

The A/B Testing Problem with Gaussian Priors

Eduardo M. Azevedo*, David Mao†, José Luis Montiel Olea‡ and Amilcar Velez§

March 22, 2020

Abstract

A risk-neutral firm can perform a randomized experiment (A/B test) to learn about the effects of implementing an idea of unknown quality. The firm's goal is to decide the experiment's sample size and whether or not the idea should be implemented at scale after observing the experiment's outcome. In this paper we study this classical problem when the firm's prior distribution over idea quality is Gaussian. We provide four results. First, there is a closed-form solution for the value of running a randomized experiment. Second, if costs increase linearly in the size of the experiment, there is a simple solution to the experiment's optimal sample size. Third, we derive comparative statics for the value of experimentation and the firm's optimal experimentation strategy. Fourth, we solve for the sample size that minimizes the firm's maximum regret and present the problem's least favorable prior over idea quality. Numerical examples confirm that indeed the firm's expected profits under the optimal experimentation strategy are higher than under standard rules of thumb for choosing sample size.

*The Wharton School, University of Pennsylvania. Email: eazevedo@wharton.upenn.edu

†Department of Economics, University of Pennsylvania. Email: davidmao@sas.upenn.edu

‡Department of Economics, Columbia University. Email: jm4474@columbia.edu

§Department of Economics, Northwestern University. Email: amilcare@u.northwestern.edu

1 Introduction

There has been a revolution in the use of randomized experiments in the last twenty years across a number of fields. One prominent example is that of large internet companies, which routinely use experiments with tens of millions of users to test almost all of their product innovations. Technology companies like Google, Facebook, and Microsoft call these experiments “A/B Tests”. A/B tests have revolutionized how these and other companies screen product innovations.

The popularity of online experimentation has reignited interest in classical questions in experimental design, such as how to best decide the size of an experiment. A prescient literature in statistical decision theory studied this and many other questions in the 1950s and 1960s. For instance, [Raiffa and Schlaifer \(1961\)](#) provide a textbook treatment of the “mathematical analysis of decision making when the state of the world is uncertain but further information about it can be obtained by experimentation”. They also provide examples that yield “charts from which the optimal sample size and expected net gain of an optimal sample can be read directly”.

In this paper we revisit this classical framework. We consider a risk-neutral firm that has an idea of unknown quality, but can perform an experiment to learn about it. The firm’s goal is to decide the experiment’s size and whether or not the idea should be implemented at scale after observing the experiment’s outcome.¹ As in [Raiffa and Schlaifer \(1961\)](#) we further assume that, i) conditional on the idea’s quality, the outcome of the experiment is Gaussian with known variance that decreases inversely proportional to the sample size; ii) the cost of experimentation increases linearly in the sample size; and iii) the firm’s prior distribution of idea quality is Gaussian. This model is a particular case of the A/B Testing Problem in [Azevedo et al. \(2020\)](#) where a firm uses scarce experimental resources to screen multiple

¹For example, if the firm runs a search engine, an idea is a potential improvement developed by engineers, quality is some scalar performance measure, and the experiment is an A/B test conducted using a share of the website’s traffic.

potential innovations in order to implement a subset of them at scale.

We present three main results, the first two of which appear *mutatis mutandis* in [Raiffa and Schlaifer \(1961\)](#) and the references therein. First, there is a closed-form solution to the value obtained from an experiment (proposition 2).² Second, there is a simple characterization of the experiment’s optimal sample size (proposition 3). Third, we derive comparative statics of the firm’s expected profits and the optimal sample size with respect to the parameters of the Gaussian prior (proposition 4). To the best of our knowledge, proposition 4 is novel. One important, counterintuitive finding from the comparative statistics is that an increase in the prior variance leads to an increase in the firm’s expected profits. Interestingly, such an increase in profits is sometimes achieved by decreasing the experiment’s sample size.

From the firm’s perspective, the principle of determining an experiment’s sample size by maximizing profits seems more appealing than the standard and prevalent practice of using power calculations for sample size determination.³ This point echoes the critiques of [Meltzer \(2001\)](#) and [Manski and Tetenov \(2016, 2019\)](#). The implementation of the optimal sample size is especially appealing in data-rich environments, such as experimentation in online firms, where information from past experiments is readily available and can be used to choose the prior following an empirical Bayes approach (see the discussion in [Azevedo et al. \(2019\)](#) and the references therein).

A common objection to the decision-theoretic analysis in this paper is that it calls for the use of a specific prior distribution over idea quality. One possible alternative is to replace the Bayesian analysis of this problem by minimax or minimax regret choices of sample size. This has been done in [Bross \(1950\)](#) and [Somerville \(1954\)](#) for the problem we study. We present their minimax regret formula for the optimal sample size using the notation of the A/B testing problem, and we complement their result by exhibiting the problem’s least favorable prior distribution (proposition 5). This prior distribution can be interpreted as the

²The value obtained from an experiment is what [Radner and Stiglitz \(1984\)](#); [Meltzer \(2001\)](#); [Chade and Schlee \(2002\)](#) refer to as the value of information.

³See for example [List et al. \(2011\)](#); [Athey and Imbens \(2017\)](#)

distribution that an adversarial nature would select to minimize the firm’s expected regret. The least favorable distribution turns out to be a two-point uniform distribution (one point positive and the other negative), where the points of support are a fraction of the optimal experiment’s standard deviation. The minimax regret choice of the sample size is the firm’s optimal choice when the prior is the least favorable distribution.

The rest of the paper is organized as follows. Section 2 presents the model, section 3 the results, and section 4 a numerical example. Section 5 concludes. Proofs are in the appendix.

2 Model

2.1 The Model

A firm has a single innovation, or idea. The firm is uncertain about the quality of the idea. The true quality of the idea is a normally distributed random variable Δ with mean M in \mathbb{R} and variance s^2 . The firm can perform an experiment (also known as an A/B test) to observe a noisy signal of quality. An experiment with n users gives a normally distributed signal $\hat{\Delta}$ with mean equal to the true quality and variance σ^2/n . The firm incurs a cost c per user in the experiment.

The firm has two choice variables. The firm can choose an *experimentation strategy* n in \mathbb{R}^+ , the number of users assigned to the experiment. We also refer to n as how much data to use. After observing the result of the experiment, the firm can choose an *implementation strategy* S equal to 0 or 1 depending on whether the firm wants to implement the idea. Thus, S is a measurable function of the signal realization, $\hat{\Delta}$. The firm’s payoff is the expected quality of the idea if implemented minus the cost of experimentation,

$$\Pi(n, S) = \mathbb{E}[S \cdot \Delta] - c \cdot n.$$

The firm's goal is to choose the experimentation and implementation strategy to maximize profits.

2.2 Related Literature

The problem described above is a particular case of what [Raiffa and Schlaifer \(1961\)](#) section 5.5 call a two-action problem with a scalar state, linear payoffs, and a Gaussian prior. Since their notation is sometimes cumbersome, we provide our own proofs for all results and give credit to the original sources in the proposition statements.

This model is a special case of [Azevedo et al. \(2020\)](#). The key restriction is that we consider a normal prior, whereas they consider general priors. Moreover, we focus on the particular case of a single idea where the only cost is a linear cost of obtaining more data. In contrast, they allow for several ideas and more general costs (see their sections 2 and 5.2). We make these simplifications to focus on the key insights from the Gaussian case.

