A Theory of Experimenters*

Abhijit Banerjee
Massachusetts Institute of
of Technology and NBER
banerjee@mit.edu
economics.mit.edu/faculty/banerjee/

Sylvain Chassang
New York University
chassang@nyu.edu
princeton.edu/∼chassang/

Sergio Montero
California Institute
of Technology
smontero@caltech.edu
hss.caltech.edu/∼smontero/

Erik Snowberg
California Institute
of Technology and NBER
snowberg@caltech.edu
hss.caltech.edu/∼snowberg/

July 19, 2016

PRELIMINARY

Abstract

This paper proposes a decision-theoretic framework for experiment design. We model experimenters as ambiguity averse decision-makers, who trade-off a subjective expected utility term with a robustness term that can be interpreted as the welfare of an adversarial audience entertaining a variety of possible priors. This framework suitably accounts for experimenters' preferences for randomization, and the circumstances in which they occur. It also provides useful guidance regarding open practical questions, such as that of rerandomization. We show that rerandomization creates a trade-off between subjective balance and robustness but that the costs of rerandomization grow very slowly with the amount of rerandomization. This lets us propose a simple and principled rule of thumb for rerandomization.

KEYWORDS: experiment design, robustness, ambiguity aversion, rerandomization

*We are grateful to Angus Deaton, Guido Imbens, Charles Manski, Alexey Tetenov, as well as audience members at the Cowles Econometric Conference (2016), the North-American Econometric Society Meeting (2016), as well as the UCL/Cemmap workshop on Econometrics for Public Policy (2016), for several helpful discussions. Chassang and Snowberg gratefully acknowledge the support of NSF grant SES-1156154.

1 Introduction

As the use of experiments spreads in academia, business, and public policy, there has been a growing need to clarify best-practices and ensure the reliability of experimental findings. The experimental community has responded through the introduction of registries, and spirited discussions of practices such as pre-analysis plans, rerandomization, and statistical techniques (Bruhn and McKenzie, 2009; Deaton, 2010; Duflo et al., 2008; Humphreys et al., 2013; Imbens, 2010; Olken, 2015). Quite surprisingly, however, there is no comprehensive decision theoretic framework for optimal experiment design to guide these efforts. This paper seeks to provide such a framework.

A framework for optimal experiment design requires a model of experimenters and their objectives. While models of information acquisition feature prominently in modern microeconomic theory (Rothschild, 1974; Grossman and Stiglitz, 1980; Aghion et al., 1991; Bergemann and Välimäki, 1996; Persico, 2000; Bergemann and Välimäki, 2002, 2006), they fail to predict a key feature of the way scientists learn: by running randomized controlled trials (RCTs; see Kasy, 2013). The reason for this is that much of applied microeconomic theory models decision-makers using subjective expected utility (Savage, 1954). Mixed strategies are never strictly optimal for such a decision-maker. Since RCTs are mixed strategies over experimental assignments, they can never be strictly optimal.

As experimenters often incur significant expense to randomize their experimental allocations, any useful decision-theoretic framework for experiment design must first account for such preferences. We propose to replace subjective expected utility with ambiguity averse preferences, specifically minmax expected utility as axiomatized by Gilboa and Schmeidler (1989). In our model, an ambiguity averse decision-maker must make a binary policy choice $a \in \{0, 1\}$ affecting a population of individuals with characteristics $x \in X \subset \mathbb{R}^m$ and conditionally independent outcomes. To improve the quality of her decision-making, the decision-maker runs an experiment that assigns a given number N of participants to either

treatment or control. Each experimental participant obtains an outcome $y \in \{0, 1\}$ observed by the decision-maker. The decision-maker's final policy choice depends on the experimental assignment and outcomes.

Under an innocuous assumption, the decision-maker can be thought of as maximizing the weighted average of a Bayesian subjective expected utility term, and an adversarial maxmin-expected-utility term capturing the robustness of decision rules. Under this decomposition the decision-maker can be interpreted as maximizing her payoffs under her own subjective view of the world, while also trying to satisfy an adversarial audience with non-common priors, and veto power. The relative weights that the decision-maker places on the subjective and maxmin termsyield informative comparative static exercises.

The paper reports two main sets of results. First, we show that RCTs can be optimal for an ambiguity-averse decision-maker, and clarify the circumstances in which it is the case. If the decision-maker places non-zero weight on satisfying her adversarial audience, then, for sufficiently large sample sizes, it is always strictly optimal for the decision-maker to use a RCT. The value of randomness is that it allows audience members with non-common priors to agree on an ex ante course of action. RCTs permit robust, prior-free inference, and achieve assignment losses of order $1/\sqrt{N}$. Inversely, deterministic experiments are generically strictly optimal when the sample size is small, or when the decision-maker puts sufficiently high weight on her subjective expected utility.

Our model fits the observed heterogeneity in experimental practive well. Randomized experiments tend to be used by decision-makers who put a high value on convincing an adversarial audience (scientists, pharmaceutical companies), or when the decision-maker can afford large samples (A/B testing in online marketing). Whenever data points are expensive and the decision-maker puts little weight on satisfying an adversarial audience (private firms testing new products in select markets, politicians testing platforms in specific states, etc. . .), optimal experiments are deterministic, and finely optimize the subjective decision-making value of each acquired data point.

Our second set of results exploits our framework to shed light on an important practical question: rerandomization (Morgan and Rubin, 2012). Rerandomization consists of drawing multiple treatment assignments, and choosing one that maximizes the balance between groups on some covariates. For example, a medical researcher may want to ensure that treatment and control groups are similar in terms of gender, age, race, and baseline health variables such as blood pressure and weight. Despite the practical ease of using rerandomization to ensure balance, researchers have voiced the concern that it may affect the reliability of findings (Bruhn and McKenzie, 2009).

We show that the trade-offs at the heart of rerandomization are well captured in our framework. Successive rerandomizations improve balance, as captured by the subjective expected utility component of preferences. However, rerandomization reduces robustness, as captured by the adversarial component of preferences. In the extreme case where the allocation is rerandomized until perfect balance is achieved, the allocation is effectively deterministic and the adversarial term remains bounded away from first best.

