



The A/B testing problem with Gaussian priors [☆]

Eduardo M. Azevedo ^{a,*}, David Mao ^b, José Luis Montiel Olea ^c,
Amilcar Velez ^d

^a *The Wharton School, University of Pennsylvania, United States of America*

^b *Department of Economics, University of Pennsylvania, United States of America*

^c *Department of Economics, Cornell University, United States of America*

^d *Department of Economics, Northwestern University, United States of America*

Received 1 June 2021; final version received 24 March 2023; accepted 25 March 2023

Available online 3 April 2023

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 after observing the experiment's outcome. We show that when the distribution for idea quality is Gaussian and there are linear costs of experimentation, there are exact formulae for the firm's optimal implementation decisions, the value of obtaining more data, and optimal experiment sizes. Our formulae—which assume that companies use randomized experiments to help them maximize expected profits—provide a simple alternative to i) the standard rules-of-thumb of power calculations for determining the sample size of an experiment, and also to ii) ad hoc thresholds based on statistical significance to interpret the outcome of an experiment.

© 2023 Published by Elsevier Inc.

JEL classification: C00; C44; C9

Keywords: Statistical Decision Theory; Optimal learning; Experiment design; A/B testing

[☆] We are grateful to Modibo Camara, Navin Kartik, Andrea Prat, and Charles Manski for very helpful comments and suggestions. All errors are our own.

* Corresponding author.

E-mail addresses: eazevedo@wharton.upenn.edu (E.M. Azevedo), davidmao@sas.upenn.edu (D. Mao), jl067@cornell.edu (J.L. Montiel Olea), amilcare@u.northwestern.edu (A. Velez).

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.

We propose practical tools and formulae for determining the sample size of experiments that are used to screen innovations, and offer concrete recommendations to decide which product innovations are worthy of being adopted based on the outcome of a randomized experiment. Our formulae—which assume that companies use randomized experiments to help them maximize expected profits—provide a simple alternative to i) the standard rules-of-thumb of power calculations for determining the sample size of an experiment, and also to ii) ad hoc thresholds based on statistical significance to interpret the outcome of an experiment.

We build upon the “A/B testing problem” proposed in Azevedo et al. (2020). In their model, 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.¹ Relative to this previous work, our key contribution is to solve the A/B testing problem for the specific case in which the distribution of idea quality is Gaussian. For clarity, we also focus exclusively on the case of a single idea and experimental cost linear in the size of the experiment. This allows us to provide exact results about the optimal implementation and experimentation strategies.²

The A/B testing problem with Gaussian priors considered in this paper has already been studied in a prescient early literature in statistical decision theory. Although the terminology and the context are different, Raiffa and Schlaifer (1961) and other researchers at the time were interested in the optimal use of scarce experimental resources long before this became a commonplace problem in the internet industry. The book by Raiffa and Schlaifer (1961), for example, provides a comprehensive 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”. Here we generalize some of the known results in this classical statistical decision theory literature and explain their significance in the empirically relevant context of experimentation in technology companies.

We present four main results. First, there is a simple closed-form solution to the firm’s optimal *implementation strategy*—that is, the firm’s decision of whether to adopt a product innovation after observing the outcome of a randomized experiment (Proposition 1). According to our result, the firm should calculate the usual t -statistic based on the estimated quality obtained from the experiment, and implement the idea only if the t -statistic is above a threshold. The threshold depends on the parameters of the Gaussian prior and the experimental noise, and can be positive or negative.³ A practical takeaway is that a profit maximizing firm might find it optimal to im-

¹ 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.

² Azevedo et al. (2020) focus on approximate results without parametric assumptions.

³ A negative threshold arises when the prior mean of the idea quality distribution is positive. In this case, the firm’s optimal estimate of the true idea quality can be positive even when the signal obtained from the experiment is negative. This happens because the firm’s optimal estimator—the posterior mean of idea quality—is a convex combination between the signal and the prior.

plement some ideas for which there is no evidence of a statistically significant positive effect on expected profits.

Second, there is a closed-form solution to the value obtained from an experiment (Proposition 2). The closed-form solution can be used to show that the value obtained from an experiment as a function of its sample size (assuming the firm uses the optimal implementation strategy) is nonnegative, bounded from above, strictly increasing, and, if the prior mean is different from zero, the value is first convex and then concave. The latter property has two qualitative implications for the determination of an experiment's sample size. First, when the prior mean is different from zero, small randomized experiments have a limited scope (as the marginal value obtained from a small experiment will be close to zero).⁴ Second, the marginal value of large experiments is eventually decreasing and close to zero, which means that very large randomized experiments are unlikely to be optimal. The formula for the value obtained from an experiment was known in the case where the firm uses an optimal implementation strategy in Raiffa and Schlaifer (1961). We generalize the result to the case of an arbitrary threshold implementation strategy and we document the differences vis-a-vis the optimal production function.

Third, there is a simple characterization of the experiment's optimal sample size (Proposition 3). A firm that maximizes expected profits can find the optimal sample size for an experiment by equating its marginal value to its marginal cost. Because in the A/B testing problem the value obtained from an experiment is first convex and then concave (as a function of the sample size), there will be typically two solutions satisfying the first-order conditions for optimality. The optimal sample size corresponds to the larger of these solutions. It is straightforward to write a simple computer algorithm that solves the first-order conditions and then selects that largest solution. Thus, our results give practical and easy-to-implement alternatives to power calculations for sample size determination.