2.3 Notation and the Production Function

We now introduce notation that simplifies the solution of the A/B testing problem. We denote the realization of the random variable Δ as δ , and likewise, we denote the realization of $\hat{\Delta}$ as $\hat{\delta}$. We denote $m := M/s$, and the share of the variance of the signal explained by the prior as $\theta_n := s^2/(s^2 + \sigma^2/n)$ for $n > 0$ and $\theta_0 := 0$.

After running an experiment, the firm uses Bayes' rule to calculate the posterior mean quality. We denote the posterior mean after an experiment of sample size n with result $\hat{\delta}$ as

$$P(\hat{\delta}, n) := \mathbb{E}[\Delta | \hat{\Delta} = \hat{\delta}].$$

The optimal implementation strategy is simple: the firm should implement an idea if and only if the posterior mean quality is positive. That is, the firm's expected payoff after an

experiment equals the posterior mean if it is positive (in which case the idea is implemented), or zero if the posterior mean is negative (in which case the idea is not implemented). This means that, after running an experiment with n users, the firm's expected payoff is

$$\mathbb{E}[P(\hat{\Delta}, n)^+] - c \cdot n.$$

We follow [Azevedo et al. \(2020\)](#) and define the *production function* $f(n)$ as the net value of investing n users in an experiment. The value of an idea without an experiment is M^+ , because the idea is only valuable if it can be implemented profitably without observing any data. So, we define $f(n)$ as

$$f(n) := \mathbb{E}[P(\hat{\Delta}, n)^+] - M^+.$$

The production function can be used to write the A/B testing problem as a maximization problem. The firm's problem is to maximize the value of the data it invests into the idea net of the costs, as in neoclassical producer theory. In our setting with a single idea, we can simplify the solution as follows.

[Azevedo et al. \(2020\)](#) showed that the optimal implementation strategy is to implement the idea if and only if $P(\hat{\delta}, n) > 0$. Under the optimal implementation strategy, the firm's payoff equals the production function minus the cost of data plus the value of the idea with no data:

$$\Pi(n, S) = f(n) - c \cdot n + M^+. \tag{1}$$

Therefore, the optimal experimentation strategy maximizes $f(n) - c \cdot n$.

3 Main Results

3.1 Optimal Implementation Strategy.

The standard formula for Bayesian updating with a Gaussian prior is that the posterior mean is a convex combination between the data, $\hat{\delta}$, and the prior, M :

$$P(\hat{\delta}, n) = \theta_n \hat{\delta} + (1 - \theta_n)M.$$

From this formula we obtain the following simple result:

Proposition 1 (Optimal Implementation Strategy). *It is optimal to implement an idea if and only if the result $\hat{\delta}$ of the experiment is greater than $t_n^* \cdot \sigma/\sqrt{n}$, where we refer to t_n^* as the threshold t -statistic*

$$t_n^* := -m \cdot \frac{\sigma/\sqrt{n}}{s}.$$

The firm should calculate the standard frequentist t -statistic of quality associated with the experiment—i.e., $\hat{\Delta}/(\sigma/\sqrt{n})$ —and implement the idea only if the t -statistic is above the threshold. The threshold t -statistic tells the firm how strict it should be in implementing the idea. If t_n^* happens to be equal to 1.65 (the 95th percentile of the standard Gaussian distribution), then the optimal implementation strategy corresponds to the commonly used rule of thumb of a statistically significant positive effect with a 5% p -value. The formula makes clear that there is no reason for the rule of thumb to be optimal. The threshold p -value could be much greater if, for example, the prior about idea quality has mean close to zero, or if the experiment is very precise relative to s . And the threshold p -value could be much smaller if, for example, the prior mean idea quality is far from 0.

3.2 Production Function and Optimal Experimentation Strategy.

Under a Gaussian prior, there is a closed-form solution for the production function:⁴

Proposition 2 (Production Function (Raiffa and Schlaifer, 1961; Grundy et al., 1956)). *The production function is*

$$f(n) = s \cdot \left(\sqrt{\theta_n} \cdot \phi\left(\frac{m}{\sqrt{\theta_n}}\right) - |m| \cdot \Phi\left(\frac{-|m|}{\sqrt{\theta_n}}\right) \right) \quad (2)$$

for $n > 0$, and $f(0) = 0$.

This function is bounded, increasing, convex in an interval $[0, \hat{n}]$, and concave in an interval $[\hat{n}, \infty)$, where

$$\hat{n} = \frac{\sigma^2}{8s^2} \cdot \left(m^2 - 1 + \sqrt{m^4 + 14m^2 + 1} \right).$$


In addition, when $M \neq 0$, the marginal product $f'(n)$ satisfies $f'(0) = 0$ and $\lim_{n \rightarrow \infty} f'(n) = 0$.

Figure 1 depicts an example production function. Azevedo et al. (2020) have shown that, with thin-tailed priors, the production function is convex near zero. For the specific case of Gaussian priors, we can provide more details regarding the shape of the production function. In particular, proposition 2 shows that the production function is convex first and then concave.

The formula for the inflection point shows that when M is close to zero the region of convexity is quite small. The same is true when σ^2 is small. These effects are both consistent with the intuitions in Radner and Stiglitz (1984) and Chade and Schlee (2002). These papers give conditions under which the production function is convex near $n = 0$. For instance, suppose that $M < 0$ so that under the prior the firm prefers not to implement an idea at scale. An

⁴The earliest reference we could find for this formula is Grundy et al. (1956) equation (4), under a slightly different setup. Keppo et al. (2008) derive a closed-form solution for the value of information in the case of two possible states, two possible actions, and normally distributed signals. Moscarini and Smith (2002) derive an asymptotic formula for the value of information with a finite number of signals and actions.

Figure 1: The production function $f(n)$.



Notes: The parameters are prior mean $M = -5$, prior standard deviation $s = 5$, and experimental noise $\sigma = 30$.

extremely noisy signal for the experiment is likely to reverse this decision, so the marginal value of information will be close to zero. If, however, σ^2 is close to zero, small experiments can provide informative signals that revert the prior ordering over actions.

With our assumption of a linear cost of data and a single idea, the optimal experimentation strategy is particularly simple. When the cost of obtaining information always exceeds the value of information—i.e., $c \cdot n > f(n)$ for all $n > 0$ —it is optimal not to experiment at all and set $n^* = 0$.

When the value of information exceeds the cost of obtaining information for some n , the optimal experimentation strategy is interior and satisfies the firm’s first order condition. Using Proposition 2, we formalize this insight.

Proposition 3 (Optimal Experimentation Strategy, (Raiffa and Schlaifer, 1961; Grundy et al., 1956)). *Assume that $f(n) > c \cdot n$ for some $n > 0$. The optimal experimentation strategy n^* solves the first order condition of the firm’s profit maximization problem:*

$$f'(n) = c.$$

The first-order condition has either one or two solutions and n^ is the largest solution. n^* is greater than or equal to the inflection point defined in Proposition 2.*

The first-order condition equates the marginal product of data and the marginal cost. It defines the optimal sample size exactly, but it is an implicit equation that needs to be solved numerically. The reasoning behind the optimality result is similar to neoclassical producer theory. By Proposition 2, if $M \neq 0$, the marginal product starts at zero, increases from zero to \hat{n} and then decreases asymptotically to zero. This implies that there can be two sample sizes that equate marginal product and marginal cost: one smaller than \hat{n} , and one larger. Among these two, only the larger one has positive profits (because the average product of data $f(n)/n$ is greater than the marginal cost c).