Importantly, our framework also lets us assess the costs of rerandomization. We show that it only affects the robustness of decision making when the number of rerandomizations is exponential in the sample size. This lets us propose a rule of thumb for rerandomization that markedly improves balance, while keeping losses in robustness small.

The paper is structured as follows. Section 2 introduces our framework. Section 3 delineates the forces that determine whether running a randomized or deterministic experiment is optimal. Section 4 studies the trade-offs involved in rerandomization. Numerical evaluations of the efficiency of various experimental designs are presented in Section 5. Section 6 concludes. Proofs are contained in Appendix A unless mentiond otherwise.

2 A Framework for Optimal Experiment Design

Decisions and payoffs. A decision-maker chooses whether or not to implement a policy that provides a treatment $\tau \in \{0,1\}$ to a large finite population of individuals. Potential outcomes for a subject with treatment status $\tau \in \{0,1\}$ are random variables $Y^\tau \in \{0,1\}$; $Y = 1$ is referred to as a success. Each individual is associated with covariates $x \in X \subset \mathbb{R}^m$, where X is finite. Covariates $x \in X$ are observable and affect the distribution of outcomes Y. The distribution $q \in \Delta(X)$ of covariates in the population is known and has full support. Outcomes Y are i.i.d. conditional on covariates. The probability of success given covariate x is denoted by $p_x^\tau \equiv \text{prob}(Y^\tau = 1 | x = x)$.[1]

The state of the world is described by the finite-dimensional vector p of success probabilities conditional on covariates, $p = (p_x^0, p_x^1)_{x \in X} \in [0,1]^{2X} \equiv P$. Note that state space P is compact, convex, and finite-dimensional. Given a state p and a policy decision $a \in \{0,1\}$, the decision-maker's payoff $u(p,a)$ is

$$u(p,a) \equiv \mathbb{E}_p Y^a = \sum_{x \in X} q(x) p_x^a.$$

Experiments and strategies. To maximize her odds of making the correct policy choice, the decision-maker can run an experiment on N participants. For simplicity, we assume that N is even, and exogenously given. Formally, an experiment is a tuple $e = (x_i, \tau_i)_{i \in \{1,\dots,N\}} \in (X \times \{0,1\})^N \equiv E$. Experiment e generates outcome data $y = (y_i)_{i \in \{1,\dots,N\}} \in \{0,1\}^N \equiv \mathcal{Y}$, with y_is independent realizations of $Y_i^{\tau_i}$ given (x_i, τ_i).

The decision-maker's strategy consists of both an experimental design $\mathcal{E} \in \Delta(E)$, which is a mixed strategy over experimental assignments, and an allocation rule $\alpha : E \times \mathcal{Y} \to \Delta(\{0,1\})$, which maps experimental data (e,y) to policy decisions $a \in \{0,1\}$.[2] We denote

[1] Our framework encompasses unobservable differences z, as well. Then $p_x^\tau = \int p_{x,z}^\tau dF(z|x)$. As they do not impact the analysis, we omit them in what follows.

[2] Targeting treatment to particular sub-populations is possible as well: if the targeting is not too fine—for example, by gender, race, or economic class—then the analysis within each sub-population would be the same as the analysis here. For more on targeting by covariates see Kitagawa and Tetenov (2015), and references therein.

by $\mathcal{A}$ the set of such mappings. A standard RCT, assigning a share $\pi \in (0,1)$ of participants to treatment $\tau = 1$, corresponds to a strategy $(\mathcal{E}_{\text{RCT}}, \alpha_{\text{RCT}})$:

- $\mathcal{E}_{\text{RCT}}$ samples N exchangeable participants labelled by $i \in \{1, \cdots, N\}$, with covariates $(x_i)_{i\in\{1,\cdots,N\}}$ drawn according to q;
- $\mathcal{E}_{\text{RCT}}$ assigns treatment $\tau_i = \mathbf{1}_{i\leq \pi N}$;
- $\alpha_{\text{RCT}}(e, y) \equiv \mathbf{1}_{\widehat{y}^1 \geq \widehat{y}^0}$, where $\widehat{y}^\tau \equiv \sum_{i=1}^N y_i \mathbf{1}_{\tau_i=\tau} \big/ \sum_{i=1}^N \mathbf{1}_{\tau_i=\tau}$.

Preferences. The decision-maker is ambiguity averse with standard maxmin preferences (Gilboa and Schmeidler, 1989). She chooses a strategy $(\mathcal{E}, \alpha)$ that solves

$$\max_{\substack{\alpha\in\mathcal{A} \\ \mathcal{E}\in\Delta(E)}} U(\mathcal{E}, \alpha), \quad \text{where} \quad U(\mathcal{E}, \alpha) \equiv \min_{h\in H'} \mathbb{E}_{h,\mathcal{E}}[u(p, \alpha(e, y))] \tag{1}$$

and H' is a convex set of priors $h \in \Delta(P)$ over states $p \in P$. This can be thought of as a zero-sum game in which nature picks distribution $h \in H'$ after the decision-maker picks a strategy $(\mathcal{E}, \alpha)$. Randomizations in mixed strategies are independent of moves by nature.

We use the usual statistical distance $d(h, h') \equiv \sup_{\substack{A\subset P \\ A \text{ meas.}}} |h(A) - h'(A)|$ on distributions whenever making genericity statements. Almost-sure statements are made with respect to the Lebesgue measure on P.

Equivalent experiments. As we are assuming that participants are exchangeable conditional on covariates, experiments that differ only by a permutation of participants with identical covariates are equivalent from a decision-making perspective. It is useful to formalize this point in the context of maxmin preferences.

