow s} \left(\frac{W_u - \mu_i}{u_i^3} \right) \varphi \left(\frac{W_u - \mu_i}{u_i} \right) \right], \\ &= \mathbb{E} \left[\left(\frac{W - \mu_i}{s_i^3} \right) \varphi \left(\frac{W - \mu_i}{s_i} \right) \right] = \frac{\partial^2 g}{\partial s_i^2}(s), \end{aligned}$$

where the exchange of limit and expectation follows from the dominated convergence theorem and the following integrable bound satisfied by all u within an ϵ -ball of s , with $\epsilon < s_i$,

$$\left| \left(\frac{W_u - \mu_i}{u_i^3} \right) \varphi \left(\frac{W_u - \mu_i}{u_i} \right) \right| \leq \left| \frac{W_u - \mu_i}{u_i^3} \right| \varphi(0) \leq \frac{|W - \mu_i| + \epsilon \max_{k \neq i} |Z_k|}{(s_i - \epsilon)^3} \varphi(0).$$

Now consider $\frac{\partial^2 g}{\partial s_i \partial s_j}$. We have

$$\begin{aligned} \lim_{u \rightarrow s} \frac{\partial^2 g}{\partial s_i \partial s_j}(u) &= \lim_{u \rightarrow s} \mathbb{E} \left[\left(\frac{-W_u + \mu_i}{u_i} \right) \varphi \left(\frac{W_u - \mu_i}{u_i} \right) Z_j \mathbf{1}_{\{\mu_j + u_j Z_j \geq W'_u\}} \right] \\ &= \mathbb{E} \left[\lim_{u \rightarrow s} \left(\frac{-W_u + \mu_i}{u_i} \right) \varphi \left(\frac{W_u - \mu_i}{u_i} \right) Z_j \mathbf{1}_{\{\mu_j + u_j Z_j \geq W'_u\}} \right] \\ &= \mathbb{E} \left[\left(\frac{-W + \mu_i}{s_i} \right) \varphi \left(\frac{W - \mu_i}{s_i} \right) Z_j \mathbf{1}_{\{\mu_j + s_j Z_j \geq W'\}} \right] = \frac{\partial^2 g}{\partial s_i \partial s_j}(s), \end{aligned}$$

where the exchange of limit and expectation follows from the dominated convergence theorem, and the following integrable bound is satisfied by all u within an ϵ -ball of s , with $\epsilon < s_i$,

$$\begin{aligned} \left| \left(\frac{-W_u + \mu_i}{u_i} \right) \varphi \left(\frac{W_u - \mu_i}{u_i} \right) Z_j \mathbf{1}_{\{\mu_j + u_j Z_j \geq W'_u\}} \right| &\leq \left| \frac{-W_u + \mu_i}{u_i} \right| \varphi(0) |Z_j| \\ &\leq \frac{|-W + \mu_i| + \epsilon \max_{k \neq i} |Z_k|}{s_i - \epsilon} \varphi(0) |Z_j| \end{aligned}$$

In taking the limit, we have also noted that the event $\{\mu_j + s_j Z_j = W'\}$ has probability 0, and so the event $\lim_{u \rightarrow s} \mathbf{1}_{\{\mu_j + u_j Z_j \geq W'_u\}} = \mathbf{1}_{\{\mu_j + s_j Z_j \geq W'\}}$ has probability 1. \square

Proof of Proposition 1

First, recall that $v(ne_i) = g(\tilde{\sigma}_i(n)e_i)$. Thus, to show the claimed expressions for $v(ne_i)$ and its derivative, it is enough to provide expressions for $g(se_i)$ and its derivative.

For $s > 0$, by Lemma 3,

$$\frac{\partial}{\partial s} g(se_i) = \varphi \left(\frac{\max_{j \neq i} \mu_j - \mu_i}{s} \right) = \varphi \left(\frac{\Delta_i}{s} \right),$$

where we have used the symmetry of φ . This expression together with the chain rule shows the claimed expression for $\frac{\partial}{\partial n} v(ne_i)$.

We now provide an analytic expression for $g(se_i)$. For $s = 0$, $g(se_i) = 0$. For $s > 0$,

$$\begin{aligned} g(se_i) &= \mathbb{E} \left[\max(\mu_i + sZ_i, \max_{j \neq i} \mu_j) \right] - \max_j \mu_j \\ &= \begin{cases} \mathbb{E} [\max(sZ_i, -\Delta_i)] & \text{if } \mu_i = \max_j \mu_j, \\ \mathbb{E} [\max(-\Delta_i + sZ_i, 0)] & \text{if } \mu_i < \max_j \mu_j, \end{cases} \\ &= \mathbb{E} [\max(sZ_i, -\Delta_i)] = \mathbb{E} [sZ_i \mathbf{1}_{\{sZ_i \geq -\Delta_i\}}] + \mathbb{E} [-\Delta_i \mathbf{1}_{\{sZ_i < -\Delta_i/s\}}] \\ &= s\varphi(-\Delta_i/s) - \Delta_i \Phi(-\Delta_i/s) = sf(-\Delta_i/s). \end{aligned}$$

The first equality on the third line follows from noting that $\mathbb{E}[sZ_i] = 0$, which allows us to subtract sZ_i from both terms in the maximum in the second case to obtain $\mathbb{E}[\max(-\Delta_i, -sZ_i)]$, which is equal to $\mathbb{E}[\max(-\Delta_i, sZ_i)]$ because of the equality in distribution between Z_i and $-Z_i$. The first equality on the fourth line uses Remark 1. Combining this expression for $g(se_i)$ with $v(ne_i) = g(\tilde{\sigma}_i(n)e_i)$ shows the claimed expression for $v(ne_i)$.