Fourth, we derive comparative statics of the firm's optimal expected profits and the optimal sample size (Proposition 4). One qualitatively interesting result in our comparative statics is that the relation between the size of an experiment and the variance of the prior is not monotone. We show that, when the prior mean is (in absolute value) smaller than the prior standard deviation, a higher prior variance can lead to smaller or larger experiments.⁵

By construction, the formulae for the implementation strategy and the experimentation strategy are optimal if and only if the distribution of idea quality is indeed Gaussian. As a robustness check, we derive similar results for the case of *regret* minimization.⁶ Our results provide similar practical alternatives to using power calculations for determining an experiment's sample size and to basing implementation decisions on statistical significance, although there are some important qualitative differences relative to model with a Gaussian prior. For instance, it is optimal to implement a product innovation whenever its *t*-statistic is positive. Also, in this adversarial setting it is never optimal not to experiment. Our results generalize results from the statistical decision theory literature in Bross (1950) and Somerville (1954). One advantage of the regret minimization approach is that it can be applied to a setting with little data from prior experiments, where estimating a reasonable prior is difficult.

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

⁴ We note, however, that the interval where the value of an experiment is convex becomes smaller as the prior mean gets closer to zero.

⁵ In Section 3, we provide more details explaining the intuition behind this result.

⁶ In Section 4 we explain what regret is, and motivate its use.

calculations for sample size determination.⁷ 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).

The rest of the paper is organized as follows. Section 2 presents the model, section 3 the results, section 4 the minimax regret case, and section 5 a numerical example. Section 6 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 $\hat{\Delta}$, where the mean is equal to the (unknown) true quality, and the variance σ^2/n is known. 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.⁸ 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. 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 expectation in the display above is taken jointly over the signal realization and the idea's unknown quality. The firm's goal is to choose the experimentation and implementation strategy to maximize its payoff.

2.2. Related literature

This problem is a particular case of the A/B testing problem from Azevedo et al. (2020). The key restrictions are the Gaussian prior, single idea, and linear cost. The restrictions let us focus on additional insights from the Gaussian case. This is also 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.

Some results below are known from prior work. We include known results for clarity and credit prior work in detail along with each result. Notably, the formulas for the production function are known for the case of an optimal implementation strategy. Propositions (2) and (6) generalize the known formulas to arbitrary threshold implementation strategies.

⁷ See for example List et al. (2011); Athey and Imbens (2017).

⁸ We take n to be a real number for simplicity, but in practice the size of an experiment has to be an integer. In almost any practical application n is sufficiently large so that rounding to an integer is not an issue.

2.3. Notation

We denote the following pieces of notation. First, we denote the realization of the random variable Δ as δ and the realization of $\hat{\Delta}$ as $\hat{\delta}$. Next, we denote the normalized mean as $m := M/s$ 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 $\theta_0 := 0$.

Finally, we denote the posterior mean of quality after an experiment of sample size n with result $\hat{\delta}$ as

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

3. Results

3.1. Optimal implementation strategy

A standard formula for Bayesian updating with a Gaussian prior is that the posterior mean is an average of the signal $\hat{\delta}$ and the prior mean M :

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

Moreover, it is optimal for the firm to implement an idea iff the posterior mean is positive. This implies that the optimal implementation strategy is the following threshold rule.

Proposition 1 (*Optimal implementation strategy*). *It is optimal to implement an idea iff the signal $\hat{\delta}$ 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 t_n^* . 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 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—which we define as the probability that a standard normal exceeds t_n^* —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 sufficiently larger than 0.

A practical takeaway is that, in the case of positive prior mean, it is optimal to implement ideas even if there is no statistically significant evidence that the idea works. For example, DellaVigna and Linos (2022) study experiments performed by nudge units in the US government. They find that the average nudge has a positive effect. Assuming a Gaussian prior, the optimal strategy is to implement ideas by default, and only stick to the status quo if an experiment reveals bad news. The optimal strategy is strikingly different from the standard practice to implement only ideas for which the experiment reveals good news with a t -statistic above 1.65.

Another practical takeaway is that the optimal implementation strategy can be quite aggressive, even in the case of a negative prior mean. For example, consider the case where the absolute

value of the prior mean is smaller than s . Then, as long as the variance of experimental noise is smaller than the variance of the prior, the optimal threshold t -statistic is below one. This is considerably more aggressive than the standard practice of 1.65.

3.2. Optimal experimentation strategy

The optimal experimentation strategy depends on the value of allocating n users to the experiment. Azevedo et al. (2020) term this the production function. We extend their original definition to allow for cases in which the firm uses a decision rule that implements an idea whenever the t -statistic is above an arbitrary threshold τ . That is, we define the production function $f_\tau(n)$ as

$$f_\tau(n) := \mathbb{E} \left[\Delta \cdot \mathbf{1}_{\{\hat{\Delta} \geq \tau \cdot \sigma / \sqrt{n}\}} \right] - M^+.$$

Denote the optimal production function as $f(n) := f_{t_n^*}(n)$, where t_n^* is the optimal implementation threshold in Proposition 1.