Definition (equivalent experiments)**.** *Two experiments* $e = (x_i, \tau_i)_{i\in\{1,\ldots,N\}}$ *and* $e' = (x'_i, \tau'_i)_{i\in\{1,\ldots,N\}}$ *are* equivalent*, denoted by* $e \sim e'$*, if there exists a permutation* $\sigma : \{1, \ldots, N\} \to \{1, \ldots, N\}$ *of the participants' labels such that* $(x_i, \tau_i) = (x'_{\sigma(i)}, \tau'_{\sigma(i)})$ *for all* i*. The equivalence class of an experiment* e *is denoted by* $[e]$*.*[3] *We denote by* $[E]$ *the partition of possible experiments in*

[3]It is convenient to include distributions $\mathcal{E}$ with support in $[e]$ in the equivalence class of e.

equivalence classes. We say that two experimental designs $\mathcal{E}$ and $\mathcal{E}'$ are equivalent, denoted by $\mathcal{E} \sim \mathcal{E}'$, if they induce the same distribution over $[E]$

Lemma 1. *Whenever $\mathcal{E} \sim \mathcal{E}'$, $\max_{\alpha \in \mathcal{A}} U(\mathcal{E}, \alpha) = \max_{\alpha \in \mathcal{A}} U(\mathcal{E}', \alpha)$.*

Thus, equivalent experiments guarantee the decision-maker the same utility.

2.1 Key assumptions

We place two additional assumptions on the model of Section 2. The first is innocuous, and allows the decision-maker's objective to be expressed as a weighted average of a Bayesian subjective expected utility term and a maxmin expected utility term.

The second is more substantial: it ensures that the set of possible priors entertained by the decision-maker is rich enough that for any given experimental assignment, there exists a prior under which this assignment does not permit efficient decision-making. Note that the order of quantifiers is important: given a realized experimental assignment, we can find such a prior h. Indeed, we show in Proposition 3 that randomized experiments yield approximately efficient decisions for all priors.

2.1.1 Decomposition of Maxmin Preferences

The following assumption leads to a useful decomposition of the decision-maker's preferences.

Assumption 1 (absolute continuity)**.** *There exist h_0 and $\lambda \in (0, 1)$ such that for every prior $h \in H'$ and almost every state $p \in P$,*

$$h(p) \geq \lambda h_0(p). \tag{2}$$

Absolute continuity requirement (2) implies that every prior $\widehat{h} \in H'$ can be written as $\widehat{h} = \lambda h_0 + (1 - \lambda)h$, where $h \in H \equiv \frac{1}{1-\lambda}(H' - \lambda h_0)$. Condition (2) moreover implies that

elements $h \in H$ are themselves probability distributions over states $p \in P$, and the set H is also compact and convex.

Altogether, this implies that the decision-maker's objective (1) can be rewritten as

$$U(\mathcal{E}, \alpha) \equiv \lambda \mathbb{E}_{h_0,\mathcal{E}}[u(p, \alpha(e, y))] + (1 - \lambda) \min_{h \in H} \mathbb{E}_{h,\mathcal{E}}[u(p, \alpha(e, y))]. \tag{3}$$

Keeping h_0 and H fixed, parameter λ provides a convenient and continuous measure of the decision-maker's degree of ambiguity aversion. For $\lambda = 1$ this nests standard subjective expected utility maximization. We sometimes refer to the case of $\lambda = 1$ as Bayesian.

This yields a useful interpretation: The decision-maker wants to make a decision that is successful under her own subjective prior h_0, but also satisfies an audience of players with priors $h \in H$. Weights λ and $1 - \lambda$ represent the respective weights that the decision-maker places on her own subjective utility and that of her audience. In a traditional model of ambiguity aversion, the possible priors are all internal to the decision-maker; here we add an additional interpretation, that they come from an adversarial audience.[4]

2.1.2 Limited Extrapolation

Throughout we assume that $N \leq 2|X|$ so that, even though there are finitely many covariate profiles $x \in X$, assigning each of them to treatment and control is not feasible. This condition is assumed to hold even as we take N to be large.[5]

This allows us to impose the following *limited extrapolation* condition on X, N, and H. Denote by $p^a \equiv \sum_{x \in X} q(x) p^a_x$ the expected probability of success given policy $a \in \{0, 1\}$. Given an experiment $e = (\tau_i, x_i)_{i \in \{1,\cdots,N\}}$, denote by $\overline{p}_e \equiv (p^{\tau_i}_{x_i})_{i \in \{1,\cdots,N\}}$ the subset of success rates for participants in the experiment. Vector $\overline{p}_e$ is an upper bound to the information

[4]Note that if the audience has veto power and enjoys some outside option, then the weight ratio $\frac{1-\lambda}{\lambda}$ would be interpreted as the appropriate Lagrange multiplier on satisfying the audience's individual rationality constraint.

[5]A natural case is that in which $N = |X|$ and $q(x) = \frac{1}{|X|}$: individuals are unique, and there are as many configurations of characteristics as potential subjects.

generated by experiment e.

Assumption 2 (limited extrapolation)**.** *There exists $\xi > 0$ such that, for all $e \in E$, there exists a prior $h \in \arg\min_{h \in H} \mathbb{E}_h(\max_{a \in \{0,1\}} p^a)$ such that, for almost every $\overline{p}_e$,*

$$\min \left\{ \mathbb{E}_h \left[\max_{a \in \{0,1\}} p^a - p^0 \middle| \overline{p}_e \right], \mathbb{E}_h \left[\max_{a \in \{0,1\}} p^a - p^1 \middle| \overline{p}_e \right] \right\} > \xi.$$

Limited extrapolation implies that for any realized experimental assignment there exists a prior $h \in H$ under which that assignment does not allow for first-best decision-making. That is, conditional on the data generated by any experiment, their exists a prior such that the residual uncertainty about which policy maximizes population-level outcomes remains bounded away from 0.[6]

3 Optimal Design and Randomization

We now use our framework to characterize optimal experimental design.

3.1 Bayesian Experimentation

When $\lambda = 1$, the decision-maker is a standard subjective expected utility maximizer. In this case, it is well known that deterministic experiments are weakly optimal. We restate this result and show that for generically every prior (that is, for an open and dense set of priors under statistical distance d), deterministic experiments are strictly optimal when λ is close to one.

Proposition 2 (near-Bayesians do not randomize)**.** *If $\lambda = 1$, then for every prior h_0, there exists a deterministic experiment e^* solving* (3)*.*

[6] We could do away with this assumption if we considered a decision-maker solving for minmax regrets rather than minmax expected utility.

For generically every prior h_0, there exist $\underline{\lambda} \in (0,1)$ and a unique equivalence class of experiments $[e^*]$ such that for all $\lambda > \underline{\lambda}$, a (potentially mixed) experiment $\mathcal{E} \in \Delta(E)$ solves (3) *if and only if* $\text{supp}\ \mathcal{E} \subset [e^*]$.