Proof of Proposition 2

Before proving Proposition 2, we state and prove the following lemma.

Lemma 5. *For $\Delta \in \mathbf{R}_+$ and any alternative i , the function*

$$n, z \mapsto \varphi(z) \max\left(0, \frac{-\Delta}{\tilde{\sigma}_i(n)} + z\right),$$

is log-concave on $\mathbf{R}_+ \times \mathbf{R}$, where $\log(0)$ is understood to be $-\infty$.

Proof. First,

$$x \mapsto \log(\max(0, x)) = \begin{cases} -\infty & \text{if } x \leq 0, \\ \log(x) & \text{if } x > 0, \end{cases}$$

is concave and non-decreasing.

Also, since $n \mapsto \tilde{\sigma}_i(n)$ is concave (by Lemma 2) and strictly positive when $n > 0$, we have that $n \mapsto -\Delta/\tilde{\sigma}_i(n)$ is concave on \mathbf{R}_{++} , and that

$$n \mapsto \begin{cases} -\frac{\Delta}{\tilde{\sigma}_i(n)} & \text{if } n > 0, \\ -\infty & \text{if } n \leq 0, \end{cases}$$

is concave on \mathbf{R} . Since the sum of two concave functions is concave,

$$n, z \mapsto \begin{cases} -\frac{\Delta}{\tilde{\sigma}_i(n)} + z & \text{if } n > 0, \\ -\infty & \text{if } n \leq 0, \end{cases}$$

is concave on \mathbf{R}^2 . Thus, by the composition rules for concave functions (Boyd and Vandenberghe, 2004),

$$n, z \mapsto \begin{cases} \log\left(\max\left(0, -\frac{\Delta}{\tilde{\sigma}_i(n)} + z\right)\right) & \text{if } n > 0, \\ -\infty & \text{if } n \leq 0, \end{cases}$$

is concave on \mathbf{R}^2 and hence also on $\mathbf{R}_+ \times \mathbf{R}$. □

We now prove Proposition 2.

Proof of Proposition 2. From the definition of v , (2), we have that

$$v(ne_i) = \mathbf{E} \left[\max \left(\mu_i + \tilde{\sigma}_i(n)Z_i, \max_{j \neq i} \mu_j \right) \right] \quad (12)$$

where Z_i is a standard normal random variable. Letting $\Delta = |\mu_i - \max_{j \neq i} \mu_j|$ we will see by considering two cases that

$$v(ne_i) = \mathbf{E} [\max(0, -\Delta + \tilde{\sigma}_i(n)Z_i)]. \quad (13)$$

In the first case, suppose $\mu_i \leq \max_{j \neq i} \mu_j$, and then (13) follows from (12) when we bring $\max_j \mu_j = \max_{j \neq i} \mu_j$ inside the expectation, and then inside both arguments of the inner maximum. In the second case, suppose $\mu_i > \max_{j \neq i} \mu_j$. Then subtracting and adding $\mu_i + \tilde{\sigma}_i(n)Z_i$ inside the expectation gives

$$\begin{aligned} v(ne_i) &= \mathbf{E} \left[\max \left(0, \max_{j \neq i} \mu_j - (\mu_i + \tilde{\sigma}_i(n)Z_i) \right) \right] + \mathbf{E} [\mu_i + \tilde{\sigma}_i(n)Z_i] - \max_j \mu_j \\ &= \mathbf{E} [\max(0, -\Delta - \tilde{\sigma}_i(n)Z_i)] = \mathbf{E} [\max(0, -\Delta + \tilde{\sigma}_i(n)Z_i)]. \end{aligned}$$

where in the second line we use that $\mathbf{E} [\tilde{\sigma}_i(n)Z_i] = 0$ and that $\mu_i - \max_j \mu_j = 0$, and in the third line we use that Z_i is equal in distribution to $-Z_i$. This shows (13).