Proposition 3 can be used to choose a sample size for experiments. In practice, the most common rule of thumb for choosing sample size is a power calculation. For example, in medical trials, one typically specifies a “minimum medically effective” treatment effect. The experiment size is then chosen to guarantee a power of 0.8 at this treatment size. Similar procedures are often used by researchers and by companies performing A/B tests.

This standard power calculation approach has been criticized because it has no reason to be optimal, or even close to optimal (Manski and Tetenov, 2016, 2019). Proposition 3 makes clear that power calculations are not optimal in a practical setting that is well-approximated by our assumptions. In particular, the optimal experimental size does not depend on an arbitrary “minimum medically effective” effect size, or on an arbitrary power level. Instead, the optimal experimental size depends on the marginal cost of data c , on the experimental noise σ , and on the parameters M and s of the prior.

Next, we present comparative statics for the production function $f(n)$ and optimal experimental size n^* .

Proposition 4 (Comparative Statics). *Suppose that the conditions of Proposition 3 hold. Then, the comparative statics for the production function $f(n)$ and optimal experimental size*

n^* are:

1. $\frac{\partial f}{\partial M} > 0$ iff $M < 0$
2. $\frac{\partial f}{\partial s} > 0$
3. $\frac{\partial n^*}{\partial M} > 0$ iff $M < 0$
4. $\frac{\partial n^*}{\partial s} > 0$ if $|M| > s$
 > 0 if $|M| < s$ and $f'(\tilde{n}) > c$
 < 0 if $|M| < s$ and $f'(\tilde{n}) < c$,

where

$$\tilde{n} = \frac{\sigma^2 (3m^2 + 2 + \sqrt{m^4 + 20m^2 + 4})}{2s^2 (1 - m^2)}.$$

The comparative statics offer qualitative principles for experimentation. (1) and (2) say that the level of the production function is higher when either the mean is close to zero or prior variance is high. These comparative statics are useful for determining whether to incur a fixed cost to set up an experimental infrastructure. For example, if innovations are very likely to be useful (large positive M and small s) then there is little point for a firm to set up an experimentation platform. It would be better to simply implement all ideas without experimenting and save on fixed costs.

Comparative statics (3) and (4) are about the optimal sample size of an experiment. (3) says that experiments should be larger when the prior mean is close to zero. The practical takeaway is similar to comparative statics (1): when the mean is far from zero, it is better to do smaller experiments. Comparative statics (4) is about how scale depends on the variance of the prior. Interestingly, the relationship is not monotone. The only relatively clear case is when $|M| > s$, so that most of the mass of the prior is on the same side of 0. In that case, a higher prior variance leads to larger experiments. This is an apparent contradiction with the recent finding from [Azevedo et al. \(2020\)](#) that fatter-tailed priors lead to small optimal

experiment sizes. The results do not mathematically contradict each other because our result is about the variance of a Gaussian prior, whereas the [Azevedo et al. \(2020\)](#) result is about the thickness of the tail of the prior. Nevertheless, comparative statics (4) shows that it is not possible to conclude that a more spread out prior always leads to smaller (or larger) experiments.

The intuition behind these comparative statics is as follows. (1) is true because a smaller value of $|M|$ (ideas that are closer to being marginal ex-ante) pushes towards both a higher value of experimentation $f(n)$ and a greater marginal value $f'(n)$. The higher marginal value of data in turn implies the comparative statics (3) for n^* , because n^* is the greater root of $f'(n) = c$, at a point where $f'(n)$ is decreasing in n . Comparative statics (2) holds because more uncertainty about quality increases the value of experimentation.

The most subtle result is comparative statics (4). The key point is a rescaling argument. If we multiply the parameters M , s , and σ by a constant, the problem is unchanged, so the production function is multiplied by the constant. Abusing notation and denoting the production function as a function of n and the parameters, we have

$$f(n|M, s, \sigma) = s \cdot f(n|M/s, 1, \sigma/s).$$

That is, the production function for any given parameters equals the production function for a normalized prior with $s = 1$ and M and σ scaled down by a factor of s . We show this formally in the proof for [Proposition 4](#).

Consider now the effect of increasing s on $f'(n|M, s, \sigma)$, which equals $s \cdot f'(n|M/s, 1, \sigma/s)$. There are three effects: s increases, M/s decreases, and σ/s decreases. Increasing s always increases f' , which pushes towards greater n^* . However, the effect of decreasing σ/s can decrease f' . For example, in the case of large n , decreasing σ/s moves into the range where information is almost perfect, so that the marginal value f' is small.⁵ This effect may

⁵The simplest intuition comes from the large n approximation to the production function from [Azevedo](#)

dominate, which is why the sign of the comparative statics depends on the parameters as described in the proposition.

3.3 Minimax regret in the A/B Testing Problem

Our analysis above relies on the choice of a specific prior distribution for idea quality. One alternative to this approach is to use the *minimax regret strategy*, the strategy that minimizes the maximum regret of the firm's payoff. When the quality of the innovation is δ , the firm's expected payoff from the strategy (n, S) is

$$u(n, S; \delta) = \delta \mathbb{E}[S] - cn,$$

where the expectation is taken over the experimental noise, $n \geq 0$ is the size of the experiment, and $S \in \{0, 1\}$ is the implementation strategy that depends on the result of the experiment, $\hat{\delta}_n$.

The regret of the strategy (n, S) is defined as the difference between the optimal expected payoff if δ were observable, minus the expected payoff from choosing (n, S) . If δ were observable, the optimal strategy would be to implement the innovation if and only if its quality is positive, without conducting any experimentation. Therefore, regret is

$$\mathcal{R}(n, S, \delta) = \delta \cdot 1\{\delta > 0\} - u(n, S; \delta). \tag{3}$$

The minimax regret strategy (n_{MMR}^*, S_{MMR}^*) is the solution to

$$\inf_{n \geq 0, S} \sup_{\delta} \mathcal{R}(n, S, \delta).$$

The minimax regret strategy has been proposed by [Savage \(1951\)](#) as a guideline for making

et al. (2020) Theorem 1. The marginal product for large n is approximately proportional to $(\sigma/s)^2$, which is decreasing in s .

decisions under uncertainty, and is perhaps the most prominent existing alternative for determining sample sizes without considering any additional assumptions (see [Manski \(2019\)](#)). [Bross \(1950\)](#) and [Somerville \(1954\)](#) studied the minimax regret strategy for a problem closely related to ours. We adapt their solution and solve for the experimentation (n) and implementation (S) strategies that minimize the maximum regret for the A/B Testing Problem. For our setting, we additionally find a least favorable prior distribution that minimizes the firm's expected regret. To this end, define the firm's average regret with respect to prior distribution G as

$$r(n, S, G) \equiv \int \mathcal{R}(n, S, \Delta) dG.$$

Let \mathcal{G} denote the set of all possible real-valued distributions for idea quality. The *least favorable distribution* is defined as the prior over the idea quality that maximizes the optimized average regret. That is, the least favorable prior solves the problem

$$\sup_{G \in \mathcal{G}} \inf_{n \geq 0, S} r(n, S, G). \quad (4)$$

Consider the following experimentation and implementation strategies:

$$n_{MMR}^* = k_1 \cdot \left(\frac{\sigma}{c}\right)^{2/3} \quad \text{and} \quad S_{MMR}^* = 1\{\hat{\delta} > 0\},$$

where k_1 is a constant approximated by 0.193, and the prior distribution G_{MMR}^* with probability mass function

$$g_{MMR}^*(\delta) = \begin{cases} \frac{1}{2} & \text{if } \delta = -k_2 \cdot \frac{\sigma}{\sqrt{n_{MMR}^*}} \\ \frac{1}{2} & \text{if } \delta = k_2 \cdot \frac{\sigma}{\sqrt{n_{MMR}^*}} \\ 0 & \text{otherwise,} \end{cases}$$

where k_2 is a constant approximated by 0.752.