Proof. The proof that deterministic experiments are always weakly optimal for a Bayesian decision-maker is instructive. The decision-maker's payoff from running experiment $\mathcal{E}$ can be written as

$$\max_{\alpha \in \mathcal{A}} \mathbb{E}_{h_0, \mathcal{E}} \left[u(p, \alpha(e, y)) \right] = \sum_{e \in E} \mathcal{E}(e) v(h_0, e).$$

where $v(h, e) \equiv \sum_{y \in \mathcal{Y}} \max_{a \in \{0,1\}} \mathbb{E}_{p \sim h} \left[\text{prob}(y|p, e) u(p, a) \right]$. Therefore, any deterministic experiment e^* solving $\max_{e \in E} v(h_0, e)$ is optimal. More strongly, $\mathcal{E}$ solves (3) if and only if $\text{supp}\ \mathcal{E} \subset \underset{e \in E}{\text{argmax}}\, v(h_0, e)$.

The proof that deterministic experiments are generically strictly optimal is provided in Appendix A. □

In recent work, Kasy (2013) uses a version of Proposition 2 when $\lambda = 1$ to argue that RCTs are suboptimal. We believe that rather than invalidating the use of RCTs, Proposition 2 highlights the limits of subjective expected utility maximization as a suitable positive model of experimenters. We argue instead that the adversarial framework of (3) is more successful in explaining the range of information acquisition strategies observed in practice.

3.2 Adversarial Experimentation

We now assume that the decision-maker puts a fixed positive weight on satisfying her audience, and study comparative statics as the sample size becomes large.

Proposition 3. *Take weight $\lambda \in (0,1)$ as given. There exists $\underline{N}$ such that for all $N \geq \underline{N}$, any optimal experiment is randomized. More precisely, the following hold:*

(i) For any N, any optimal experiment $\mathcal{E}^*$ satisfies

$$\max_{\alpha} \min_{h \in H} \mathbb{E}_{h,\mathcal{E}^*} \left[u(p, \alpha(e, y)) \right] \geq \min_{h \in H} \mathbb{E}_h \left(\max_{a \in \{0,1\}} u(p, a) \right) - \sqrt{\frac{1}{N}}. \tag{4}$$

(ii) For any N, all deterministic experiments $e \in E$ are bounded away from first-best:

$$\forall e \in E, \quad \max_{\alpha \in \mathcal{A}} \min_{h \in H} \mathbb{E}_{h,e} \left[u(p, \alpha(e, y)) \right] < \min_{h \in H} \mathbb{E}_h \left(\max_{a \in \{0,1\}} u(p, a) \right) - \xi. \tag{5}$$

The first part of the proposition shows that the efficiency loss of the optimal experiment compared to the first-best decision is bounded above by $1/\sqrt{N}$. The second part shows that the loss from a deterministic experiment is bounded below by ξ, where ξ is bounded away from zero, and independent of N. Thus, as N grows, the optimal experiment cannot be deterministic: it must be randomized. Intuitively, the decision-maker is playing a zero-sum game against nature. After the decision-maker picks an experiment, nature picks the prior which maximizes the chance of picking the wrong policy, given that experimental design. If there is any clear pattern in the decision-maker's assignment of treatment, nature will exploit this pattern to lower the decision-maker's payoff. Randomization eliminates patterns that nature can exploit.[7]

Figure 1 maps out implications of Propositions 2 and 3 for practical experiment design. Proposition 2 shows that when sample points are scarce, or when the decision-maker does not put much weight on satisfying anyone else (λ close to 1), optimal experimentation will be deterministic, driven by prior h_0. That is, the experimenter will focus on assigning treatment and control observations to the participants from whom she expects to learn the most. This is the case, for example, when a firm is implementing a costly new process in a handful of

[7]This point is related to Saito (2015) who emphasizes that ambiguity averse agents may have preferences for randomization even if they exhibit risk-aversion over known lotteries.

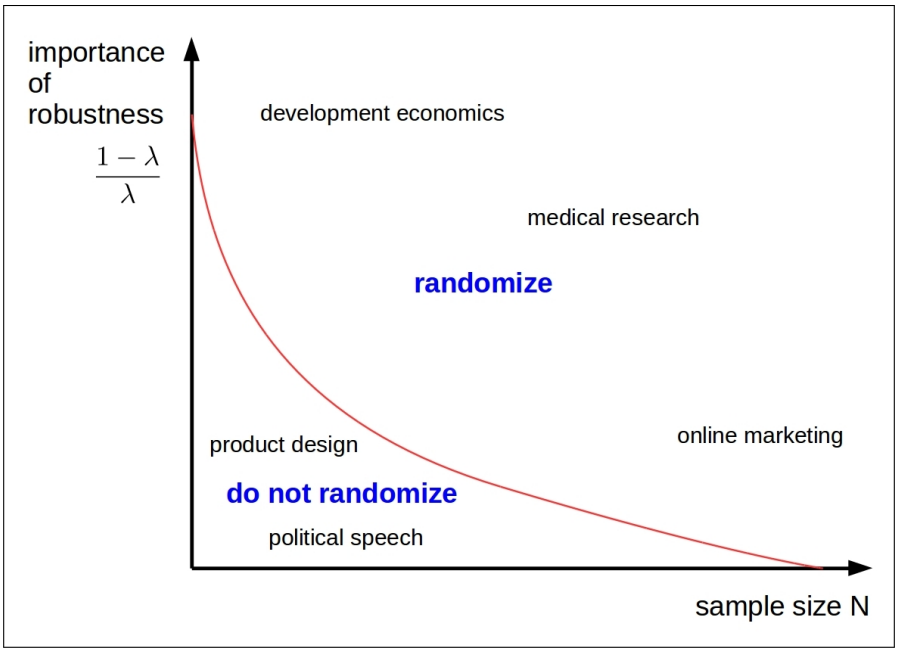


Figure 1: Maximizing subjective decision-making value, vs adversarial, prior-free inference.

production sites: The firm will focus on a few teams where it can learn the most. Similarly, a politician trying out platforms will do so at a few carefully chosen venues in front of carefully chosen audiences.