Under a Gaussian prior, there is a closed-form solution for the production function. Our formula generalizes known results on the production function for an optimal implementation strategy to an arbitrary threshold strategy.⁹

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

$$f_\tau(n) = \begin{cases} s \left\{ \sqrt{\theta_n} \phi \left(m \sqrt{\theta_n} - \tau \sqrt{1 - \theta_n} \right) + m \Phi \left(m \sqrt{\theta_n} - \tau \sqrt{1 - \theta_n} \right) \right\}, & \text{if } m < 0 \\ s \left\{ \sqrt{\theta_n} \phi \left(m \sqrt{\theta_n} - \tau \sqrt{1 - \theta_n} \right) - m \Phi \left(-m \sqrt{\theta_n} + \tau \sqrt{1 - \theta_n} \right) \right\}, & \text{if } m \geq 0 \end{cases} \quad (1)$$

for $n > 0$. In addition, the optimal production function is

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

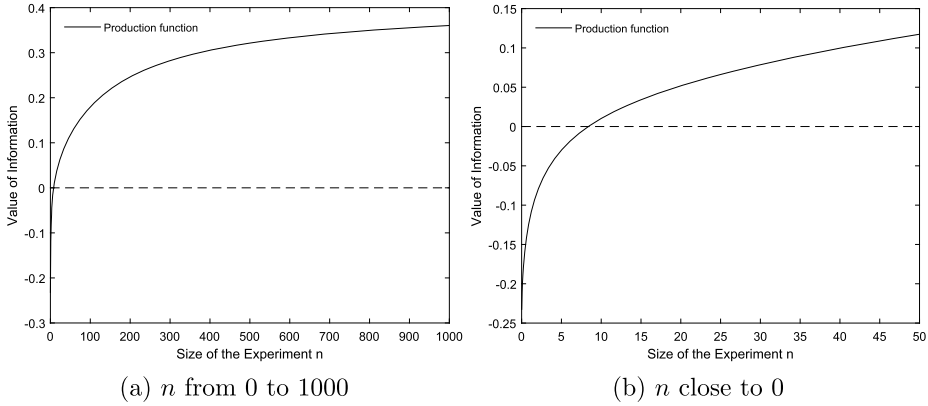
Both production functions are bounded and increasing. In addition, when $m \neq 0$, the production functions, $f_\tau(n)$ and $f(n)$, and their marginal product, $f'_\tau(n)$ and $f'(n)$, satisfy:

1. $\lim_{n \rightarrow 0} f_\tau(n) = -s|m|\Phi(|m|\tau/m) < 0$ and $\lim_{n \rightarrow 0} f(n) = 0$,
2. $\lim_{n \rightarrow 0} f'_\tau(n) = +\infty$ and $\lim_{n \rightarrow 0} f'(n) = 0$,
3. $\lim_{n \rightarrow \infty} f'_\tau(n) = 0$ and $\lim_{n \rightarrow \infty} f'(n) = 0$

Furthermore, there exists a threshold \hat{n} such that the optimal production function is convex on the interval $[0, \hat{n}]$ and concave on the interval $[\hat{n}, \infty]$.

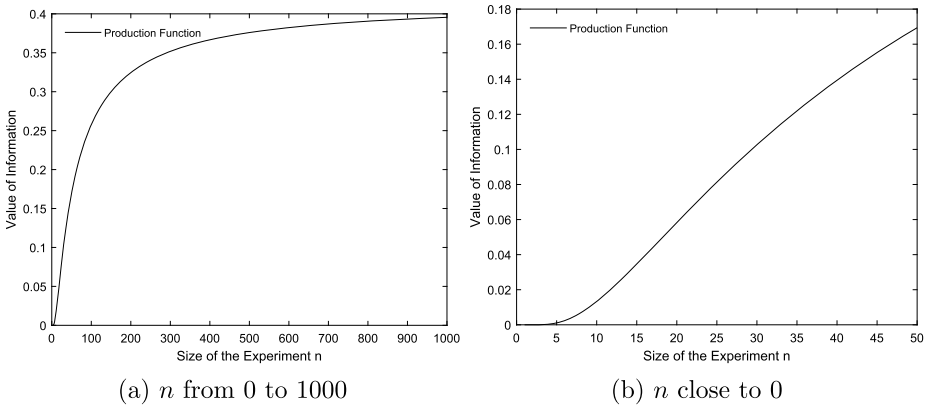
Fig. 1 and 2 present examples of production functions for a given threshold $\tau = 1.65$ and optimal threshold $\tau = t_n^*$. The production function $f_{1.65}$ is negative and concave near zero, while the optimal production function f is positive and convex near zero (as long as $M \neq 0$). For large values of n , both production functions are eventually concave.

⁹ This formula is known in the literature for the case where τ is the optimal threshold t_n^* . The earliest reference we could find is Grundy et al. (1956) equation (4), under a slightly different setup. The formula also appears in Raiffa and Schlaifer (1961). The formula is related but different from other formulas in the value of information literature. 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.



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

Fig. 1. The production function $f_{1.65}(n)$.



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

Fig. 2. The optimal production function $f(n)$.

For a given production function $f_\tau(n)$ with a given threshold τ , the optimal experimentation strategy is simple. When the cost of obtaining information always exceeds the value of information—i.e., $c \cdot n > f_\tau(n)$ for all $n > 0$ —it is optimal not to experiment at all and set $n_\tau^* = 0$. When the value of information exceeds the cost of obtaining information for some n , the optimal sample size equates the marginal cost and marginal product. The same principle holds for the optimal experimentation strategy n^* for the optimal production function $f(n) = f_{\tau_n^*}^*(n)$.

