The A/B Testing Problem with Gaussian Priors

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

December 15, 2020

Abstract

The A/B testing problem (Azevedo et al., 2020) considers how to optimally use scarce experimental resources to screen potential innovations with unknown quality. Azevedo et al. (2020) derive asymptotic results for the value of experimentation under general priors for innovation quality. In this paper, we consider the special case of Gaussian priors. We find two main results. First, there is a closed-form solution for the value of running a randomized experiment. Second, under certain assumptions on costs, there is a simple solution to the experiment's optimal size. We show, in numerical examples, that the firm's expected profits under the optimal experimentation strategy can be higher than under standard rules of thumb for choosing sample size in randomized experiments (such as power calculations).

*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. The most 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. Experimentation on ideas and innovations has also grown more common for public policy evaluation and in academic research from anti-poverty interventions to behavioral nudges.

Azevedo et al. (2020) analyzed the "A/B Testing Problem" of how a firm should best use scarce experimental resources to screen potential innovations. In their set-up, a firm has a set of potential ideas of unknown quality, and can perform experiments to learn about each idea, subject to some limitations. The goal is to maximize the expected sum of quality of implemented ideas. For example, if the firm runs a search engine, the ideas are potential improvements developed by engineers, quality is some key performance measure, and the experiments are A/B tests. They take an empirical Bayes approach, where past experimental outcomes are used to learn about the distribution of unknown quality. Their main theoretical finding is that the optimal experimentation strategy crucially depends on the thickness of the tail of the distribution of quality. They then show how to apply the model to a firm, using data from experiments on Microsoft's Bing search engine.

Our contribution is to specialize the A/B Testing Problem to the case where the prior distribution of idea quality is Gaussian. We have two findings. First, there is a closed-form solution to the value of the information obtained from an experiment (Proposition 2).[1] Second, if we specialize the model to have a fixed cost of experimentation plus a linear cost

[1] The value of information—Radner and Stiglitz (1984); Meltzer (2001); Chade and Schlee (2002)—is what Azevedo et al. call the production function of data, and plays a crucial role in choosing an experimentation strategy.

in sample size, we can make clear predictions about optimal sample sizes for the A/B tests (proposition 3). The standard approach to select a sample size for an A/B test is to perform a power calculation.[2] However, power calculations have been criticized as being sub optimal in consistent models, both from Bayesian (Meltzer (2001)) and non-Bayesian perspectives (Manski and Tetenov (2016, 2019)).[3] Indeed, we find that our optimal sample sizes can considerably outperform standard power calculations. Therefore, the proposed method can be both superior to standard power calculations and simple to implement, especially in data-rich environments such as experimentation in online firms, where it is easy to estimate the prior from past experiments.[4]

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.

[2] List et al. (2011); Athey and Imbens (2017).

[3] Manski and Tetenov (2016, 2019) consider the minimax regret approach as an alternative to power calculations to select sample size for clinical trials. In Appendix B, we solve the A/B testing problem using the minimax regret approach. Proposition 4 gives a closed form solution for optimal sample size under this approach.

[4] See Azevedo et al. (2019) for an empirical Bayes approach to estimate the prior distribution.

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, $\widehat{\Delta}$. 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 profit. This model is a special case of Azevedo et al. (2020). The key restriction is that we consider a normal prior. In contrast, they allow for general thin- and fat-tailed 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.

2.2 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 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.

Lemma 1 (Azevedo et al. (2020))**.** *The optimal implementation strategy is to implement the idea iff $P(\hat{\delta}, n) > 0$. If the firm uses an optimal implementation strategy, then its payoff equals the production function minus the cost of data plus the value of the idea with no data:*

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

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, $\widehat{\delta}$, and the prior, M:

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

From this formula we obtain:

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 to 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.

Our main result is that, with a Gaussian prior, there is a closed-form solution for the production function.