However, when the decision-maker must satisfy an adversarial audience, or has a sufficiently large sample, she will randomize as in Proposition 3. The former is the case in scientific research. The latter is the case for firms dealing online with many end users: Although the firm only needs to convince itself of the effectiveness of a particular ad or UI design, observations are so plentiful that randomization is used to effictively address internal concerns over robustness.

Proposition 3, in particular, has constructive applications to experimental practice. First, it implies that a decision-maker who randomizes even without understanding all its ramifications—why she is randomizing, what audience the experiment is meant to satisfy—will nevertheless produce an almost-optimal experiment for large values of N. Even if someone (or her own doubts) produces a particularly challenging prior, the decision rule is still likely to be close

to optimal. Further, this proposition highlights the importance of actually randomizing. An experimenter that adopts a protocol that is only "nearly" random, such as assignment based on time of day of an experimental session (see Green and Tusicisny, 2012, for a critique), or the first letter of an experimental subject's name (as was the case in the deworming study of Miguel and Kremer, 2004; see Deaton, 2010 for a critique), can always find a skeptical prior in its audience. Randomization provides a defense against the most skeptical priors, but near-randomization offers no such protection.

One interpretation of Proposition 3 is that randomized experiments let parties with different priors agree on a process. However this agreement is only possible at the ex ante stage before a specific experiment is realized. Once the experiment is realized, there will exist a prior under which this experiment is no longer informative. This corresponds to the fact that non-expected utility preferences are not dynamically consistent (Machina, 1989).

4 Rerandomization

Proposition 3 establishes that randomization is essential to guarantee successful prior-free performance. Since deterministic experiments are subjectively optimal, this must come at the cost of expected utility under the decision-maker's subjective prior h_0. More generally, randomization will come at the expense of the experimenter's subjective goals. One such prominent goal is balance, i.e. differences between the mean of covariates across treatment and control samples.[8]

A common tool to improve balance and other subjective performance criteria is the practice of rerandomization—redrawing an assignment until an acceptable subjective performance is achieved. However, concerns about the effects of rerandomization on the robustness of inference from RCTs lead many scholars to not report the fact that they have rerandomized (Bruhn and McKenzie, 2009). While scholars are often advised to use stratifi-

[8]In practice, unbalanced samples can lead to re-evaluation and criticism of a study's findings Banerjee et al. (forthcoming); Gerber and Green (2000); Imai (2005).

cation or matching to achieve balance instead, this is often quite difficult to implement when stratifying on multiple continuous covariates; thus, researchers turn to rerandomization. Our framework lets us formalize the robustness loss to rerandomization, and leads to a useful rule of thumb for experimental practice.

4.1 Balance and Rerandomization

Balance. The objective of the decision-maker, as described by (3), can be rewritten as

$$\max_{\mathcal{E}\in\Delta(E)} \lambda\mathbb{E}_{\mathcal{E}}[B(e)] + (1-\lambda)R(\mathcal{E}). \tag{6}$$

Given some assignment rule $\alpha(e,y)$, term $B(e) = \mathbb{E}_{h_0}[u(p,\alpha(e,y))]$ measures the informational value of the experimental assignment from the perspective of a Bayesian decision-maker with prior h_0.[9] Importantly, it depends on the realized experimental sample. Term $R(\mathcal{E}) \equiv \min_{h\in H} \mathbb{E}_h[u(p,\alpha(e,y))]$ captures robustness objectives.

In practice, decision-makers frequently express preferences for balanced samples, i.e., treatment and control samples with matching distributions of characteristics. This suggests that decision-makers frequently start from symmetric Laplacian priors. Formally, it correponds to an objective of the form

$$B(e) = \Gamma\left(\left|\left|\sum_{i|\tau_i=1} x^i - \sum_{i|\tau_i=0} x^i\right|\right|\right) \tag{7}$$

for some appropriate norm $||\cdot||$ in $\mathbb{R}^m$, and Γ a strictly increasing positive function.[10]

In the analysis that follows, we allow for any bounded balance function $B(e)$. Importantly, this function can come from any source: an experimenter's pre-analysis plan, an implementation partner... In fact, the objective function $B(e)$ may be chosen knowing the

[9] A natural assignment rule is $\alpha(e,y) \equiv \text{argmax}_{a\in\{0,1\}} \overline{y}^a - \overline{y}^{1-a}$.

[10] This class of objective functions includes the Mahalanobis distance commonly used in multivariate matching (Rubin, 1980; Cochrane and Rubin, 1973; Rubin, 1979).

charcteristics $(x_i)_{i\in\{1,\cdots,N\}}$ of participants.

Rerandomization. We define rerandomization as follows. Given an integer K, the $K-$rerandomized experiment $\mathcal{E}_K$ proceeds as follows:

1. For a fixed sample of xs drawn according to population distribution $q \in \Delta(X)$, independently draw a set of K assignments $\{e_1, \cdots, e_K\}$ with each $e_k = (x_i, \tau_{i,k})$ such that a fraction $\pi \in (0,1)$ of participants receives treatment $\tau = 1$;

2. Select the assignment $e_K^* \in \text{argmax}_{e\in\{e_1,\cdots,e_K\}} B(e)$ that maximizes balance function $B(e)$, breaking ties randomly;

3. Run the experiment e_K^*.

Rerandomization rules may also use a stopping time to endogenously pick the number of randomizations (Morgan and Rubin, 2012). Provided the stopping time has an upper bound K, all our results apply for this bound.[11]

Note that it follows from standard arguments based on exchangeability that the asymptotic standard error of the Wald estimator is unchanged provided the balance function B is symmetric.

4.2 A Tradeoff

It is immediate that $B(e_K^*)$ first-order stochastically dominates $B(e_{K-1}^*)$. By definition, rerandomization must mechanically improve balance. The question, therefore, is whether rerandomization can adversely affect robustness. We show that it can. For simplicity we consider the focal case where $N = |X|$ and the selection of characteristics $(x_i)_{i\in\{1,\cdots,N\}}$ is fixed.