Proposition 3 (Optimal experimentation strategy). Consider a firm that implements any idea for which the t -statistic is above a threshold τ . Assume that $f_\tau(n) > c \cdot n$ for some $n > 0$. The optimal experimentation strategy n_τ^* solves the first order condition

$$f'_\tau(n) = c.$$

In addition, if $f(n) > c \cdot n$ for some $n > 0$, the optimal experimentation strategy n^* solves the first order condition

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

Proposition 3 gives a practical method to determine the optimal sample size, especially in the case where data on previous experiments is available. Data on previous experiments can be used to estimate the parameters of the Gaussian prior. Proposition 2 then gives a formula for the marginal product.¹⁰ The optimal sample size can then be determined as a function of the marginal cost of data. This is similar to the empirical Bayes method used in Azevedo et al. (2020). However, the Gaussian case is simpler for two reasons. First, estimation of the parameters is simpler because it only depends on estimating the mean and variance of the idea distribution. Second, the formula for the production function is simpler. In particular, the formula can be implemented using standard spreadsheet software, much like a power calculation.

Our practical method differs from the standard rule of thumb of power calculations. 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 optimal production function $f(n)$ and its optimal experimental strategy 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 $M < 0$

¹⁰ See Equation (4) in Appendix for an exact formula.

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

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

The practical takeaways of the comparative statics are qualitative guidelines for experimental design. (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. This is 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 finding in 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.¹¹ This effect may dominate, which is why the sign of the comparative statics depends on the parameters as described in the proposition.

4. The minimax regret case

4.1. Model

Our analysis in the previous sections relied on the choice of a specific prior distribution for idea quality. This means that our formulae for the implementation strategy and the experimentation strategy are optimal if and only if the distribution of idea quality is indeed Gaussian.

One possible alternative to maximizing expected profits for a fixed prior is to optimize the worst-case expected profit that could be attained under a large class of possible priors. For instance, one could consider all possible Gaussian priors with parameters (M, s) . Unfortunately, the results of this exercise turn out to be quite anticlimactic. If M is a very large negative constant, Proposition 1 implies that an extremely large value of $\hat{\delta}$ is needed to implement an idea. This implies that the production function is close to zero for any n , which makes the value of experimentation to be close to zero.

The pessimistic nature of minimax approaches is a well-understood phenomenon. Savage (1951) has suggested the use of *regret* as a target criterion (instead of utility or profits). In our problem, when the quality of the innovation is δ , the firm's expected payoff from the strategy (n, S) is

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

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, $\hat{\Delta}$. Define the *regret* of the strategy (n, S) as the difference between the optimal expected payoff if δ were observable, minus the expected payoff from choosing (n, S) . *Regret* thus becomes

$$\mathcal{R}(n, S, \delta) = \delta \cdot 1_{\{\delta > 0\}} - u(n, S; \delta). \tag{2}$$

Consider the strategy (n^*, S^*) that solves the problem

$$\inf_{n \geq 0, S} \sup_{\delta \in \mathbb{R}} \mathcal{R}(n, S, \delta). \tag{3}$$

The strategies that minimize the worst-case regret are usually referred to as minimax regret strategies. These strategies have been proposed by Savage (1951) as a guideline for making decisions under uncertainty. Minimax regret is one of the most prominent alternatives in the statistics literature on how to make implementation and experiment design decisions (see Manski (2019)). In what follows, we write $(n_{\text{MMR}}^*, S_{\text{MMR}}^*)$ instead of (n^*, S^*) to emphasize that the strategies that minimize the maximum regret are minimax regret strategies.

4.2. Results

It is well known that in Gaussian experiments the minimax regret implementation decision is to implement the idea if and only if the signal is positive.¹² That is,

¹¹ The simplest intuition comes from the large n approximation to the production function from Azevedo et al. (2020) Theorem 1. The marginal product for large n is approximately proportional to $(\sigma/s)^2$, which is decreasing in s .

¹² See Proposition 1 part (iii) in Stoye (2012).

Proposition 5 (*Minimax regret optimal implementation strategy*). *It is optimal to implement an idea iff the signal δ is positive.*

This strategy is what Manski (2004) calls the empirical success rule. A practical takeaway for this rule is that, from a minimax regret perspective, there is no reason to have a status quo bias. As a concrete example, consider a firm choosing between a standard marketing email and a new marketing email. The firm has little prior information and decides instead to use the minimax regret criterion. If switching to the new marketing email is costless, then it makes sense to switch as long as the experimental point estimate is positive. This is in contrast to the standard practice, for example in internet firms and in the nudge units studied in DellaVigna and Linos (2022), of implementing only ideas with statistically significant positive results, even when implementation costs are trivial.

For a fixed sample size, Stoye (2009) characterizes the minimax regret strategies for a problem similar to ours. His results include a comparison between the minimax regret strategies and the empirical success rule. In what follows we conduct a complementary analysis. For a given implementation strategy, we study the sample size determination problem using the minimax regret criterion. We restrict our analysis to strategies motivated by decisions based on t -statistics.

Let S_τ be a decision rule with a threshold t -statistic of τ . Define the cost function $\tilde{f}_\tau(n)$ as the firm’s maximum regret with threshold strategy S_τ and sample size n :

$$\tilde{f}_\tau(n) = \sup_{\delta \in \mathbb{R}} \mathcal{R}(n, S_\tau, \delta).$$

Let $\tilde{f}(n) := \tilde{f}_0(n)$ be the optimal cost function with the optimal implementation strategy from Proposition 5.

The next proposition shows that the minimax regret production function has a simple formula.