Proposition 2 (Production Function)**.** *The production function is*

$$f(n) = s \cdot \left(\sqrt{\theta_n} \cdot \phi\left(\frac{m}{\sqrt{\theta_n}}\right) - |m| \cdot \Phi\left(\frac{-|m|}{\sqrt{\theta_n}}\right) \right) \tag{1}$$

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

This function is bounded, increasing, convex in an interval $[0, \hat{n}]$*, and concave in an interval* $[\hat{n}, \infty)$*. The inflection point* $\hat{n}$ *is*

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

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

This expression can be used for practical experimental design, such as in determining optimal sample sizes. Figure 1 depicts an example production function. Even in settings with multiple ideas and more general costs of experimentation, the production function is one of the key determinants of the optimal experimentation strategy, and the formula can be applied to these problems (see Azevedo et al. (2020) section 5.2).

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 $n^* = 0$.[5] When the value of information exceeds the cost of obtaining information for

[5] It is sufficient to verify this condition for $\tilde{n}$ - i.e. $c \cdot \tilde{n} > f(\tilde{n})$-, where $\tilde{n}$ is the sample size that maximizes the average product. In particular, $\tilde{n}$ can be defined as the solution of $f'(n) = f(n)/n$. This can be solved numerically using the closed-form solution for the production function (Proposition 2).

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

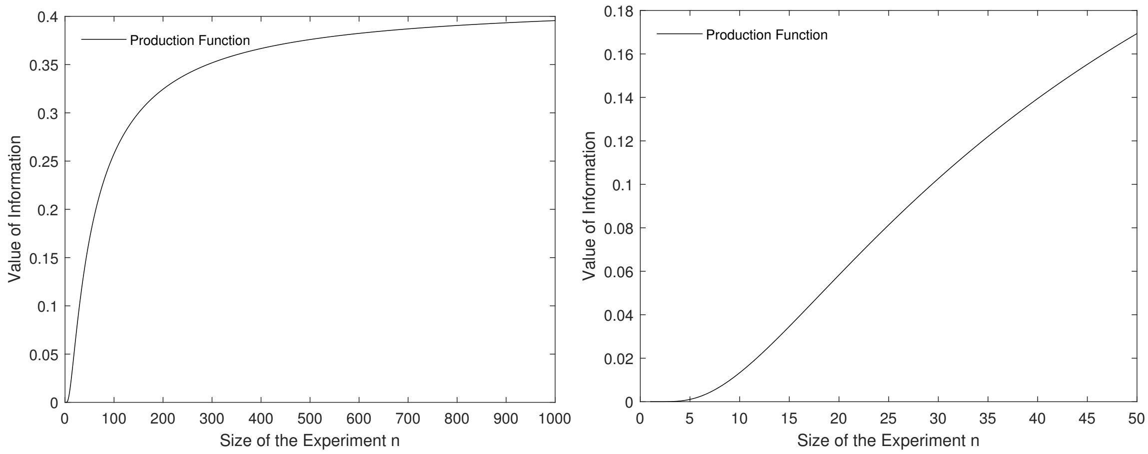

(a) n from 0 to 1000 (b) n close to 0

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

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). *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, is increasing 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.

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]

[6] See Pierce et al. (2020) for an example.

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

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 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

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

Table 1: Numerical comparison of optimal strategy, standard power calculation rule of thumb, and min-max regret.

Cost of an experiment of size 1,000	5	1	0.5	0.1
$\hat{n}$: lower-bound of n^* (Proposition 3)	18	18	18	18
t_n^*: optimal implementation strategy	-	0.572	0.458	0.289
n^*: optimal experimentation strategy	0	110	172	430
profit 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
profit under $t_{1-\alpha}$ & n_{pc}^*	-4.097%	-0.535%	-0.09%	0.266%
profit under t^* & n_{pc}^*	-4.06%	-0.5%	-0.05%	0.3%
t_{MMR}: min-max regret threshold t-statistic	0	0	0	0
n_{MMR}^*: min-max regret sample size	64	187	296	866
profit under t_{MMR} & n_{MMR}^*	-0.22%	0.11%	0.2%	0.3%

Notes: The example compares the optimal strategy, power calculation, and min-max regret experiment designs for a numerical example. Profits are measured as 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. Profit is $\mathbb{E}[S \cdot \Delta] - cn$. See Appendix B for calculations of t_{MMR} and n_{MMR}^*.