[11] Our analysis also extends to resampling procedures for which both the set of participants and the treatment assignment are redrawn at each step. Interestingly the loss or robustness from resampling may exceed the loss of robustness from rerandomization by increasing the range of experimental designs that may be picked by a misguided balance function.

Proposition 4. *Consider rerandomized experiment $\mathcal{E}_K$. There exists $\underline{\rho} > 0$ such that for every N, if $K \geq 2^N$, then*

$$\max_{\alpha} \min_{h \in H} \mathbb{E}_{h,\mathcal{E}_K} [u(p, \alpha(e, y))] < \min_{h \in H} \mathbb{E}_h \left(\max_{a \in \{0,1\}} u(p, a) \right) - \underline{\rho}\xi.$$

Intuitively, when K is sufficiently large, the allocation is essentially deterministic, which by Proposition 3 precludes first-best robustness.[12]

Quantifying the cost of rerandomization. The number of rerandomizations K needed for Proposition 4 to apply is extremely large: it is exponential in the sample size. This suggests that the cost of rerandomization may be quite small for reasonable numbers of rerandomizations. Indeed, this is the case:

Proposition 5. *Given $K \geq 2$, consider a rerandomized experiment $\mathcal{E}_K$ assigning treatment to a proportion $\pi \in (0, 1)$ of participants. Then,*

$$\min_{h \in H} \mathbb{E}_{h,\mathcal{E}_K} [u(p, \alpha(e, y))] \geq \min_{h \in H} \mathbb{E}_h \left(\max_{a \in \{0,1\}} u(p, a) \right) - \sqrt{\frac{\max\{1, \ln(K)\}}{\kappa N}},$$

where $\kappa = \min\{\pi, 1 - \pi\}$.

Comparing this with Proposition 3, the additional loss from rerandomization, $\sqrt{\ln(K)}$, is rleatively small: between 1.5 and 3 for sample sizes between 10 and 10,000. However, the additional loss to robustness due to increasing K falls off relatively slowly. On the other hand, subjective balance is established very quickly: K rerandomizations guarantee that the final sample will be within the group of 5% most balanced samples with probability $1-0.95^K$, and the improvement in balance falls off relatively quickly. Observing that $1-0.95^{100} > 0.99$, we suggest the following simple rule of thumb for rerandomization:

[12]However, it doesn't affect inference for a Bayesian, as $h_0(p \,|\, e, y, e \sim \delta_e) = h_0(p \,|\, e, y, e \sim \mathcal{E}_{\text{RCT}}) = h_0(p \,|\, e, y, e \sim \mathcal{E}_K)$, where $e \sim \delta_e$ denotes that e is drawn deterministically, and $e \sim \mathcal{E}_K$ denotes e is drawn according to the K rerandomized experiment $\mathcal{E}_K$.

Rule of Thumb. *Set* $K = \min\{N, 100\}$.

5 Simulations

We explore our results through two sets of simulations. The first is closely connected to the theory above, and provides additional intuition. The second is designed to produce high error rates, and be closer to experimental practice. The message from both is clear, however, rerandomization increases balance with very little to no increase in mistaken decisions.

5.1 Simulation for Smooth Priors

The first simulation is closely tied to the theory developed above and studies the effect of rerandomization in a model where treatment effects depend smoothly on real-valued covariates.

Covariates x are drawn i.i.d. according to $\prod_{k=1}^{5} U[0,1]$, a five-dimensional uniform distribution. These are mapped to outcomes according to a five-dimensional unknown parameter μ according to $\text{prob}(Y_i = 1|x) = \frac{\exp(\mu \cdot x)}{1+\exp(\mu \cdot x)}$. The parameter μ is drawn according to a five-dimensional normal: $\mu \sim \prod_{k=1}^{5} \mathcal{N}(0,1)$. The final treatment decision α is selected according to the treatment status that maximizes the Bayesian posterior probability of success in the population.

Loss, compared to first best, can be examined in a number of ways, as shown in the first panel of Figure 2. These are:

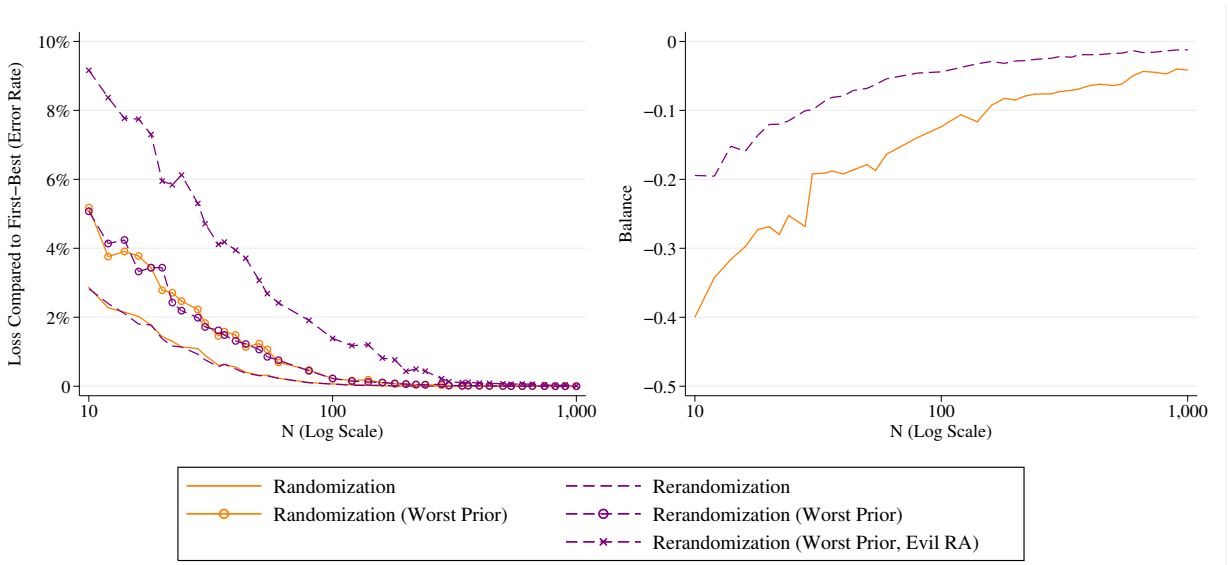
Bayes Loss given Bayes Optimal Assignment: $\mathbb{E}_{\mu,x,\tau^*}(max_{a\in\{0,1\}} u(p,a) - u(p,\alpha))$