Proceeding from (13), we have

$$v(ne_i) = \tilde{\sigma}_i(n) \mathbf{E} \left[\max \left(0, \frac{-\Delta}{\tilde{\sigma}_i(n)} + Z_i \right) \right].$$

The concavity of $\tilde{\sigma}_i$ shown in Lemma 2 implies that $n \mapsto \log(\tilde{\sigma}_i(n))$ is also concave (Boyd and Vandenberghe, 2004), and since

$$\log(v(ne_i)) = \log(\tilde{\sigma}_i) + \log \left(\mathbf{E} \left[\max \left(0, \frac{-\Delta}{\tilde{\sigma}_i(n)} + Z_i \right) \right] \right),$$

it is sufficient to show concavity on \mathbf{R}_+ of

$$n \mapsto \log \left(\mathbf{E} \left[\max \left(0, \frac{-\Delta}{\tilde{\sigma}_i(n)} + Z_i \right) \right] \right).$$

To show this, we begin with the log-concavity on $\mathbf{R}_+ \times \mathbf{R}$ of

$$n, z \mapsto \varphi(z) \max \left(0, \frac{-\Delta}{\tilde{\sigma}_i(n)} + z \right),$$

as shown in Lemma 5. This implies that

$$\mathbf{E} \left[\max \left(0, \frac{-\Delta}{\tilde{\sigma}_i(n)} + Z_i \right) \right] = \int_{\mathbf{R}} \varphi(z) \max \left(0, \frac{-\Delta}{\tilde{\sigma}_i(n)} + z \right) dz$$

is log-concave in n as the integral over \mathbf{R} of function that is log-concave in n and the integrand z (Boyd and Vandenberghe, 2004). \square

Proof of Lemma 1

Without loss of generality we may assume that $\max_i \mu_i = \mu_2$. If this is not the case we may simply rename the indices. Adding and subtracting $\mu_2 + \tilde{\sigma}_2(n_2)Z_2$ to (4) provides

$$\begin{aligned} v(n) &= \mathbf{E} \left[\max_i \mu_i + \tilde{\sigma}_i(n_i)Z_i \right] - \mu_2 \\ &= \mathbf{E} \left[(\mu_1 - \mu_2 + \tilde{\sigma}_1(n_1)Z_1 - \tilde{\sigma}_2(n_2)Z_2)^+ \right] + \mathbf{E} \left[\mu_2 + \tilde{\sigma}_2(n_2)Z_2 \right] - \mu_2 \\ &= \mathbf{E} \left[(\mu_1 - \mu_2 + \tilde{\sigma}_1(n_1)Z_1 - \tilde{\sigma}_2(n_2)Z_2)^+ \right]. \end{aligned}$$

Now, Z_1 and Z_2 are independent standard normal random variables, and so $\tilde{\sigma}_1(n_1)Z_1 - \tilde{\sigma}_2(n_2)Z_2$ is also normal with mean 0 and variance $\tilde{\sigma}_1^2(n_1) + \tilde{\sigma}_2^2(n_2) = s^2(n)$, and so is equal in distribution to $s(n)Z$. Thus,

$$v(n) = \mathbf{E} \left[(\mu_1 - \mu_2 + s(n)Z)^+ \right] = s(n) \mathbf{E} \left[\left(\frac{-|\mu_1 - \mu_2|}{s(n)} + Z \right)^+ \right] = s(n) f \left(\frac{-|\mu_1 - \mu_2|}{s(n)} \right).$$

where we recall from Proposition 1 that $\mathbf{E} [(c + Z)^+] = f(c)$ for $c \in \mathbf{R}$.

Proof of Proposition 3

Before proving Proposition 3, we state two lemmas. The first lemma defines a new quantity, called Mill's ratio (Ruben, 1963), which is the ratio $R(z) = \Phi(-z)/\varphi(z)$. The proof of this first lemma is due to Laplace (Laplace, 1820).

Lemma 6 (Laplace). *Mill's ratio, $R(z) = \Phi(-z)/\varphi(z)$, is bounded above and below by*

$$(z^2 + 2)/(z^3 + 3z) < R(z) < z/(z^2 + 1).$$

Proof. The proof is due to Laplace (Laplace, 1820), as described in (Ruben, 1963). \square

Lemma 7. $f(-z)/\varphi(z)$ is strictly decreasing in z , and $(z^2+3)^{-1} < f(-z)/\varphi(z) < (z^2+1)^{-1}$.

Proof. We begin by noting that the derivative of Mill's ratio may be written

$$\frac{\partial}{\partial z} R(z) = \frac{\partial}{\partial z} \frac{\Phi(-z)}{\varphi(z)} = \frac{-\varphi(z)^2 + z\Phi(-z)\varphi(z)}{\varphi(z)^2} = -1 + z \frac{\Phi(-z)}{\varphi(z)} = -1 + zR(z).$$

Using this expression, we may differentiate $f(-z)/\varphi(z)$ as

$$\frac{\partial}{\partial z} \frac{f(-z)}{\varphi(z)} = \frac{\partial}{\partial z} [1 - zR(z)] = -R(z) - z \frac{\partial}{\partial z} R(z) = -R(z)(1 + z^2) + z$$

which is strictly negative by the inequality $R(z) > z/(z^2 + 1)$ from Lemma 6. This shows that $z \mapsto f(-z)/\varphi(z)$ is strictly decreasing.

Then, to show the bounds on $f(-z)/\varphi(z) = 1 - zR(z)$, we use the bounds on $R(z)$ from Lemma 6 to write

$$\begin{aligned} f(-z)/\varphi(z) &> 1 - z[z/(z^2 + 1)] = 1/(z^2 + 1), \\ f(-z)/\varphi(z) &< 1 - z[(z^2 + 2)/(z^3 + 3z)] = 1/(z^2 + 3). \quad \square \end{aligned}$$

We are now ready to prove Proposition 3.