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 profit under a power calculation is 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, profit under a power calculation and the optimal implementation strategy is only 10% less than optimal.

Finally, the table describes the performance of a minimax regret strategy. This strategy minimizes the maximum regret for any value of true quality Δ (see Appendix B for the definition). The minimax regret strategy has been proposed by Savage (1951) as a guideline to make decisions under uncertainty, and is perhaps the most prominent existing alternative for determining sample sizes without considering any additional assumptions (see Manski (2019)). We find that, for our illustrative example, the minimax regret strategy performs considerably better than the power calculation rule-of-thumb, and profit is 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. Moreover, the minimax regret approach has the advantage of not depending on a prior, so it can be particularly useful in cases where there is no data from previous experiments to reliably estimate a prior.[9]

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,

[9]See Azevedo et al. (2019) for a practical guide on how to estimate the prior distribution using data on previous experiments.

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, and the model prescribes an optimal implementation and experimentation strategy that is simple to use in practice (Propositions 1 and 3).

We compare our optimal strategy with the standard rule-of-thumb. The standard approach to select a sample size is to perform a power calculation, but this approach has been criticized as arbitrary and sub-optimal from Bayesian and non-Bayesian perspectives. In an illustrative example, we demonstrate that our optimal strategy can considerably outperform the standard rule-of-thumb. The results suggest that when our assumptions are well-approximated, our optimal strategy can be useful to improve A/B testing.

A Proofs

A.1 Proof of Lemma 1

Proof. See Azevedo et al. (2020), Proposition 2. □

A.2 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.3 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; and, 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 $s\sqrt{\theta_n} \cdot \phi(\frac{M/s}{\sqrt{\theta_n}})$ 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 $M \cdot \Phi\big(\frac{M/s}{\sqrt{\theta_n}}\big)$ using the following identity for the Gaussian distribution:

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

Then, combining gives Equation 1.

The production function is strictly increasing because differentiating Equation 1 shows that:

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

a function that is positive for all n.

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

From an analysis of the second derivative, the production function is convex, then concave. Differentiating Equation 1 twice shows that:

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

The sign of $f''(n)$ depends on $-4s^6n^2 + (M^2 - s^2)\sigma^2 s^2 n + M^2\sigma^4$, a second order polynomial with a negative principal coefficient. Using the quadratic formula, the smaller root is $\frac{\sigma^2\left((M^2-s^2)-\sqrt{M^4+14M^2s^2+s^4}\right)}{8s^4}$ 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:

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

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

Finally, by plugging in 0 into Equation 2, we find that $f'(0) = 0$ when $M \neq 0$. And by taking the limit of Equation 2 as $n \to \infty$, we find that $\lim_{n\to\infty} f'(n) = 0$. $\square$

A.4 Proof of Proposition 3

Proof. By Lemma 1, 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 $\lim_{n\to\infty} f'(n) = 0$, and that $f'(n)$ is increasing over $(0, \hat{n})$ and decreasing over $(\hat{n}, \infty)$. 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 $M = 0$, $\lim_{n\to 0} f'(n) = \infty$ and $\lim_{n\to\infty} f'(n) = 0$ by taking the limit of Equation 2. $f'(n)$ is strictly decreasing, since $\hat{n} = 0$ when $M = 0$. Therefore, $f'(n)$ crosses c once this critical point is a local maximum, and must be the optimal experimentation strategy. $\square$

B Minimax Regret calculations

Proposition 4 (Strategy to minimize the maximal regret)**.** *The strategy that minimizes the maximal regret in the A/B testing problem has a closed form solution:*

$$S^*_{MMR} = 1\{\hat{\delta} > 0\} \quad \textit{and} \quad n^*_{MMR} = \Big[-0.5 x^*_0 \, \Phi(x^*_0) \, \frac{\sigma}{c} \Big]^{2/3},$$

where x^*_0 is the unique, negative solution to $\Phi(x) = -x\phi(x)$, and Φ and ϕ are the CDF and PDF of a standard normal distribution.