Proposition 6 (*Minimax regret cost function*). *The cost function is*

$$\tilde{f}_\tau(n) = \frac{\sigma}{\sqrt{n}} b(|\tau|) + cn,$$

for $n > 0$, and $\tilde{f}_\tau(0) = 0$. The function b is defined by

$$b(\tau) = \max_x x \Phi(\tau - x).$$

Bross (1950) and Somerville (1954) studied the minimax regret strategy for a problem closely related to ours. The production function for the optimal threshold $\tau = 0$ is derived in their work.

As in the Gaussian case, the closed-form solution to the cost function gives a simple method to determine the optimal experimentation strategy:

Proposition 7 (*Minimax regret optimal experimentation strategy*). *Consider a firm with a decision rule with threshold t -statistic of τ and minimax regret objective. The optimal experimentation strategy solves the first order condition*

$$\tilde{f}'_\tau(n) = c.$$

That is,

$$n_\tau^{mmr} = \left\{ \frac{\sigma b(|\tau|)}{2c} \right\}^{2/3}.$$

The optimal experiment size is strictly positive.

These propositions offer a practical alternative to power calculations for choosing an experimentation strategy. Practical implementation is even simpler than the Gaussian case because it requires no information on the prior. The optimal strategy depends only on the variable cost of experimentation and on the experimental noise. So this is a practical method even in cases where there is too little data on prior experiments to estimate a prior. Moreover, the formulae can be implemented easily, much like a traditional power calculation.

5. Illustrative numerical example

We now consider an illustrative numerical example. The example shows that, in plausible settings, our optimal experimentation strategy can perform considerably better than the standard rule-of-thumbs of power calculations and statistical significance. We based our setting on typical experiments run by brick-and-mortar firms with relatively small sample sizes.¹³

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 $s = 5$.¹⁴ 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 30.¹⁵

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.¹⁶ 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 Fig. 2.

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 110 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

¹³ See Pierce et al. (2021) for an example.

¹⁴ In practice the prior distribution might have been estimated from data on previous experiments. See Azevedo et al. (2019).

¹⁵ In our model, the parameters M , s and σ are known. In practice, the parameters have to be estimated. The most direct approach is the empirical Bayes method in Azevedo et al. (2019). They make a frequentist point estimate of the parameters, and use it for counterfactuals. This approach works well in their setting of a large online product with extensive data on previous experiments, where the parameters can be estimated with some precision. In settings with less data, uncertainty about the parameters may be a first-order concern. In this case, one possibility is to also place a prior on σ^2 and consider a model where the available data are the estimated effect and its standard error.

¹⁶ Consequently, c is between .1/1000 and 5/1000.

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^* & n^*	0%	0.158%	0.226%	0.327%
$t_{1-\alpha}$: 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%	-0.090%	0.266%
profits under t_n^* & n_{pc}^*	-4.058%	-0.497%	-0.052%	0.304%
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}^*	-0.220%	0.115%	0.195%	0.305%

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

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 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.

6. Conclusion

We study the A/B testing problem in the case where the prior distribution of idea quality is Gaussian. We propose practical tools and formulae for determining the sample size of experiments that are used to screen innovations, and offer concrete recommendations to decide which product innovations should be adopted based on the outcome of a randomized experiment. There is a simple closed-form solution to the firm’s optimal implementation strategy (Proposition 1). There is also a closed-form solution to the value of information obtained from an experiment (Proposition 2). There is a simple recommendation for deciding the size of a randomized experiment based on equating the marginal value of the experiment to its marginal cost (Proposition 3). We also derive qualitative principles for experimentation from comparative statics (Proposition 4). As a robustness check, we provide similar results for the case of expected regret minimization for an adversarial prior.

We compare the optimal strategy under the Gaussian prior, 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 rules-of-thumb. The results suggest that, in a setting that fits the assumptions well, the Bayesian optimal strategy and the minimax regret strategy can improve practical A/B testing.

Data availability

No data was used for the research described in the article.

Appendix A. Proofs

A.1. Proof of Proposition 1

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

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

A.2. Proof of Proposition 2

Proof. Denote $\tau_n := \tau \cdot \sigma / \sqrt{n}$. The innovation is implemented if and only if the signal exceeds τ_n . Then, the production function equals the expected value of the innovation times the probability it is implemented. Therefore,

$$\begin{aligned} f_\tau(n) &= \int \delta \cdot P(\hat{\Delta} \geq \tau_n | \Delta = \delta) \cdot g(\delta) d\delta - M^+ \\ &= \int \delta \cdot \Phi\left(\frac{\delta - \tau_n}{\sigma/\sqrt{n}}\right) \cdot g(\delta) d\delta - M^+ \\ &= \int (\delta - M) \cdot \Phi\left(\frac{\delta - \tau_n}{\sigma/\sqrt{n}}\right) \cdot g(\delta) d\delta + M \int \Phi\left(\frac{\delta - \delta_n^*}{\sigma/\sqrt{n}}\right) \cdot g(\delta) d\delta - M^+ \end{aligned}$$

The first term can be simplified to $s\sqrt{\theta_n} \cdot \phi((m - \tau_n/s) \sqrt{\theta_n})$ using integration by parts and setting:

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

The second term can be simplified to $M \cdot \Phi((m - \tau_n/s) \sqrt{\theta_n})$ using the following identity for the Gaussian distribution:

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

Note that $(m - \tau_n/s)\sqrt{\theta_n} = m\sqrt{\theta_n} - \tau\sqrt{1 - \theta_n}$. Then, combining gives Equation (1). The optimal production function follows by substituting $\tau = \tau_n^* = -m\sigma/(s\sqrt{n})$.