Proof of Proposition 3. First consider the case $\Delta_x = 0$. In this case,

$$v(me_x)/m = \tilde{\sigma}_x(m)f(0)/m = \frac{1}{\sqrt{2\pi}} \frac{1}{m} \sqrt{\frac{\sigma_x^2 m}{(\lambda_x/\sigma_x^2) + m}} = \frac{\sigma_x}{\sqrt{2\pi}} \frac{1}{\sqrt{(\lambda_x/\sigma_x^2)m + m^2}}.$$

This is a strictly decreasing function of m . This shows the result for the case $\Delta_x = 0$.

Now consider the case $\Delta_x > 0$. Differentiating $v(me_x)/m$ with respect to m at $m > 0$ provides

$$\begin{aligned} \frac{\partial}{\partial m} \frac{v(me_x)}{m} &= -\frac{1}{m^2} v(me_x) + \frac{1}{m} \frac{\partial}{\partial m} v(me_x) \\ &= -\frac{1}{m^2} v(me_x) + \frac{1}{m} \tilde{\sigma}'_x(m) \varphi\left(-\frac{\Delta_x}{\tilde{\sigma}_x(m)}\right) \\ &= -\frac{1}{m^2} \tilde{\sigma}_x(m) f\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right) + \frac{1}{m} \varphi\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right) \frac{\lambda_x}{2} \frac{\tilde{\sigma}_x(m)}{m(\lambda_x + \sigma_x^2 m)} \\ &= \varphi\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right) \frac{\tilde{\sigma}_x(m)}{m^2(\lambda_x + \sigma_x^2 m)} \left[-\varphi\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right)^{-1} f\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right) (\lambda_x + \sigma_x^2 m) + \frac{\lambda}{2} \right]. \end{aligned}$$

In the first line we have used the chain rule. In the second line we have used Proposition 1 to differentiate v . In the third line we have used the expression from Proposition 1 for $v(m_e x)$ and Lemma 2 to differentiate $\tilde{\sigma}_x(m)$. In the fourth line we have rearranged terms.

Since $\varphi\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right)\tilde{\sigma}_x(m)/m^2$ is strictly positive, the sign of $\frac{\partial}{\partial m}v(m_e x)/m$, i.e., whether it is strictly positive, strictly negative, or 0, is equal to the sign of the expression inside the parentheses, which we will call $h(m)$,

$$h(m) = -\varphi\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right)^{-1} f\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right) (\lambda_x + \sigma_x^2 m) + \frac{\lambda}{2}. \quad (14)$$

Since $\varphi(-z)^{-1}f(-z)$ is strictly decreasing in z by Lemma 7, and $\tilde{\sigma}_x(m)$ is strictly increasing in m , both $\varphi\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right)^{-1} f\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right)$ and $\varphi\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right)^{-1} f\left(\frac{-\Delta_x}{\tilde{\sigma}_x(m)}\right) (\lambda_x + \sigma_x^2 m)$ are strictly increasing in m . Thus, $h(m)$ is strictly decreasing in m .

Since h is continuous on \mathbf{R}_{++} with $\lim_{m \rightarrow 0} h(m) = \lambda/2 > 0$ (this limit relies on $\Delta_x > 0$) and $\lim_{m \rightarrow \infty} h(m) = -\infty$, h has a unique root in \mathbf{R}_{++} . Call this root m^* . In addition, both $h(m)$ and $\frac{\partial}{\partial m}v(m_e x)$ are strictly positive for $m < m^*$ and strictly negative for $m > m^*$. Thus, $m \mapsto v(m_e x)$ is strictly unimodal with a unique maximum at m^* .

We now bound m^* above and below. By Lemma 7 and (14), $h(m)$ is bounded above by $h(m, 3)$ and below by $h(m, 1)$, where $h(m, a)$ is defined by

$$h(m, a) = -\frac{\lambda_x + \sigma_x^2 m}{(\Delta_x/\tilde{\sigma}_x(m))^2 + a} + \frac{\lambda}{2}.$$

Using the definition of $\tilde{\sigma}_x(m)$, we may write $(\Delta_x/\tilde{\sigma}_x(m))^2$ as

$$(\Delta_x/\tilde{\sigma}_x(m))^2 + a = \Delta_x^2 \frac{(\lambda_x/\sigma_x^2) + m}{\sigma_x^2 m} + a \frac{(\Delta_x^2 \lambda_x/\sigma_x^2) + (\Delta_x^2 + \sigma_x^2 a)m}{\sigma_x^2 m},$$

and then substitute this into the definition of $h(m, a)$ to obtain

$$\begin{aligned} h(m, a) &= -\frac{(\lambda_x + \sigma_x^2 m)\sigma_x^2 m}{(\Delta_x^2 \lambda_x/\sigma_x^2) + (\Delta_x^2 + \sigma_x^2 a)m} + a \\ &= \frac{-2\sigma_x^4 m^2 + \lambda_x(\sigma_x^2(a-2) + \Delta_x^2)m + (\Delta_x^2 \lambda_x^2)/(\sigma_x^2)}{2[(\Delta_x^2 \lambda_x/\sigma_x^2) + (\Delta_x^2 + \sigma_x^2 a)m]}. \end{aligned}$$

Let $\tilde{h}(m, a) = -2\sigma_x^4 m^2 + \lambda_x(\sigma_x^2(a-2) + \Delta_x^2)m + (\Delta_x^2 \lambda_x^2)/(\sigma_x^2)$ be the numerator of this expression, so that the sign of $h(m, a)$ is the same as the sign of $\tilde{h}(m, a)$.