Proof. The expected payoff of the firm is

$$\Delta \, \mathbb{E}[S] - cn,$$

where the expectation is taken over the experimental noise, and $S \in \{0, 1\}$ is the implementation strategy that depends on the result of the experiment $\hat{\delta}_n$. In the minimax regret approach, Δ is not a random variable, but an unobserved value. Since $\hat{\delta}_n \sim \mathcal{N}(\Delta, \sigma_n^2)$, Lemma 3 in Karlin and Rubin (1956) implies that (admissible) implementation strategies must be monotone decision rules. Therefore, $S = 1\{\hat{\delta}_n \geq \psi\}$ for some $\psi \in \mathbb{R}$. Then, the expected payoff of the firm is:

$$\Delta \, \Phi\left(\frac{\Delta - \psi}{\sigma_n}\right) - cn.$$

The firm's regret is defined as:

$$R(\Delta, \psi, n) := \Delta^+ - \left(\Delta \, \Phi\left(\frac{\Delta - \psi}{\sigma_n}\right) - cn\right),$$

the optimal expected payoff if Δ were observable, minus the expected payoff attained from choosing (ψ, n).

The maximal regret from the strategy (ψ, n) is:

$$\max_{\Delta} R(\Delta, \psi, n),$$

which can be expressed as:

$$\max\{\max_{\Delta\geq 0} R(\Delta, \psi, n),\ \max_{-\Delta>0} R(-\Delta, -\psi, n)\}, \tag{3}$$

since $R(\Delta, \psi, n) = R(-\Delta, -\psi, n)$ if $\Delta < 0$. Then, it is sufficient to solve

$$\max_{\Delta\geq 0} R(\Delta, \psi, n). \tag{4}$$

For $\Delta \geq 0$, we have $R(\Delta, \psi, n) = \Delta\Phi((\psi - \Delta)/\sigma_n) + cn$. Further, a change of variable $x = (\psi - \Delta)/\sigma_n$ in gives the maximization problem:

$$\max_{x\leq\psi/\sigma_n} (\psi - x\sigma_n)\,\Phi(x) + cn.$$

Denote by $G(x)$ the derivative of the objective function with respect to x. This is:

$$G(x) := -\sigma_n\,\Phi(x) + (\psi - x\sigma_n)\phi(x).$$

Define $x_1^* := [\psi/\sigma_n - \sqrt{\psi^2/\sigma_n^2 + 4}]/2$ and $x_2^* := [\psi/\sigma_n + \sqrt{\psi^2/\sigma_n^2 + 4}]/2$. Note that $x_1^* < \psi/\sigma_n < x_2^*$, and $x_1^* < 0$.

Claim 1: First, $G(x) > 0$ if $x \leq x_1^*$. Second, there exists a $x^* \in (x_1^*, \psi/\sigma_n)$ such that $G(x^*) = 0$. Third, x^* is the unique solution to the maximization problem.

Proof. First, we prove part 1. It can be shown that $\Phi(x) < \frac{-1}{x}\phi(x)$ for $x < 0$ (since $x \leq x_1^* < 0$). Then, $G(x)$ is strictly bounded from below by

$$\frac{\sigma_n}{x}\phi(x) + (\psi - x\sigma_n)\phi(x) = \phi(x)\frac{-\sigma_n}{x}Q(x),$$

where $Q(x) = -1 - \frac{\psi}{\sigma_n}x + x^2$, a polynomial with two roots, x_1^* and x_2^*. Since $x \leq x_1^*$, it follows that $Q(x) \geq 0$.

For the second part, observe that $G(x_1^*) > 0$ and $G(\psi/\sigma_n) = -\sigma_n\Phi(\psi/\sigma_n) < 0$. Then, there exists a zero in $(x_1^*, \psi/\sigma_n)$.