Loss with worst prior, most balanced assignment: $\max_{\mu} \mathbb{E}_{x,\tau^*}(\max_{a\in\{0,1\}} u(p,a) - u(p,\alpha))$

Loss with worst prior, τ: $\max_{\mu} \mathbb{E}_{x} \max_{\tau} \mathbb{E}(\max_{a\in\{0,1\}} u(p,a) - u(p,\alpha))$

The first of these is the ex-ante expected loss, which is essentially identical under randomization and rerandomization. The second chooses the prior that maximizes the error rate

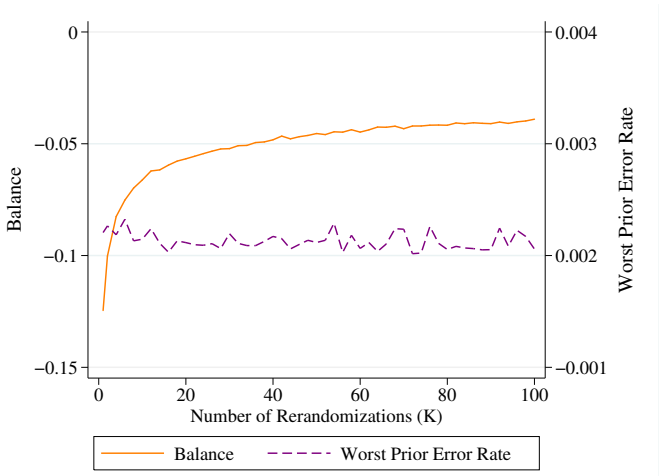
Figure 2: Rerandomization substantially increase balance with no cost to robustness.



given the experimental strategy $\mathcal{E}$ of the experimenter. While this is substantially higher, as expected, it is not substantially different between randomization and rerandomization. The final stacks the deck against the experimenter, and assumes that the experimenter has an "evil RA" who chooses the experimental assignment τ from e_K that maximizes the expected loss. This has no application in the case of randomization, but in the case of rerandomization it substantially increases error rates. However, it is important to note even under this highly unrealistic scenario—the evil RA must know the data-generating process—the error rate is about one-tenth of 1% for $N \geq 300$.

In the simulations above, we vary K, the number of rerandomizations according to our rule of thumb, $K = \min\{N, 100\}$. This suggests that the simulation may be masking some decision-making cost of rerandomization by increasing N simultaneously. Figure 3 shows this is not the case by plotting worst-prior loss and balance with K, holding N fixed at 100. Balance improves substantially, especially for the first 20 rerandomizations, but the error rate is essentially flat.

Figure 3: Rerandomizaton increases balance with no robustness cost with fixed N.



5.2 Simulation for Non-Smooth Priors

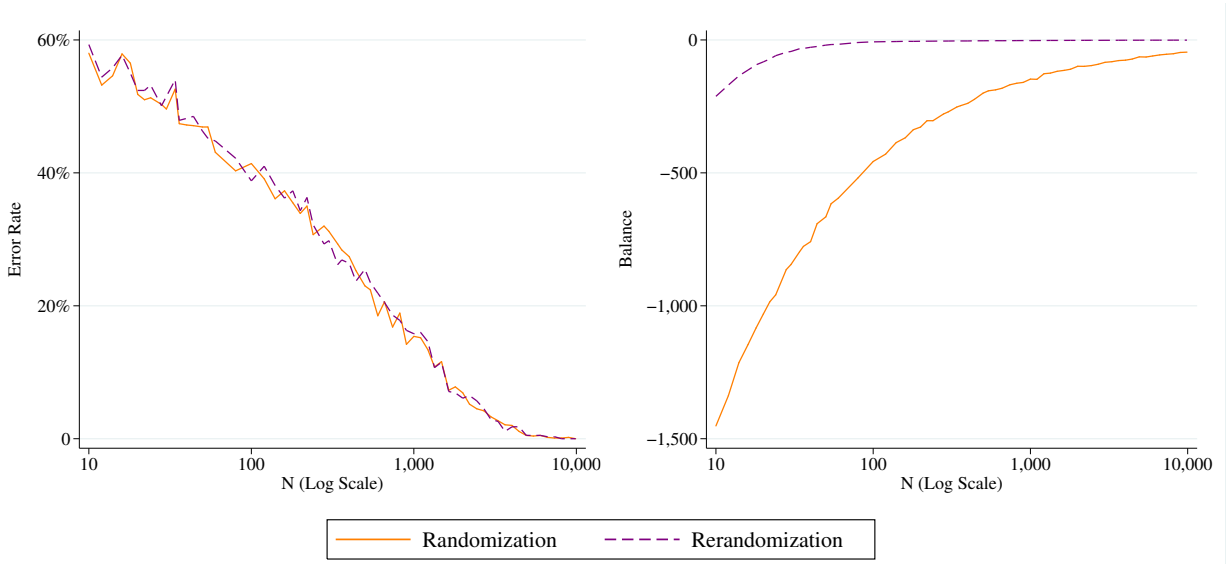
Generally, the assignment τ will be chosen according to the balance function in (7), or by someone without particular knowledge of the data-generating process. To provide a conservative numerical assessment of losses from rerandomization, we consider experimental design problems where nature gets to pick very non-smooth prior, and assignment is chosen according to $\alpha(e, y) \equiv \text{argmax}_{a \in \{0,1\}} \overline{y}^a - \overline{y}^{1-a}$.