This expression $\tilde{h}(m, a)$ is quadratic in m and has roots

$$\left[\lambda_x(\sigma_x^2(a-2) + \Delta_x^2) \pm \sqrt{\lambda_x^2(\sigma_x^2(a-2) + \Delta_x^2)^2 + 8\Delta_x^2\lambda_x^2\sigma_x^2} \right] / 4\sigma_x^4.$$

For $a \in \{1, 3\}$, $(a-2)^2 = 1$ and these roots may be written

$$\left[(a-2) + r \pm \sqrt{1 + [2(a-2) + 8]r + r^2} \right] / 4\sigma_x^2.$$

where $r = \Delta_x^2/\sigma_x^2$. Of these roots when $a \in \{1, 3\}$, only one is strictly positive and the other is strictly negative. Let $\tilde{h}_+(a)$ denote the strictly positive root, so

$$\begin{aligned} \tilde{h}_+(1) &= \left[-1 + r + \sqrt{1 + 6r + r^2} \right] / 4\sigma_x^2, \\ \tilde{h}_+(3) &= \left[1 + r + \sqrt{1 + 10r + r^2} \right] / 4\sigma_x^2. \end{aligned}$$

Since $\lim_{m \rightarrow 0} \tilde{h}(m, a) = \Delta_x^2\lambda_x^2/\sigma_x^2 > 0$ and $\lim_{m \rightarrow \infty} \tilde{h}(m, a) = -\infty$, $\tilde{h}(m, a)$ is strictly positive on $[0, \tilde{h}_+(a))$ and strictly negative on $(\tilde{h}_+(a), \infty)$. Furthermore, $h(m) \leq h(m, 3)$ and $\text{sgn}(h(m, 3)) = \text{sgn}(\tilde{h}(m, 3))$ implies $\tilde{h}(m) < 0$ for $m > \tilde{h}_+(3)$, which implies $m^* \leq \tilde{h}_+(3)$. Similarly, $h(m) \geq h(m, 1)$ and $\text{sgn}(h(m, 1)) = \text{sgn}(\tilde{h}(m, 1))$ implies $\tilde{h}(m) > 0$ for $m < \tilde{h}_+(1)$, which implies $m^* \geq \tilde{h}_+(1)$. Thus, $m^* \in [\tilde{h}_+(1), \tilde{h}_+(3)]$. \square

Proof of Theorem 2

Using the chain rule, we compute the gradient and Hessian of v as

$$\begin{aligned} \nabla v(n) &= \nabla g(\tilde{\sigma}(n)) \text{diag}(\nabla \tilde{\sigma}(n)), \\ \nabla^2 v(n) &= \text{diag}(\nabla \tilde{\sigma}(n)) \nabla^2 g(\tilde{\sigma}(n)) \text{diag}(\nabla \tilde{\sigma}(n)) + \nabla^2 \tilde{\sigma}(n) \text{diag}(\nabla g(\tilde{\sigma}(n))). \end{aligned} \quad (15)$$

Consider the term $\nabla^2 \tilde{\sigma}(n) \text{diag}(\nabla g(\tilde{\sigma}(n)))$ from (15). Since $\nabla^2 \tilde{\sigma}(n)$ is a diagonal matrix with entries $\tilde{\sigma}_i''(n_i)$, this term is a diagonal matrix with entries $\tilde{\sigma}_i''(n_i) \frac{\partial g}{\partial s_i}(\tilde{\sigma}(n)) = -\tilde{\sigma}_i'(n_i)^2 w_i(n_i) \frac{\partial g}{\partial s_i}(\tilde{\sigma}(n))$, where we define $w_i(n_i) = -\tilde{\sigma}_i''(n_i)/\tilde{\sigma}_i'(n_i)^2$. Defining a vector $w(n) = [w_i(n_i)]_i$, we then have an expression for the term in (15),

$$\nabla^2 \tilde{\sigma}(n) \text{diag}(\nabla g(\tilde{\sigma}(n))) = -\text{diag}(\nabla \tilde{\sigma}(n)) \text{diag}(w(n)) \text{diag}(\nabla g(\tilde{\sigma}(n))) \text{diag}(\nabla(\tilde{\sigma}(n))),$$

and (15) becomes

$$\nabla^2 v(n) = \text{diag}(\nabla \tilde{\sigma}(n)) \left[\nabla^2 g(\tilde{\sigma}(n)) - \text{diag}(w(n)) \text{diag}(\nabla g(\tilde{\sigma}(n))) \right] \text{diag}(\nabla \tilde{\sigma}(n)).$$

Thus, to show that $\nabla^2 v(n)$ is negative semi-definite and that v is concave at n , it is sufficient to show that

$$H(n) = \nabla^2 g(\tilde{\sigma}(n)) - \text{diag}(w(n)) \text{diag}(\nabla g(\tilde{\sigma}(n)))$$

is negative semi-definite.