Proposition 5 (Minimax Regret). *The experimentation and implementation strategy (n_{MMR}^*, S_{MMR}^*) minimizes maximum regret and the distribution G_{MMR}^* is the least favorable prior distribution.*

The least favorable prior distribution G_{MMR}^* can be interpreted as the distribution that an adversarial nature would select to minimize the firm’s expected profits. The firm’s minimax regret implementation and experimentation strategy are in turn Bayes optimal with respect to this prior.

4 Illustrative Numerical Example

We now consider an illustrative numerical example. The example shows that our optimal experimentation strategy can perform considerably better than the standard rule-of-thumb in plausible settings. We based our setting on typical experiments run by brick-and-mortar firms with relatively small sample sizes.⁶

Consider a firm with 1,000 business units. The firm can choose a number $n \leq 1000$ of identical business units to include in an experiment to test an innovation with unknown quality Δ . We define quality as the percent gain in revenue due to the innovation. The firm has a prior that quality follows a normal distribution with mean $M = -5$ and standard deviation $s = 5$.⁷ The experiment is noisy due to variance in revenue from each business unit. We assume that the experimental error is normal as in Section 2, and we set the parameter σ to 30.

We consider several specifications for the marginal cost c of experimentation. We set c so that the total cost of running an experiment with all 1,000 business units, $1000c$, is between 5% of revenue and 0.1% of revenue.⁸ Thus, we consider a range of costs spanning relatively

⁶See [Pierce et al. \(2021\)](#) for an example.

⁷In practice the prior distribution might have been estimated from data on previous experiments. See [Azevedo et al. \(2019\)](#).

⁸Consequently, c is between $.1/1000$ and $5/1000$.

high-cost experimentation and relatively low-cost experimentation. The production function for this example is illustrated in Figure 1.

Table 1 displays the optimal experimentation and implementation strategies under these different costs. In the highest-cost scenario, it is never optimal to implement the idea. When the cost of experimenting on all 1,000 business units is 1% of revenue, the optimal sample size is about 100 business units. As costs decrease, the optimal sample size increases; in the lowest-cost scenario, the optimal sample size is about 430 business units. The optimal implementation strategy is to accept the innovation if the experiment’s t -statistic exceeds a small, positive threshold. Profits, when positive, range from 0.16% of revenue to 0.33% of revenue, when the size of the experiment is chosen optimally. This is a significant number. A firm that runs ten experiments in a year would then have an expected revenue gain between 1.6% and 3.3% of revenue.

We compare our optimal implementation and experimentation strategy with the standard rule-of-thumb. The standard rule-of-thumb implementation strategy is to implement an idea if and only if it is statistically significant at the 5% level in a one-sided t -test. This means that an idea is implemented if the experiment’s t -statistic is larger than $t_{0.95} = 1.645$, which is considerably more strict than our threshold t_n^* . Table 1 reports that the t_n^* in our numerical example is small and positive, ranging from 0.57 to 0.29. Compared to our implementation strategy, the rule-of-thumb will reject a profitable innovation more frequently.

The standard rule-of-thumb experimentation strategy is to select sample size based on a power calculation. In a power calculation, the experimenter starts from a “minimum significant effect size”. The experiment size is then chosen to guarantee some minimum power if the effect is greater than the minimum significant effect size, in a one-sided 5% t -test. We calculated the power calculation sample size with a required power of 80% and minimum significant effect size of a 2.5% gain in revenue. This yielded a sample of about 890.

Table 1 compares profits under the optimal strategy with profits under the standard rules of thumb. When the experimental cost is large, the power calculation performs poorly, and

Table 1: Numerical comparison of Bayesian optimal strategy, standard power calculation rule of thumb, and minimax regret strategy.

Cost of an experiment of size 1,000	5	1	0.5	0.1
\hat{n} : lower-bound of n^* (Proposition 2)	18	18	18	18
t_n^* : optimal implementation strategy	-	0.572	0.458	0.289
n^* : optimal experimentation strategy	0	110	172	430
profits under t^* & n^*	0%	0.16%	0.23%	0.33%
$t_{1-\alpha}$: rule of thumb threshold t -statistic	1.645	1.645	1.645	1.645
n_{pc}^* : rule of thumb sample size	890	890	890	890
profits under $t_{1-\alpha}$ & n_{pc}^*	-4.097%	-0.535%	-0.09%	0.266%
profits under t^* & n_{pc}^*	-4.06%	-0.5%	-0.05%	0.3%
t_{MMR} : minimax regret threshold t -statistic	0	0	0	0
n_{MMR}^* : minimax regret sample size	64	187	296	866
profits under t_{MMR} & n_{MMR}^*	-0.22%	0.11%	0.2%	0.3%

Notes: The numerical example compares the Bayesian optimal strategy, the minimax regret strategy, and the power calculation for selecting sample size. Profits are measured in terms of fraction of gains in revenue. The parameters are prior mean $M = -5$, prior standard deviation $s = 5$, experimental noise $\sigma = 30$, and statistical significance level $\alpha = 5\%$. The implementation strategy is $S = 1\{\hat{t} > T\}$, where \hat{t} is the experiment's t -statistic and T is either the optimal threshold t_n^* (Proposition 1) or the p -value threshold $t_{1-\alpha}$ (the $(1 - \alpha)$ th percentile of a standard normal distribution). Denote by n the sample size of the experiment. Profits are $\mathbb{E}[S \cdot \Delta] - cn$.

profits under a power calculation are negative. The reason is that, in this case, the optimal sample size is relatively small, whereas the power calculation suggests running a large and costly experiment. On the other hand, under the smallest experimental cost of 0.1%, the power calculation performs well. When costs are low, it is optimal to experiment on many business units, and this is what the power calculation suggests. In this case, profits under a power calculation and the optimal implementation strategy are only 10% less than optimal. Finally, the table describes the performance of a minimax regret strategy. We find that, for our illustrative example, the minimax regret strategy performs considerably better than the power calculation rule-of-thumb, and profits are relatively close to optimal. This suggests that, as argued by [Manski and Tetenov \(2016, 2019\)](#), the minimax regret strategy can be a good candidate for practical applications.

There are two caveats to the illustrative example. First, the fact that a power calculation suggests a sample size that is too large is an artifact of the example parameters. In general, the power calculation could give a sample size either above or below the optimal experiment size. The reason is that the power calculation depends only on how noisy the experiment is (σ) and on the arbitrarily chosen minimum significant effect and power level. In contrast, the optimum depends the level of noise σ , as well as the cost and the prior. Second, throughout this paper we have maintained the assumption that the only cost of experimentation is the cost of acquiring more data. In practice, there can be other costs, such as the opportunity cost of resources that could be used to test other ideas ([Azevedo et al., 2020](#)). Numerically, these examples can look quite different. However, it is still true that the power calculations do not depend on the same variables as the optimal experimentation strategy. Therefore, the power calculation may or may not have good performance, much like in the simple case we consider.