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

$$f'_\tau(n) = \frac{s^3}{2\sigma^2} \cdot \frac{(1 - \theta_n)^2}{\sqrt{\theta_n}} \phi\left(m\sqrt{\theta_n} - \tau\sqrt{1 - \theta_n}\right) \left(1 + \left(m\sqrt{1 - \theta_n} + \tau\sqrt{\theta_n}\right)^2\right), \tag{4}$$

a function that is positive for all n since $\theta_n = s^2/(s^2 + \sigma^2/n) \in (0, 1)$. Similarly, the derivative of the optimal production function is:

$$f'(n) = \frac{1}{2n^2} \cdot \frac{\sqrt{\theta_n}}{s} \cdot \phi\left(\frac{m}{\sqrt{\theta_n}}\right) \cdot \sigma^2 \cdot \theta_n = \frac{s^3}{2\sigma^2} \cdot \frac{(1 - \theta_n)^2}{\sqrt{\theta_n}} \cdot \phi\left(\frac{m}{\sqrt{\theta_n}}\right), \tag{5}$$

a function that is positive for all n . Therefore, the optimal production function is also strictly increasing.

These production functions are bounded because they are strictly increasing, and as $n \rightarrow \infty$, $\theta_n \rightarrow 1$, $f_\tau(n) \rightarrow s(\phi(m) - |m|\Phi(-|m|))$, 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 concave near the origin and in the interval $[n_2, +\infty)$. Differentiating Equation (1) twice shows that:

$$f''_\tau(n) = -f'_\tau(n) \cdot \left(\frac{1 + 3\theta_n}{2} + (m\sqrt{\theta_n} - \tau\sqrt{1 - \theta_n})x_n(\theta_n(1 - \theta_n))^{1/2} \left(\frac{2 + x_n^2}{1 + x_n^2}\right)\right),$$

where $x_n = (m\sqrt{1 - \theta_n} + \tau\sqrt{\theta_n})$. From Equation (4), $f'_\tau(n) > 0$. Then, the sign of the second derivative is the opposite of the sign of the following expression

$$\frac{1 + 3\theta_n}{2} + (m\sqrt{\theta_n} - \tau\sqrt{1 - \theta_n})x_n(\theta_n(1 - \theta_n))^{1/2} \left(\frac{2 + x_n^2}{1 + x_n^2} \right),$$

which is positive for $n \rightarrow 0$ and $n \rightarrow \infty$, since in those cases $\theta_n \rightarrow 0$ and $\theta_n \rightarrow 1$, respectively. Thus, we can conclude that $f''_\tau(n)$ is negative for values of n near zero and in an interval $[n_2, \infty)$ for some $n_2 > 0$.

For the optimal production function, we can conclude from an analysis of the second derivative that the production function is convex, then concave. That is

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

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:

$$n_1 := \frac{\sigma^2 \left((M^2 - s^2) + \sqrt{M^4 + 14M^2s^2 + s^4} \right)}{8s^4}. \tag{6}$$

The production function is convex over $[0, n_1)$, and concave over (n_1, ∞)

Finally, by taking the limits of Equations (4) and (5), we find that $\lim_{n \rightarrow 0} f'_\tau(n) = +\infty$, $\lim_{n \rightarrow 0} f'(n) = 0$, $\lim_{n \rightarrow \infty} f'_\tau(n) = 0$, and $\lim_{n \rightarrow \infty} f'(n) = 0$. \square

A.3. Proof of Proposition 3

Proof. The optimal experimentation strategy n^* maximizes $f_\tau(n) - cn$ for $n \geq 0$. The first order condition is $f'_\tau(n) = c$. Note that, in this case, the threshold t-statistic τ does not depend on n .

There must be at least one critical point that satisfies the first order condition, since (1) $\lim_{n \rightarrow 0} f_\tau(n) - c \cdot n \leq 0$, since $\lim_{n \rightarrow 0} f_\tau(n)$ is equal to $s \cdot m \cdot \Phi(-\tau)$ if $m < 0$ and $-s \cdot m \cdot \Phi(\tau)$ if $m \geq 0$ from Proposition 2; (2) $f_\tau(n) - c \cdot n > 0$ for some $n > 0$ by assumption; and (3) $f_\tau(n) - c \cdot n < 0$ for large n , since $f_\tau(n)$ is bounded, as shown in Proposition 2, and $c \cdot n$ is not. Further, from (1) and (2), the solution n^*_τ cannot be a boundary solution and must be a critical point that satisfies the first order condition.

The previous explanation applies for the optimal production function $f(n)$ as well: (1) is true since $\lim_{n \rightarrow 0} f(n) = 0$, and (2) and (3) follow from the same reasoning. The optimal experimentation strategy n^* must be a critical point that satisfies the first order condition. In this particular case, we can provide additional analysis and claim that there are either one or two solutions to the first-order condition, and n^* is the largest solution.

From Proposition 2 and its proof, 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 $(0, n_1)$ and decreasing over (n_1, ∞) , where n_1 is defined in Equation (6). Therefore, 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 n_1 and one larger. In this case, the sign analysis of $f''(n)$ from the proof of 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 $n_1 = 0$ 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.