From the expressions for $\tilde{\sigma}'_i$ and $\tilde{\sigma}''_i$ in Lemma 2,

$$w_i(n_i) = \frac{\sigma_i}{\lambda_i} n_i^{-1/2} \left(\frac{\lambda_i}{\sigma_i^2} + n_i \right)^{1/2} \left(\frac{\lambda_i}{\sigma_i^2} + 4n_i \right).$$

Letting maxeig and mineig be functions that returns the maximum and minimum eigenvalues of a matrix respectively, we have from the fact that $\text{maxeig}(A + B) \leq \text{maxeig}(A) + \text{maxeig}(B)$ that

$$\text{maxeig}(H(n)) \leq \text{maxeig}(\nabla^2 g(\tilde{\sigma}(n))) - \text{mineig}(\text{diag}(w(n)) \text{diag}(\nabla g(\tilde{\sigma}(n)))).$$

Furthermore, $\text{mineig}(\text{diag}(w(n)) \text{diag}(\nabla g(\tilde{\sigma}(n)))) = \min_i w_i(n_i) \frac{\partial g}{\partial s_i}(\tilde{\sigma}(n))$.

Let $B_\epsilon(\sigma) = \{u \in \mathbf{R}^M : |u_i - \sigma_i| < \epsilon\}$, and let $\epsilon > 0$ be such that constants b and c defined by

$$b = \sup_{u \in B_\epsilon(\sigma)} \text{maxeig}(\nabla^2 g(u)), \quad c = \min_i \inf_{u \in B_\epsilon(\sigma)} \frac{\partial g}{\partial s_i}(u)$$

satisfy $b < \infty$ and $c > 0$. That such a strictly positive ϵ exists is guaranteed by the continuity of maxeig , the continuity of $\nabla^2 g$ and ∇g (Lemma 4), and the fact due to Lemma 3 that $\nabla g(\sigma) \succ 0$. Here, \succ indicates componentwise inequality, so $a \succ a'$ means $a_i > a'_i$ for each i .

Then let $N \in \mathbf{R}_+^M$ be such that $\tilde{\sigma}(N) \in B_\epsilon(\sigma_i)$. Such an N exists because $\lim_{N_i \rightarrow \infty} \tilde{\sigma}_i(N_i) = \sigma_i$. For this N , we also have $\tilde{\sigma}(n) \in B_\epsilon$ for all $n \succ N$ because $\tilde{\sigma}$ is increasing in each component. Thus, for all $n \succ N$, $\text{maxeig}(\nabla^2 g(\tilde{\sigma}(n))) \leq b$, $\min_i \frac{\partial g}{\partial s_i}(\tilde{\sigma}(n)) \geq c$, and

$$\text{maxeig}(H(n)) \leq b - \min_i c w_i(n_i). \tag{16}$$

Since $\lim_{n_i \rightarrow \infty} w_i(n_i) = \infty$ for each i , there exists an $N^* \succ N$ such that $\min_i w_i(n_i) \geq b/c$ for all $n \succ N^*$. From (16), we then have that $\text{maxeig}(H(n)) \leq 0$ for all $n \succ N^*$.

Proof of Theorem 3

Optimality at $N = 1$ follows from noting that the decision of the KG policy is identical to that of the KG(*) policy when $n = N - 1$. The decision of the KG policy is shown to be optimal when $N = 1$ in Frazier et al. (2008).

To show optimality in the limit as $N \rightarrow \infty$, we first augment our existing notation. To indicate the dependence upon the posterior distribution we write $v(n; \mu, \sigma)$ instead of $v(n)$ to indicate the single-stage value of the allocation n under the posterior distribution given by $\mu \in \mathbb{R}^M$ and $\sigma \in \mathbb{R}_+^M$. We also let $\nu_x^{\text{KG}}(\mu, \sigma) = v(e_x; \mu, \sigma)$ and $\nu_x^{\text{KG}^*}(\mu, \sigma) = \max_{1 \leq m \leq N-n} v(me_x; \mu, \sigma)/m$ be the KG and KG(*) factors, respectively, for measuring alternative x . Thus, the KG policy measures $\arg \max_x \nu_x^{\text{KG}}(\mu, \sigma)$, and the KG(*) policy measures $\arg \max_x \nu_x^{\text{KG}^*}(\mu, \sigma)$,

We now provide upper and lower bounds on $\nu_x^{\text{KG}^*}(\mu, \sigma)$. By noting that $m = 1$ is in the set of m over which the maximum in the definition of $\nu_x^{\text{KG}^*}(\mu, \sigma)$ is taken, we obtain the lower bound,

$$\nu_x^{\text{KG}^*}(\mu, \sigma) \geq v(e_x; \mu, \sigma) = \nu_x^{\text{KG}}(\mu, \sigma).$$