Finally, to prove the third part, note that $G'(x) = \big(-2\sigma_n + (-x)(\psi - x\sigma_n)\big)\phi(x)$. The sign of $G'(x)$ depends on $\sigma_n x^2 - \psi x - 2\sigma_n$, a second order polynomial over x. Since $G'(\psi/\sigma_n) < 0$, it must be that $G(x)$ is increasing, then decreasing over the area relevant for maximization. From the first part of the claim, $G(x) > 0$ for small x, and so the solution to the maximization problem is unique. □

Denote by $x^*(\psi, n)$ this optimal solution, and define $\Delta(\psi, n) := \psi - x^*(\psi, n)\sigma_n$.

Claim 2: $R(\Delta(-\psi, n), -\psi, n) \leq R(\Delta(\psi, n), \psi, n)$ if and only if $\psi \geq 0$.

Proof. By definition of R and $\Delta(-\psi, n) \geq 0$, it follows

$$\begin{aligned} R(\Delta(-\psi, n), -\psi, n) &= \Delta(-\psi, n)\Phi\Big(\frac{-\psi - \Delta(-\psi, n)}{\sigma_n}\Big) + cn \\ &\leq \Delta(-\psi, n)\Phi\Big(\frac{\psi - \Delta(-\psi, n)}{\sigma_n}\Big) + cn \\ &= R(\Delta(-\psi, n), \psi, n), \end{aligned}$$

where the inequality above holds if and only if $\psi \geq 0$. Since $\Delta(\psi, n)$ maximizes regret by definition, we have $R(\Delta(-\psi, n), \psi, n) \leq R(\Delta(\psi, n), \psi, n)$. Moreover, equality holds only if $\psi = 0$. □

Claim 2 and Equation (3) imply that maximum regret is equal to $R(\Delta(|\psi|, n), |\psi|, n)$. This means that it is sufficient to restrict the monotone strategies to $\psi \geq 0$ to study the minimization of maximum regret.

The minimax regret implementation strategy is the optimal $\psi^*(n) \geq 0$ that solves

$$\min_{\psi \geq 0} R(\Delta(\psi, n), \psi, n),$$

where n is assumed to be given. We can find the derivative of this objective function by taking the derivative of Equation (4) and applying the envelope condition. This gives us $\frac{\Delta}{\sigma_n}\phi\left(\frac{\psi-\Delta(\psi,n)}{\sigma_n}\right)$. Then, since $G(x) = 0$ must hold at the maximal regret, the derivative is equal to

$$\Phi\left(\frac{\psi - \Delta(\psi, n)}{\sigma_n}\right),$$

which is positive. Therefore, for any n, the minimization of the maximum regret is achieved at the left boundary, $\psi^*(n) = 0$.

By Claim 1 and the definition of $\Delta(\psi, n)$, we have $\Delta(0, n) = -x_0^*\sigma_n$ where $x_0^* := x^*(0, n)$ is a negative constant that does not depend on n, since it is the unique solution to $\Phi(x) = -x\phi(x)$.

Then, for a given n and strategy $(\psi^*(n), n)$, the maximum regret is equal to $R(-x_0^*\sigma_n, 0, n)$. This is equal to

$$\frac{\sigma}{\sqrt{n}}\Phi(x_0^*)(-x_0^*) + cn,$$

which is a convex function on n. We can use this expression to find the optimal sample size that minimize the maximum regret, n^*_{MMR}. By first order condition, it follows that:

$$n^*_{MMR} \equiv \Big[-0.5x_0^* \, \Phi(x_0^*) \, \frac{\sigma}{c} \Big]^{2/3}.$$

Thus, $S^*_{MMR} = 1\{\hat{\delta}_n \geq 0\}$ and n^*_{MMR} are the minimax regret strategies that solve the A/B testing problem using the minimax regret approach. □

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.

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.

Karlin, S. and H. Rubin

1956. The theory of decision procedures for distributions with monotone likelihood ratio. In *The Annals of Mathematical Statistics*, volume 27, Pp. 272–299.

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.

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

2020. 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.

Savage, L.

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