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

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

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:

$$f(n, M, s, \sigma) = s \cdot f(n, M/s, 1, \sigma/s). \tag{9}$$

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

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

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

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

where

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

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 $\tilde{\theta}_1 < \tilde{\theta}_2$,

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

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

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

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

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 $|M| < s$. \square

A.5. Proof of Proposition 5

This proof uses the game-theoretic approach described in Section 2.2 of Stoye (2009) to solve the minimax regret problem. Define the firm’s average regret with respect to prior distribution G as

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

Let \mathcal{G} be the set of integrable distributions in \mathbb{R} . Consider the strategy (n^*, S^*) that solves the problem

$$\inf_{n \geq 0, S} \sup_{G \in \mathcal{G}} r(n, S, G). \tag{8}$$

Algebra shows that the strategy (n^*, S^*) that solves (8) is equivalent to the strategy that solves

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

To verify that $\hat{S} = 1\{\hat{\delta} > 0\}$ is an optimal implementation strategy under minimax regret for a given sample size $n > 0$, we consider the following three steps. First, we find a least favorable prior distribution \hat{G} for the firm’s decision rule, $\hat{G} \in \arg \sup_{G \in \mathcal{G}} r(n, \hat{S}, G)$. Second, we find the optimal (Bayesian) implementation strategy $S_{\hat{G}}$ for the least favorable prior \hat{G} , which implies that $r(n, S, \hat{G}) \geq r(n, S_{\hat{G}}, \hat{G})$ for any S . Third, we verify that $r(n, \hat{S}, \hat{G}) = r(n, S_{\hat{G}}, \hat{G})$.

Then, we conclude that \hat{S} is a minimax regret strategy since

$$\sup_{G \in \mathcal{G}} r(n, S, G) \geq r(n, S, \hat{G}) \geq r(n, S_{\hat{G}}, \hat{G}) = r(n, \hat{S}, \hat{G}) = \sup_{G \in \mathcal{G}} r(n, \hat{S}, G),$$

where the first inequality follows by definition, the second inequality by second step, the third equality by third step, and the last equality by the first step.

In what follows we verify the three steps mentioned above.

Step 1: Consider the prior \hat{G} with probability mass function

$$\hat{g}(\delta) = \begin{cases} \frac{1}{2} & \text{if } \delta = -k \cdot \frac{\sigma}{\sqrt{n}} \\ \frac{1}{2} & \text{if } \delta = k \cdot \frac{\sigma}{\sqrt{n}} \\ 0 & \text{otherwise,} \end{cases}$$

where k is the unique value that maximizes $\max_{x>0} x \cdot \Phi(-x)$. Under this prior, the firm’s average regret $r(n, \hat{S}, \hat{G})$ is equal to $(\sigma/\sqrt{n}) \cdot k \cdot \Phi(-k) + cn$.

Now, we verify that \hat{G} is a least favorable prior for the decision rule \hat{S} . Note that for $\delta > 0$, the firm’s regret is

$$\mathcal{R}(n, \hat{S}, \delta) = \delta \cdot \Phi\left(\frac{-\delta}{\sigma/\sqrt{n}}\right) + cn \leq \frac{\sigma}{\sqrt{n}} \max_{x>0} x \cdot \Phi(-x) + cn.$$

And for $\delta \leq 0$, the firm's regret is

$$\mathcal{R}(n, \hat{S}, \delta) = -\delta \cdot \Phi\left(\frac{\delta}{\sigma/\sqrt{n}}\right) + cn \leq \frac{\sigma}{\sqrt{n}} \max_{x>0} x \cdot \Phi(-x) + cn.$$

Thus, we conclude that

$$r(n, \hat{S}, G) = \int \mathcal{R}(n, S, \delta) dG(\delta) \leq \frac{\sigma}{\sqrt{n}} \max_{x>0} x \cdot \Phi(-x) + cn = r(n, \hat{S}, \hat{G}).$$

Step 2: Before computing the optimal (Bayesian) implementation strategy, note that it is equivalent to find an implementation strategy that maximize firm's expected profits,

$$S_{\hat{G}} \in \arg \max_S r(n, S, G) \iff S_{\hat{G}} \in \arg \max_S E[\Delta S] - cn.$$

The firm implements the idea if and only if the posterior mean quality is positive. When the prior is \hat{G} , using equation (A.2) in Azevedo et al. (2020), the firm implements the idea if and only if

$$\begin{aligned} & \int_{-\infty}^{\infty} \delta \cdot \phi(\hat{\delta} | \delta, \sigma^2/n) \cdot \hat{g}(\delta) d\delta \geq 0 \\ \iff & \frac{1}{2} k \frac{\sigma}{\sqrt{n}} \left\{ \phi\left(\hat{\delta} \mid k \frac{\sigma}{\sqrt{n}}, \frac{\sigma^2}{n}\right) - \phi\left(\hat{\delta} \mid -k \frac{\sigma}{\sqrt{n}}, \frac{\sigma^2}{n}\right) \right\} \geq 0. \end{aligned}$$

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

$$\begin{aligned} & \phi\left(\hat{\delta} \mid k \frac{\sigma}{\sqrt{n}}, \frac{\sigma^2}{n}\right) \geq \phi\left(\hat{\delta} \mid -k \frac{\sigma}{\sqrt{n}}, \frac{\sigma^2}{n}\right) \\ \iff & -\frac{(\hat{\delta} - k\sigma/\sqrt{n})^2}{2\sigma^2/n} \geq -\frac{(\hat{\delta} + k\sigma/\sqrt{n})^2}{2\sigma^2/n} \\ \iff & k\hat{\delta} \geq -k\hat{\delta}. \end{aligned}$$

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

$$S_{\hat{G}} = 1\{\hat{\delta} > 0\}.$$

Step 3: Since $S_{\hat{G}} = \hat{S}$, it follows that $r(n, \hat{S}, \hat{G}) = r(n, S_{\hat{G}}, \hat{G})$.