5 Conclusion

We study the A/B testing problem in the case where the prior distribution of idea quality is Gaussian. There is a closed-form solution to the value of information obtained from an experiment (Proposition 2) that can be used in extensions to our setting, the model prescribes optimal implementation and experimentation strategies that are simple to use in practice (Propositions 1 and 3), and qualitative principles from comparative statics (Proposition 4). We also solve for the sample size that minimizes the firm's maximum regret and present the problem's least favorable prior over idea quality (Proposition 5).

We compare the Bayesian optimal strategy, the minimax regret strategy and the standard power calculation for selecting sample size. In an illustrative example, we demonstrate that these optimal strategies can considerably outperform the standard rule-of-thumb. The results suggest that when our assumptions are well-approximated, the Bayesian optimal strategy and the minimax regret strategy can be useful to improve A/B testing.

A Proofs

A.1 Proof of Proposition 1

Proof. The firm implements the idea if and only if the posterior mean quality is positive:

$$P(\hat{\delta}, n) = \theta_n \hat{\delta} + (1 - \theta_n)M > 0 \iff \hat{\delta} > -M \cdot \frac{\sigma^2/n}{s^2} = t_n^* \cdot \sigma/\sqrt{n}$$

□

A.2 Proof of Proposition 2

Proof. Denote $\delta_n^* := t_n^* \cdot \sigma/\sqrt{n}$, the unique threshold signal such that the posterior mean is zero, given n . Then, the production function equals the expected value of the innovation times the probability it is implemented; moreover, the innovation is implemented if and only if the signal exceeds δ_n^* (Azevedo et al. (2020)). Therefore,

$$\begin{aligned} f(n) &= \int \delta \cdot \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot g(\delta) d\delta - M^+ \\ &= \int (\delta - M) \cdot \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot g(\delta) d\delta + M \int \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot g(\delta) d\delta - M^+ \end{aligned}$$

The first term can be simplified to $s\sqrt{\theta_n} \cdot \phi\left(\frac{M/s}{\sqrt{\theta_n}}\right)$ using integration by parts and setting:

$$\begin{aligned} dv &= (\delta - M) \cdot g(\delta) d\delta \\ u &= \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right). \end{aligned}$$

The second term can be simplified to $M \cdot \Phi\left(\frac{M/s}{\sqrt{\theta_n}}\right)$ using the following identity for the Gaussian

distribution:

$$\int \Phi(a + bx) \phi(x) dx = \Phi\left(\frac{a}{\sqrt{1 + b^2}}\right).$$

Then, combining gives Equation 2.

The production function is strictly increasing because differentiating Equation 2 shows that:

$$f'(n) = \frac{1}{2n^2} \cdot \frac{\sqrt{\theta_n}}{s} \cdot \phi\left(\frac{m}{\sqrt{\theta_n}}\right) \cdot \sigma^2 \cdot \theta_n, \quad (5)$$

a function that is positive for all n .

The production function is bounded because the production function is strictly increasing, and as $n \rightarrow \infty$, $\theta_n \rightarrow 1$ and $f(n) \rightarrow s(\phi(m) - |m| \cdot \Phi(-|m|))$, a finite value.

From an analysis of the second derivative, the production function is convex, then concave.

Differentiating Equation 2 twice shows that:

$$f''(n) = f'(n) \cdot \left(\frac{-4s^6n^2 + (M^2 - s^2)\sigma^2s^2n + M^2\sigma^4}{s^2\theta_n}\right).$$

The sign of $f''(n)$ depends on $-4s^6n^2 + (M^2 - s^2)\sigma^2s^2n + M^2\sigma^4$, a second order polynomial over n with a negative principal coefficient. Using the quadratic formula, the smaller root $\frac{\sigma^2((M^2 - s^2) - \sqrt{M^4 + 14M^2s^2 + s^4})}{8s^4}$ is negative, since $M^2 - s^2 = \sqrt{M^4 - 2M^2s^2 + s^4} < \sqrt{M^4 + 14M^2s^2 + s^4}$. The larger of the two roots is positive and the inflection point:

$$\hat{n} := \frac{\sigma^2\left((M^2 - s^2) + \sqrt{M^4 + 14M^2s^2 + s^4}\right)}{8s^4}.$$

The production function is convex over $[0, \hat{n})$, and concave over (\hat{n}, ∞)

Finally, by plugging in 0 into Equation 5, we find that $f'(0) = 0$ when $M \neq 0$. And by taking the limit of Equation 5 as $n \rightarrow \infty$, we find that $\lim_{n \rightarrow \infty} f'(n) = 0$. \square

A.3 Proof of Proposition 3

Proof. The optimal experimentation strategy n^* maximizes $f(n) - cn$ for $n \geq 0$. The first order condition is $f'(n) = c$.

There must be at least one critical point that satisfies the first order condition, since (1) $f(0) - c \cdot 0 = 0$, since $f(0) = 0$ from Proposition 2; (2) $f(n) - c \cdot n > 0$ for some $n > 0$ by assumption; and (3) $f(n) - c \cdot n < 0$ for large n , since $f(n)$ is bounded, as shown in Proposition 2, and $c \cdot n$ is not. Further, from (1) and (2), the solution cannot be a boundary solution and must be a critical point that satisfies the first order condition.

From Proposition 2, we know that when $M \neq 0$, $f'(0) = 0$ and $\lim_{n \rightarrow \infty} f'(n) = 0$, and that $f'(n)$ is increasing over $(0, \hat{n})$ and decreasing over (\hat{n}, ∞) . From above, we know the maximum exists and must satisfy the first order condition.

If there is only one solution to the first order condition, then this must be the optimal experimentation strategy. Otherwise, there are two solutions to $f'(n) = c$, one smaller than \hat{n} and one larger. In this case, the sign analysis of $f''(n)$ from Proposition 2 shows that the smaller critical point is a local minimum and the larger critical point is a local maximum, and so the larger critical point must be the optimal experimentation strategy.

When $M = 0$, $\lim_{n \rightarrow 0} f'(n) = \infty$ and $\lim_{n \rightarrow \infty} f'(n) = 0$ by taking the limit of Equation 5. $f'(n)$ is strictly decreasing, since $\hat{n} = 0$ when $M = 0$. Therefore, $f'(n)$ crosses c once and this critical point is a local maximum, and must be the optimal experimentation strategy. \square

A.4 Proof of Proposition 4

The comparative statics results with respect to the production function $f(n)$ hold under a general prior, and we prove these results under a general prior. We prove the comparative statics results with respect to the optimal sample size n^* under a Gaussian prior.