To provide an upper bound, we note that $v(me_x; \mu, \sigma)$ is bounded above by $u_x(\mu, \sigma) = \sigma_x f(-|\mu_x - \max_{x' \neq x} \mu_{x'}|/\sigma_x)$. This can be seen by noting that $s \mapsto sf(-a/s)$ is increasing in s for any $a \geq 0$, and $\tilde{\sigma}_x(m) \leq \sigma_x$. Then, letting m^* be a maximand in the definition of $\nu_x^{\text{KG}^*}(\mu, \sigma)$, and noting $m^* \geq 1$,

$$\nu_x^{\text{KG}^*}(\mu, \sigma) = v(m^* e_x; \mu, \sigma)/m^* \leq v(m^* e_x; \mu, \sigma) \leq u_x(\mu, \sigma).$$

As shown in Frazier and Powell (2009) (Theorem 1 and Lemmas 5 and 6), a sufficient condition for optimality as $N \rightarrow \infty$ is that, for each $(\mu, \sigma) \in \mathbb{R}^M \times \mathbb{R}_+^M$ satisfying $\sigma \neq 0$, we have a corresponding set U containing (μ, σ) that is open in $\mathbb{R}^M \times \mathbb{R}_+^M$ and on which the policy only measures alternatives x with $\sigma_x > 0$.

To show that this sufficient condition is met, we fix $\mu \in \mathbb{R}^M$ and $\sigma \in \mathbb{R}_+^M$ with $\sigma \neq 0$. We define $A = \{x : \sigma_x = 0\}$, and

$$U = \{\mu', \sigma' : \nu_x^{\text{KG}}(\mu', \sigma') > u_{x'}(\mu', \sigma') \forall x \notin A, x' \in A\}.$$

Note that $U = \mathbf{R}^M \times \mathbf{R}_+^M$ when $A = \emptyset$. We check that U satisfies the sufficient conditions. First, the upper and lower bounds for $\nu^{\text{KG}(\ast)}$ imply for all $x \in A$, $x' \notin A$ and $(\mu', \sigma') \in U$ that

$$\nu_x^{\text{KG}(\ast)}(\mu', \sigma') \geq \nu_x^{\text{KG}}(\mu', \sigma') > u_{x'}(\mu', \sigma') \geq \nu_{x'}^{\text{KG}}(\mu', \sigma').$$

Since A has a non-empty complement (although it may be empty itself), this implies that the $\text{KG}(\ast)$ policy measures outside of A for each $(\mu', \sigma') \in U$.

Second, the upper and lower bounds for $\nu^{\text{KG}(\ast)}$, together with the facts $\nu_i^{\text{KG}}(\mu, \sigma) = 0$ and $u_i(\mu, \sigma) = 0$ iff $\sigma_i = 0$ iff $i \in A$, imply for all $x \in A$ and $x' \notin A$ that

$$\begin{aligned} \nu_x^{\text{KG}(\ast)}(\mu, \sigma) &\geq \nu_x^{\text{KG}}(\mu, \sigma) > 0 \\ \nu_{x'}^{\text{KG}(\ast)}(\mu, \sigma) &\leq \nu_{x'}^{\text{KG}}(\mu, \sigma) = 0. \end{aligned}$$

The inequality $\nu_{x'}^{\text{KG}(\ast)}(\mu, \sigma) \geq 0$ then shows $\nu_{x'}^{\text{KG}(\ast)}(\mu, \sigma) = 0$. Thus, $\nu_x^{\text{KG}(\ast)}(\mu, \sigma) > 0 = \nu_{x'}^{\text{KG}(\ast)}(\mu, \sigma)$, and $(\mu, \sigma) \in U$.

Third and finally, both ν_x^{KG} and u_x are continuous functions, and so U is an open set. Thus, U satisfies the sufficient conditions from Frazier and Powell (2009).

References

- Berger, J. and Deely, J. (1988). A Bayesian Approach to Ranking and Selection of Related Means With Alternatives to Analysis-of-Variance Methodology. *Journal of the American Statistical Association*, 83(402):364–373.
- Bickel, J. and Smith, J. (2006). Optimal sequential exploration: A binary learning model. *Decision Analysis*, 14(15):16.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press, Cambridge.
- Chade, H. and Schlee, E. (2002). Another look at the Radner-Stiglitz nonconcavity in the value of information. *Journal of Economic Theory*, 107(2):421–452.