This completes the verification of the three steps in our proof.

A.6. Proof of Proposition 6

First, note that for $\delta > 0$, the firm's regret is

$$\mathcal{R}(n, S_{\tau}, \delta) = \delta \cdot \Phi\left(\tau - \frac{\delta}{\sigma/\sqrt{n}}\right) + cn \leq \frac{\sigma}{\sqrt{n}} \max_{x>0} x \cdot \Phi(\tau - x) + cn.$$

And for $\delta \leq 0$, the firm's regret is

$$\mathcal{R}(n, S_{\tau}, \delta) = -\delta \cdot \Phi\left(-\tau - \frac{-\delta}{\sigma/\sqrt{n}}\right) + cn \leq \frac{\sigma}{\sqrt{n}} \max_{x>0} x \cdot \Phi(-\tau - x) + cn.$$

This implies that $\mathcal{R}(n, S_\tau, \delta) \leq (\sigma/\sqrt{n})b(\tau) + cn$ if $\delta > 0$ and $\mathcal{R}(n, S_\tau, \delta) \leq (\sigma/\sqrt{n})b(-\tau) + cn$ if $\delta \leq 0$, where

$$b(\tau) = \max_x x\Phi(\tau - x).$$

Then, combining these two expressions, we conclude $\mathcal{R}(n, S_\tau, \delta) \leq (\sigma/\sqrt{n})b(|\tau|) + cn$ since $b(|\tau|) = \max\{b(\tau), b(-\tau)\}$. Finally, note that for any τ , there exists $x(\tau) \in \arg \max_x x \cdot \Phi(\tau - x)$. Define $\delta^*(\tau) = (\sigma/\sqrt{n})x^*(\tau)$ where $x^*(\tau) = x(\tau)$ if $b(\tau) > b(-\tau)$ and $x^*(\tau) = x(-\tau)$ if $b(\tau) \geq b(-\tau)$. Since $\mathcal{R}(n, S_\tau, \delta^*(\tau)) = (\sigma/\sqrt{n})b(|\tau|) + cn$, we conclude that the firm's cost function is equal to the maximum regret: $\tilde{f}'_\tau(n) = \mathcal{R}(n, S_\tau, \delta^*(\tau)) = (\sigma/\sqrt{n})b(|\tau|) + cn$.

A.7. Proof of Proposition 7

By Proposition 6, the cost function is a strictly convex function since $b(|\tau|) > 0$ for any τ . Then, the first order conditions characterize the minimax regret optimal experimentation strategy of the firm. Since

$$\tilde{f}'_\tau(n) = -\frac{\sigma}{2}n^{-3/2}b(|\tau|) + c.$$

It follows that $n^{mmr} = (\sigma b(|\tau|)/(2c))^{2/3}$.

References

Athey, S., Imbens, G.W., 2017. The Econometrics of Randomized Experiments. Handbook of Economic Field Experiments, vol. 1. Elsevier, pp. 73–140.

Azevedo, E.M., Alex, D., Montiel Olea, J., Rao, J.M., Weyl, E.G., 2020. A/B testing with fat tails. J. Polit. Econ. 128 (12), 4614–4672.

Azevedo, E.M., Deng, A., Montiel Olea, J.L., Weyl, E.G., 2019. Empirical Bayes estimation of treatment effects with many A/B tests: an overview. In: AEA Papers and Proceedings, vol. 109, pp. 43–47.

Bross, I., 1950. Two-choice selection. J. Am. Stat. Assoc. 45 (252), 530–540.

DellaVigna, S., Linos, E., 2022. RCTs to scale: comprehensive evidence from two nudge units. Econometrica 90 (1), 81–116.

Grundy, P., Healy, M., Rees, D., 1956. Economic choice of the amount of experimentation. J. R. Stat. Soc., Ser. B, Methodol. 18 (1), 32–49.

Keppo, J., Moscarini, G., Smith, L., 2008. The demand for information: more heat than light. J. Econ. Theory 138 (1), 21–50.

List, J.A., Sadoff, S., Wagner, M., 2011. So you want to run an experiment, now what? Some simple rules of thumb for optimal experimental design. Exp. Econ. 14 (4), 439.

Manski, C.F., 2004. Statistical treatment rules for heterogeneous populations. Econometrica 72 (4), 1221–1246.

Manski, C.F., 2019. Treatment choice with trial data: statistical decision theory should supplant hypothesis testing. Am. Stat. 73 (sup1), 296–304.

Manski, C.F., Tetenov, A., 2016. Sufficient trial size to inform clinical practice. Proc. Natl. Acad. Sci. 113 (38), 10518–10523.

Manski, C.F., Tetenov, A., 2019. Trial size for near-optimal choice between surveillance and aggressive treatment: reconsidering MSLT-II. Am. Stat. 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. J. Health Econ. 20 (1), 109–129.

Moscarini, G., Smith, L., 2002. The law of large demand for information. Econometrica 70 (6), 2351–2366.

Pierce, L., Rees-Jones, A., Blank, C., 2021. The negative consequences of loss-framed performance incentives. Technical report. National Bureau of Economic Research.

Raiffa, H., Schlaifer, R., 1961. Applied statistical decision theory.

Savage, L., 1951. The theory of statistical decision. J. Am. Stat. Assoc. 89, 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. *J. Econom.* 151 (1), 70–81.

Stoye, J., 2012. Minimax regret treatment choice with covariates or with limited validity of experiments. *J. Econom.* 166 (1), 138–156.