Proof. 1. $\frac{\partial f}{\partial M} > 0$ iff $M < 0$

By Lemma A.2 in [Azevedo et al. \(2020\)](#), the production function is

$$f(n, M, s, \sigma) = \max_{\bar{\delta}} \int_{-\infty}^{+\infty} \delta \cdot \Phi\left(\frac{\delta - \bar{\delta}}{\sigma/\sqrt{n}}\right) \cdot \frac{1}{s} \cdot h\left(\frac{\delta - M}{s}\right) d\delta - M^+.$$

By the envelope theorem, we have

$$f_M(n, M, s, \sigma) = \int \delta \cdot \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot \frac{1}{s} \cdot h'\left(\frac{\delta - M}{s}\right) \cdot \frac{(-1)}{s} d\delta - \mathbb{1}\{M > 0\}.$$

Then, we set $u = \delta \cdot \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot \frac{1}{s}$ and $v = h\left(\frac{\delta - M}{s}\right)$, and integrate by parts.

$$\begin{aligned} f_M(n, M, s, \sigma) &= - \int u dv - \mathbb{1}\{M > 0\} \\ &= \int v du - (uv) \Big|_{-\infty}^{+\infty} - \mathbb{1}\{M > 0\} \\ &= \int h\left(\frac{\delta - M}{s}\right) \cdot \left\{ \frac{1}{s} \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) + \frac{\delta}{s} \cdot \phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot \frac{1}{\sigma/\sqrt{n}} \right\} d\delta - \mathbb{1}\{M > 0\} \\ &= \int h\left(\frac{\delta - M}{s}\right) \frac{1}{s} \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) d\delta - \mathbb{1}\{M > 0\}, \end{aligned}$$

where the last equality holds because $\int \delta \cdot \phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot h\left(\frac{\delta - M}{s}\right) d\delta = 0$ by definition of δ_n^* .

If $M < 0$, then $f_M(n, M, s, \sigma) > 0$ because each term inside of the integrand is positive. If $M > 0$, then since $\Phi(\cdot) \in [0, 1]$ and $h(\cdot)/s$ is a p.d.f., the integral $\int h\left(\frac{\delta - M}{s}\right) \frac{1}{s} \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) d\delta$ is smaller than 1, and so $f_M(n, M, s, \sigma) < 0$.

2. $\frac{\partial f}{\partial s} > 0$

First, we prove that, under a general prior distribution, the production function is homogeneous of degree one over the prior mean, prior standard deviation, and experimental noise:

$$f(n, M, s, \sigma) = s \cdot f(n, M/s, 1, \sigma/s). \quad (6)$$

Again, by Lemma A.2 from [Azevedo et al. \(2020\)](#), the production function is

$$\begin{aligned}
f(n, M, s, \sigma) &= \int \delta \cdot \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot \frac{1}{s} \cdot h\left(\frac{\delta - M}{s}\right) d\delta - M^+ \\
&= s \left(\int (\delta/s) \cdot \Phi\left(\frac{(\delta/s) - (\delta_n^*/s)}{(\sigma/s)/\sqrt{n}}\right) \cdot h\left((\delta/s) - (M/s)\right) d(\delta/s) - (M/s)^+ \right) \\
&= s \cdot f(n, M/s, 1, \sigma/s).
\end{aligned}$$

We have essentially re-scaled the production function, setting the prior standard deviation to 1. Then, the scaled prior mean is M/s and the scaled experimental noise is σ/s .

Next, we show that the production function is decreasing over the experimental noise:

$$f(n, M, s, \sigma) > f(n, M, s, \sigma').$$

The formula for the production function shows that σ and n each appear only once, and it must be that

$$f(n, M, s, \sigma) = f(\lambda^2 \cdot n, M, s, \lambda \cdot \sigma).$$

Then, take any $\sigma' > \sigma$ and denote $\lambda := \sigma'/\sigma > 1$. It follows that $f(n, M, s, \sigma) = f(\lambda^2 \cdot n, M, s, \sigma') > f(n, M, s, \sigma')$, where the inequality is true because the production function is increasing over n (by Proposition 2).

Finally, we prove the comparative statics result. By equation (6) and the chain rule, $f_s(n, M, s, \sigma)$ is equal to

$$f(n, M/s, 1, \sigma/s) + f_M(n, M/s, 1, \sigma/s) \cdot \left(\frac{-M}{s^2}\right) + f_\sigma(n, M/s, 1, \sigma/s) \cdot \left(\frac{-\sigma}{s^2}\right).$$

The first term is positive since by Proposition 2, the production function is always positive. The second term is positive since M and f_M have opposite signs (as proven

above in the first comparative statics). Finally, the third term is positive because the production function is decreasing over the experimental noise.

$$3. \frac{\partial n^*}{\partial M} > 0 \text{ iff } M < 0$$

Proposition 3 shows that the optimal experimental scale n^* is the largest solution to the first order condition, $f'(n) = c$. By the implicit function theorem, we have

$$\frac{\partial n^*}{\partial M} = -\frac{f_{nM}(n^*)}{f_{nn}(n^*)}.$$

Since n^* is greater than the inflection point defined in Proposition 2, $f_{nn}(n^*) < 0$. and the comparative statics for M is described by the sign of $f_{nM}(n^*)$.

In the proof of Proposition 2, we derived the marginal product,

$$f'(n) = \frac{1}{2n^2} \cdot \frac{\sqrt{\theta_n}}{s} \cdot \phi\left(\frac{M/s}{\sqrt{\theta_n}}\right) \cdot \sigma^2 \cdot \theta_n.$$

Taking the derivative with respect to M , we obtain

$$f_{nM}(n) = f'(n) \cdot \frac{-M}{s^2 \theta_n}.$$

Since $f'(n)$ is always positive, we conclude that for any n , $f_{nM}(n) > 0$ if and only if $M < 0$.

$$4. \frac{\partial n^*}{\partial s} > 0 \text{ if } |M| > s$$

$$> 0 \text{ if } |M| < s \text{ and } f'(\tilde{n}) > c$$

$$< 0 \text{ if } |M| < s \text{ and } f'(\tilde{n}) < c,$$

where

$$\tilde{n} = \frac{\sigma^2 (3m^2 + 2 + \sqrt{m^4 + 20m^2 + 4})}{2s^2 (1 - m^2)}.$$

Proposition 3 shows that the optimal experimental scale n^* is the largest solution to

the first order condition, $f'(n) = c$. By the implicit function theorem, we have

$$\frac{\partial n^*}{\partial s} = -\frac{f_{ns}(n^*)}{f_{nn}(n^*)}.$$

Since n^* is greater than the inflection point defined in Proposition 2, $f_{nn}(n^*) < 0$ and the comparative statics of s is described by the sign of $f_{ns}(n^*)$.

In the proof of Proposition 2, we derived the marginal product,

$$f'(n) = \frac{1}{2n^2} \cdot \frac{\sqrt{\theta_n}}{s} \cdot \phi\left(\frac{M/s}{\sqrt{\theta_n}}\right) \cdot \sigma^2 \cdot \theta_n,$$

Taking the derivative with respect to s , we obtain

$$f_{ns}(n) = f'(n) \cdot \frac{M^2(2 - \theta_n) + s^2\theta_n(2 - 3\theta_n)}{s^3\theta_n}.$$

The sign of $f_{ns}(n)$ is the same as the sign of the following quadratic polynomial over θ_n :

$$Q(\theta_n) \equiv -3s^2\theta_n^2 + (2s^2 - M^2)\theta_n + 2M^2,$$

which has two roots $\tilde{\theta}_1 < \tilde{\theta}_2$,

$$\tilde{\theta}_1 = \frac{2s^2 - M^2 - \sqrt{M^4 + 20M^2s^2 + 4s^4}}{6s^2} \quad \text{and} \quad \tilde{\theta}_2 = \frac{2s^2 - M^2 + \sqrt{M^4 + 20M^2s^2 + 4s^4}}{6s^2},$$

It must be that $\tilde{\theta}_1 < 0$, since $2s^2 - M^2 = \sqrt{M^4 - 4M^2s^2 + 4s^4} < \sqrt{M^4 + 20M^2s^2 + 4s^4}$, and since $Q(\theta)$ is a quadratic polynomial with a negative principle coefficient, $Q(\theta) > 0$ if and only if $\theta \in (\tilde{\theta}_1, \tilde{\theta}_2)$.