- Chen, C. (1995). An effective approach to smartly allocate computing budget for discrete event simulation. *IEEE Conference on Decision and Control, 34 th, New Orleans, LA*, pages 2598–2603.
- Chen, C. (1996). A lower bound for the correct subset-selection probability and its application to discrete-event system simulations. *IEEE Transactions on Automatic Control*, 41:1227–1231.
- Chen, C., Dai, L., and Chen, H. (1996). A gradient approach for smartly allocating computing budget for discrete event simulation. In Charnes, J., Morrice, D., Brunner, D., and Swain, J., editors, *Proceedings of the 1996 Winter Simulation Conference*, pages 398–405, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.
- Chen, C., Lin, J., Yücesan, E., and Chick, S. (2000a). Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Discrete Event Dynamic Systems*, 10(3):251–270.
- Chen, H., Chen, C., and Yucesan, E. (2000b). Computing efforts allocation for ordinal optimization and discrete event simulation. *Automatic Control, IEEE Transactions on*, 45(5):960–964.
- Chen, H., Dai, L., Chen, C., and Yücesan, E. (1997). New development of optimal computing budget allocation for discrete event simulation. *Proceedings of the 29th conference on Winter simulation-Volume 00*, pages 334–341.
- Chick, S., Branke, J., and Schmidt, C. (2007). New greedy myopic and existing asymptotic sequential selection procedures: preliminary empirical results. In *Proceedings of the 2007 Winter Simulation Conference*, Piscataway, NJ. Winter Simulation Conference, IEEE.
- Chick, S., Branke, J., and Schmidt, C. (2009). Sequential sampling to myopically maximize the expected value of information. *INFORMS J. on Computing*. to appear.
- Chick, S. and Frazier, P. (2009a). The conjunction of the knowledge gradient and economic approach to simulation selection. *Winter Simul. Conf. Proc., 2009*.

- Chick, S. and Frazier, P. (2009b). Sequential sampling for selection: Beyond the knowledge gradient. submitted.
- Chick, S. and Gans, N. (2009). Economic analysis of simulation selection problems. *Management Sci.*, to appear.
- Chick, S. and Inoue, K. (2001). New two-stage and sequential procedures for selecting the best simulated system. *Operations Research*, 49(5):732–743.
- De Lara, M. and Gilotte, L. (2007). A tight sufficient condition for Radner–Stiglitz nonconvexity in the value of information. *Journal of Economic Theory*, 137(1):696–708.
- DeGroot, M. H. (1970). *Optimal Statistical Decisions*. John Wiley and Sons.
- Frazier, P. and Powell, W. (2009). Convergence to global optimality with sequential bayesian sampling policies. submitted.
- Frazier, P., Powell, W. B., and Dayanik, S. (2008). A knowledge gradient policy for sequential information collection. *SIAM J. on Control and Optimization*, 47(5):2410–2439.
- Goel, A., Khanna, S., and Null, B. (2008). The Ratio Index for Budgeted Learning, with Applications. *Arxiv preprint arXiv:0810.0558*.
- Guha, S. and Munagala, K. (2007). Approximation algorithms for budgeted learning problems. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 104–113. ACM New York, NY, USA.
- Gupta, S. and Miescke, K. (1996). Bayesian look ahead one-stage sampling allocations for selection of the best population. *Journal of statistical planning and inference*, 54(2):229–244.
- Guttman, I. and Tiao, G. (1964). A bayesian approach to some best population problems. *The Annals of Mathematical Statistics*, 35:825–835.
- He, D., Chick, S., and Chen, C. (2007). Opportunity cost and ooba selection procedures in ordinal optimization for a fixed number of alternative systems. *IEEE Transactions on Systems Man and Cybernetics Part C-Applications and Reviews*, 37(5):951–961.

- Howard, R. (1966). Information Value Theory. *Systems Science and Cybernetics, IEEE Transactions on*, 2(1):22–26.
- Kaelbling, L., Littman, M., and Cassandra, A. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134.
- Kapoor, A. and Greiner, R. (2005). Learning and Classifying Under Hard Budgets. In *16th European Conference on Machine Learning, Porto, Portugal, October 3-7, 2005*. Springer-Verlag New York Inc.
- Laplace, P. (1820). *Théorie analytique des probabilités*. Courcier.
- Lizotte, D., Madani, O., and Greiner, R. (2003). Budgeted learning of naive-bayes classifiers. In *UAI*, volume 3, pages 378–385. Citeseer.
- Mattsson, L., Voorneveld, M., and Weibull, J. (2004). Information-Constrained Discrete Choice. Working Paper 558, Stockholm School of Economics/EFI.
- Moscarini, G. and Smith, L. (2002). The law of large demand for information. *Econometrica*, pages 2351–2366.
- Radner, R. and Stiglitz, J. (1984). A Nonconcavity in the Value of Information. *Bayesian models in economic theory*, 5:33–52.
- Raiffa, H. and Schlaifer, R. (1968). *Applied Statistical Decision Theory*. M.I.T. Press.
- Ruben, H. (1963). A convergent asymptotic expansion for Mill’s ratio and the normal probability integral in terms of rational functions. *Mathematische Annalen*, 151(4):355–364.
- Schwartz, B. (2004). *The paradox of choice: Why more is less*. Harper Perennial.
- Stigler, G. (1961). The economics of information. *The Journal of Political Economy*, 69(3).
- Tiao, G. and Afonja, B. (1976). Some bayesian considerations of the choice of design for ranking, selection and estimation. *Annals of the Institute of Statistical Mathematics*, 28:167–185.

Weibull, J., Mattsson, L., and Voorneveld, M. (2007). Better May be Worse: Some Monotonicity Results and Paradoxes in Discrete Choice Under Uncertainty. *Theory and Decision*, 63(2):121–151.

Yokota, F. and Thompson, K. (2004). Value of information literature analysis: a review of applications in health risk management. *Medical Decision Making*, 24(3):287.