Recall that $\theta_n = s^2/(s^2 + \sigma^2/n)$ for $n > 0$ and $\theta_0 = 0$, so $\theta_n \in [0, 1)$. Therefore, the sign of $Q(\theta_n)$ will depend on whether θ_n is greater than or less than $\tilde{\theta}_2$.

If $M^2 > s^2$, then $\tilde{\theta}_2 > 1$. In this case, $Q(\theta_n) > 0$ for all n since $\theta_n \in [0, 1) \subseteq (\tilde{\theta}_1, \tilde{\theta}_2)$.

This implies that $f_{ns}(n) > 0$ if $M^2 > s^2$, which proves the first part of this comparative

statics result, since dn^*/ds and $f_{ns}(n) > 0$ share the same sign.

If $M^2 < s^2$, then $\tilde{\theta}_2 < 1$. This implies that for any $\theta \in [0, \tilde{\theta}_2)$, we have $Q(\theta) > 0$, and for any $\theta \in (\tilde{\theta}_2, 1)$, we have $Q(\theta) < 0$. Since $\theta_n = s^2/(s^2 + \sigma^2/n)$ is an increasing function over n that takes $n \in [0, \infty)$ into $\theta_n \in [0, 1)$, there exists a \tilde{n} such that $\theta_{\tilde{n}} = \tilde{\theta}_2$. Solving for \tilde{n} shows that

$$\tilde{n} = \frac{\sigma^2 (3m^2 + 2 + \sqrt{m^4 + 20m^2 + 4})}{2s^2 (1 - m^2)}, \text{ where } m = \frac{M}{s}.$$

Finally, note that n^* solves the first order condition $f'(n) = c$, and is in the decreasing region of $f'(n)$ by Proposition 2. Therefore, when $f'(\tilde{n}) > c$, $n^* < \tilde{n}$, which implies that $\theta_{n^*} < \theta_{\tilde{n}} = \tilde{\theta}_2$ and $Q(\theta_{n^*}) > 0$. And conversely, when $f'(\tilde{n}) < c$, $n^* > \tilde{n}$, which implies that $\theta_{n^*} > \theta_{\tilde{n}} = \tilde{\theta}_2$ and $Q(\theta_{n^*}) < 0$. This shows the remaining parts of this comparative statics result, since dn^*/ds and $Q(\theta_{n^*})$ share the same sign when $|M| < s$.

□

A.5 Proof of Proposition 5

Proof. Let us recall the definition of regret (3):

$$\mathcal{R}(n, S, \delta) = \delta^+ - (\delta \cdot \mathbb{E}[S] - c \cdot n).$$

The minimax regret strategy (n_{MMR}^*, S_{MMR}^*) is the solution to

$$\inf_{n \geq 0, S} \sup_{\delta} \mathcal{R}(n, S, \delta).$$

This problem can be solved using a game theoretic approach (See [Stoye \(2009\)](#) and the references therein). We can imagine a simultaneous zero-sum game between the firm and an adversarial nature. The firm chooses an implementation and experimentation strategy,

(n, S) to minimize regret $\mathcal{R}(n, S, \delta)$. Nature chooses a prior distribution for δ , G , to maximize regret $\mathcal{R}(n, S, \delta)$.

This approach allows us to solve the minimax regret problem defined above using guess and verify. We will show that the following strategy profile is a Nash Equilibrium, where $k \approx -0.752$ is the unique negative solution to $\Phi(k) + k\phi(k) = 0$:

1. The firm chooses

$$n_{MMR}^* = \left[\frac{1}{2} k^2 \cdot \phi(k) \cdot \frac{\sigma}{c} \right]^{2/3} \quad \text{and} \quad S_{MMR}^* = 1\{\hat{\delta} > 0\},$$

2. Nature chooses the prior distribution G_{MMR}^* with probability mass function

$$g_{MMR}^*(\delta) = \begin{cases} \frac{1}{2} & \text{if } \delta = -k \cdot \frac{\sigma}{\sqrt{n_{MMR}^*}} \\ \frac{1}{2} & \text{if } \delta = k \cdot \frac{\sigma}{\sqrt{n_{MMR}^*}} \\ 0 & \text{otherwise.} \end{cases}$$

We will verify that each player's strategy satisfies their best response condition:

$$\begin{aligned} (n_{MMR}^*, S_{MMR}^*) &\in \arg \inf_{n \geq 0, S} \int \mathcal{R}(n, S, \delta) dG_{MMR}^* \\ \iff (n_{MMR}^*, S_{MMR}^*) &\in \arg \sup_{n \geq 0, S} \int u(n, S; \delta) dG_{MMR}^*, \end{aligned}$$

and

$$\delta^* \in \arg \sup_{\delta} R(n_{MMR}^*, S_{MMR}^*, \delta) \quad \text{for all } \delta^* \text{ in the support of } G_{MMR}^*.$$

Part 1: The Firm's Best Response. The firm implements the idea if and only if the posterior mean quality is positive. When the prior is G_{MMR}^* , using equation (A.2) in [Azevedo](#)

et al. (2020), the firm implements the idea if and only if

$$\int_{-\infty}^{\infty} \delta \cdot \phi(\hat{\delta} \mid \delta, \sigma^2/n) \cdot g_{MMR}^*(\delta) d\delta \geq 0$$

$$\iff \frac{1}{2} k \frac{\sigma}{\sqrt{n_{MMR}^*}} \left\{ \phi\left(\hat{\delta} \mid k \frac{\sigma}{\sqrt{n_{MMR}^*}}, \frac{\sigma^2}{n}\right) - \phi\left(\hat{\delta} \mid -k \frac{\sigma}{\sqrt{n_{MMR}^*}}, \frac{\sigma^2}{n}\right) \right\} \geq 0.$$

Since $k < 0$, the inequality above is equivalent to

$$\phi\left(\hat{\delta} \mid -k \frac{\sigma}{\sqrt{n_{MMR}^*}}, \frac{\sigma^2}{n}\right) \geq \phi\left(\hat{\delta} \mid k \frac{\sigma}{\sqrt{n_{MMR}^*}}, \frac{\sigma^2}{n}\right)$$

$$\iff -\frac{\left(\hat{\delta} + k\sigma/\sqrt{n_{MMR}^*}\right)^2}{2\sigma^2/n} \geq -\frac{\left(\hat{\delta} - k\sigma/\sqrt{n_{MMR}^*}\right)^2}{2\sigma^2/n}$$

$$\iff -k\hat{\delta} \geq k\hat{\delta}.$$

Since $k < 0$, the posterior mean quality is positive if and only if $\hat{\delta}$ is positive. Therefore,

$$S_{MMR}^* = 1\{\hat{\delta} > 0\}$$

satisfies the firm's best response condition.

Next, the optimal experimentation strategy in the A/B testing problem equates the marginal product and the marginal cost (see Proposition 3). This result holds under a general prior. Lemma A.2 from Azevedo et al. (2020) gives the formula for the marginal product for any prior distribution. For the least favorable prior G_{MMR}^* , the marginal product is

$$\frac{1}{2n} \int \delta^2 \cdot \phi(\delta^* \mid \delta, \sigma^2/n) \cdot g_{MMR}^*(\delta) d\delta$$

where δ^* is the $\hat{\delta}$ that sets the posterior mean equal to 0. From above, $\delta^* = 0$ under G_{MMR}^* .

Therefore, the marginal product is equal to

$$\begin{aligned} & \frac{1}{2n} k^2 \cdot \frac{\sigma^2}{n_{MMR}^*} \cdot \phi \left(0 \mid -k \frac{\sigma}{\sqrt{n_{MMR}^*}}, \sigma^2/n \right) \\ &= \frac{1}{2\sqrt{n}} \cdot k^2 \cdot \frac{\sigma}{n_{MMR}^*} \cdot \phi \left(k \cdot \left(\frac{n}{n_{MMR}^*} \right)^{1/2} \right). \end{aligned}$$

The marginal product above is decreasing over n (recall that $k < 0$). This implies that there is a unique value for n such that marginal product equals the constant marginal cost c . We can verify that $n = n_{MMR}^*$ is the solution. Substituting $n = n_{MMR}^*$ into the formula above, we have

$$(n_{MMR}^*)^{-3/2} \cdot \frac{1}{2} \cdot k^2 \cdot \sigma \cdot \phi(k).$$

This is equal to c since

$$n_{MMR}^* = \left[0.5k^2 \phi(k) \frac{\sigma}{c} \right]^{2/3}.$$

Therefore, n_{MMR}^* satisfies the firm's best response condition. This concludes the first part of the proof.

Part 2: Nature's Best Response. When the firm plays (n_{MMR}^*, S_{MMR}^*) , the expected payoff for the firm is

$$\begin{aligned} u(n_{MMR}^*, S_{MMR}^*; \delta) &= \delta \cdot \mathbb{E}[S_{MMR}^*] - c \cdot n_{MMR}^* \\ &= \delta \cdot \mathbb{E}[1\{\hat{\delta} > 0\}] - c \cdot n_{MMR}^* \\ &= \delta \Phi \left(\frac{\delta \cdot \sqrt{n_{MMR}^*}}{\sigma} \right) - c \cdot n_{MMR}^*, \end{aligned}$$

since $\hat{\delta} \sim \mathcal{N}(\delta, \sigma^2/n_{MMR}^*)$. Then, regret —nature's payoff— is

$$\mathcal{R}(n_{MMR}^*, S_{MMR}^*, \delta) = \delta^+ - \left(\delta \Phi \left(\frac{\delta \cdot \sqrt{n_{MMR}^*}}{\sigma} \right) - c \cdot n_{MMR}^* \right).$$

Denoting $x := \delta \cdot \sqrt{n_{MMR}^*}/\sigma$, this is equal to

$$\frac{\sigma}{\sqrt{n_{MMR}^*}} \left(x^+ - x\Phi(x) + c \cdot \frac{(n_{MMR}^*)^{3/2}}{\sigma} \right).$$

The expression above is symmetric over x , since for any $x > 0$, $x^+ - x\Phi(x) = x(1 - \Phi(x)) = (-x)^+ - (-x)\Phi(-x)$. For $x < 0$, regret is equal to

$$\frac{\sigma}{\sqrt{n_{MMR}^*}} \left(-x \cdot \Phi(x) + c \cdot \frac{(n_{MMR}^*)^{3/2}}{\sigma} \right).$$

Nature chooses $x < 0$ to maximize this expression. The second derivative is $\phi(x)(x^2 - 2)$, and so the expression is strictly convex over $(-\infty, -\sqrt{2})$, and strictly concave over $(-\sqrt{2}, 0)$. Since $\lim_{x \rightarrow -\infty} -x \cdot \Phi(x) = 0 < -k \cdot \Phi(k) \approx 0.169$, there is a unique maximum that is an interior solution. Therefore, $x = k \approx -0.752$ is the unique solution, where we had defined k to solve the first order condition of this expression.

By symmetry, for $x > 0$, regret is maximized at $x = -k$, and the maximal values at the two solutions are equal.

Since $x := \delta \cdot \sqrt{n_{MMR}^*}/\sigma$, the values $k\sigma/\sqrt{n_{MMR}^*}$ and $-k\sigma/\sqrt{n_{MMR}^*}$ both satisfy nature's best response condition. These are the two points in the support of nature's prior distribution G_{MMR}^* , which concludes the second part of the proof. \square

References

Athey, S. and G. W. Imbens

2017. The econometrics of randomized experiments. In *Handbook of economic field experiments*, volume 1, Pp. 73–140. Elsevier.

Azevedo, E. M., D. Alex, J. Montiel Olea, J. M. Rao, and E. G. Weyl

2020. A/b testing with fat tails. *Journal of Political Economy*, 128 (12):4614–4672.

Azevedo, E. M., A. Deng, J. L. Montiel Olea, and E. G. Weyl

2019. Empirical bayes estimation of treatment effects with many a/b tests: An overview. In *AEA Papers and Proceedings*, volume 109, Pp. 43–47.

Bross, I.

1950. Two-choice selection. *Journal of the American Statistical Association*, 45(252):530–540.

Chade, H. and E. Schlee

2002. Another look at the radner–stiglitz nonconcavity in the value of information. *Journal of Economic Theory*, 107(2):421–452.

Grundy, P., M. Healy, and D. Rees

1956. Economic choice of the amount of experimentation. *Journal of the Royal Statistical Society: Series B (Methodological)*, 18(1):32–49.

Keppo, J., G. Moscarini, and L. Smith

2008. The demand for information: More heat than light. *Journal of Economic Theory*, 138(1):21–50.

List, J. A., S. Sadoff, and M. Wagner

2011. So you want to run an experiment, now what? some simple rules of thumb for optimal experimental design. *Experimental Economics*, 14(4):439.

Manski, C. F.

2019. Treatment choice with trial data: Statistical decision theory should supplant hypothesis testing. *The American Statistician*, 73(sup1):296–304.

Manski, C. F. and A. Tetenov

2016. Sufficient trial size to inform clinical practice. *Proceedings of the National Academy of Sciences*, 113(38):10518–10523.

Manski, C. F. and A. Tetenov

2019. Trial size for near-optimal choice between surveillance and aggressive treatment: Reconsidering mslt-ii. *The American Statistician*, 73(sup1):305–311.

Meltzer, D.

2001. Addressing uncertainty in medical cost–effectiveness analysis: implications of expected utility maximization for methods to perform sensitivity analysis and the use of cost–effectiveness analysis to set priorities for medical research. *Journal of health economics*, 20(1):109–129.

Moscarini, G. and L. Smith

2002. The law of large demand for information. *Econometrica*, 70(6):2351–2366.

Pierce, L., A. Rees-Jones, and C. Blank

2021. The negative consequences of loss-framed performance incentives. Technical report, National Bureau of Economic Research.

Radner, R. and J. Stiglitz

1984. A nonconcavity in the value of information. *Bayesian models in economic theory*, 5:33–52.

Raiffa, H. and R. Schlaifer

1961. Applied statistical decision theory.

Savage, L.

1951. The theory of statistical decision. In *Journal of the American Statistical Association*, volume 89, Pp. 55–67.

Somerville, P. N.

1954. Some problems of optimum sampling. *Biometrika*, 41(3/4):420–429.

Stoye, J.

2009. Minimax regret treatment choice with finite samples. *Journal of Econometrics*, 151(1):70–81.