Here, there a single covariate $x \in X = \{1, 2, \ldots, 10{,}000\}$. Even covariates are twice as likely as odd covariates, and the treatment effect is small and negative for even covariates, and large and positive for odd covariates. Specifically, for $n \in \{1, 2, \ldots, 5{,}000\}$,

$$q(2n-1) = \frac{q(2n)}{2} = \frac{2}{3|X|}, \quad p^1_{2n-1} = 4p^0_{2n-1} = \frac{4}{5}, \quad \text{and} \quad p^1_{2n} = \frac{p^0_{2n}}{2} = \frac{1}{4},$$

Thus, on aggregate, $u(p, 1) = \frac{13}{30} > \frac{2}{5} = u(p, 0)$, so treatment is beneficial, and $\alpha = 1$ is the "correct" decision. This setup is meant to make attempts to balance the sample likely to cause inferential mistakes—balancing will tend to pair odd observations with the more

Figure 4: In simulations, resampling and rerandomization substantially increase balance with no cost to robustness.



numerous even observations, which are not an appropriate comparison group.[13]

Figure 4 examines the error rates and balance according to (7) using the l_1 norm, of three different ways of selecting a sample: randomization, rerandomization, and resampling. As can be seen in the first panel, all three give roughly the same error rate. This is because the chosen balance function, $B(e)$, in these simulations is very unlikely to select a more biased sample allocation. While in any specific application the interaction of the model parameters and the balance function may produce different results, it appears quite difficult to find a balance function that 1) might actually be used and 2) is particularly pernicious.

On the other hand, once again, rerandomization substantially improves the balance of the samples. This is particularly true for small and moderate sample sizes, up to the order of 1,000, although even with 10,000 participants there is an improvement in balance, even though we only re-draw the experimental allocation τ 100 times.

[13]Indeed, using pairwise matching to assign treatment and control status increases inferential errors, but does so equally for randomization and rerandomization

6 Discussion

We have shown that a maxmin framework for experimental design produces reliable positive results, and can be used to inform practice. Elsewhere, we informally apply this framework to discuss other aspects of experimental design, including registration, pre-analysis plans, and external validity (Banerjee et al., forthcoming).

Before closing, we make a final comment that can tie together the findings from both sets of our results. In Section 3 we discuss the issues with quasi-random assignment, and discuss as an example the seminal work on field trials in development by Miguel and Kremer (2004), and the criticism of that work by Deaton (2010). At issue was the authors use of treatment assignment to every third village, alphabetically.[14] This was done to satisfy implementing partners, who were uncomfortable with randomization. Our proposal for rerandomization could provide a middle ground: specifically, Miguel and Kremer could have drawn up $K = N = 75$ lists of possible allocations, and allowed the implementing partner to pick which one they preferred. This would come at little cost of robustness to the trial, and may have satisfied the implementing partners need for a feeling of control. However, to be sure of this, one would need a theory of implementing partners, which is beyond the scope of this paper.

[14] All villages eventually received treatment, but in three waves. The first, fourth, seventh, etc. villages were assigned to Wave I, and the second, fifth, eighth, etc. to Wave II, and the third, sixth, ninth, etc. to Wave III.

A Proofs

Proof of Lemma 1: The decision-maker's indirect utility from running experiment $\mathcal{E}$, $V(\mathcal{E}) \equiv \max_{\alpha \in \mathcal{A}} U(\mathcal{E}, \alpha) = \max_{\alpha \in \mathcal{A}} \min_{h \in H'} \mathbb{E}_{h,\mathcal{E}}[u(p, \alpha(e, y))]$, can be viewed as the value of a zero-sum game where player 1, the decision-maker, chooses an allocation rule $\alpha \in \mathcal{A}$, and player 2, nature, chooses a prior $h \in H'$. By standard arguments, this game has a Nash equilibrium, implying that $V(\mathcal{E}) = \min_{h \in H'} \max_{\alpha \in \mathcal{A}} \mathbb{E}_{h,\mathcal{E}}[u(p, \alpha(e, y))]$.

Given h, the decision-maker's payoff from running experiment $\mathcal{E}$ can be written as

$$\begin{aligned}
\max_{\alpha \in \mathcal{A}} \mathbb{E}_{h,\mathcal{E}} \left[u(p, \alpha(e, y))\right] &= \max_{\alpha \in \mathcal{A}} \sum_{e \in E} \mathcal{E}(e) \mathbb{E}_{p \sim h} \left[\sum_{y \in \mathcal{Y}} \text{prob}(y|p, e) u(p, \alpha(e, y))\right] \\
&= \sum_{e \in E} \mathcal{E}(e) \sum_{y \in \mathcal{Y}} \max_{a \in \{0,1\}} \mathbb{E}_{p \sim h} \left[\text{prob}(y|p, e) u(p, a)\right] \\
&= \sum_{e \in E} \mathcal{E}(e) v(h, e),
\end{aligned}$$

where $v(h, e) \equiv \sum_{y \in \mathcal{Y}} \max_{a \in \{0,1\}} \mathbb{E}_{p \sim h} \left[\text{prob}(y|p, e) u(p, a)\right]$. Since $v(h, e) = v(h, e') \equiv v(h, [e])$ for all $e' \in [e]$, it follows that $V(\mathcal{E}) = \min_{h \in H'} \sum_{[e] \in [E]} \mathcal{E}([e]) v(h, [e])$. Thus, if $\mathcal{E}$ and $\mathcal{E}'$ induce the same distribution over $[E]$, $V(\mathcal{E}) = V(\mathcal{E}')$. ■

Proof of Proposition 2: We now show that $\text{argmax}_{[e] \in [E]} v(h_0, [e])$ is generically a singleton. We first show that the set of priors h_0 such that there is a uniquely optimal equivalence class of experiments is open. Suppose that $[e_0]$ is uniquely optimal under h_0. Since E is finite, there exists $\eta > 0$ such that $v(h_0, [e]) < v(h_0, [e_0]) - \eta$ for all $[e] \neq [e_0]$. Since $v(h, e)$ is continuous in h, this implies that there exists a neighborhood H_0 of h_0 such that, for all $h \in H_0$, $v(h, [e]) < v(h, [e_0]) - \eta/2$. Hence, $[e_0]$ is the uniquely optimal design for all priors $h \in H_0$.

We now prove that the set of priors h_0 such that there is a uniquely optimal equivalence class of experiments is dense. The proof is by induction on the number of equivalence classes

$[e_0] \in \operatorname{argmax}_{[e]\in[E]} v(h_0, [e])$. We show that if there exist n such equivalence classes, then in any neighborhood of h_0 there exists a prior h such that there are at most $n-1$ equivalence classes in $\operatorname{argmax}_{[e]\in[E]} v(h, [e])$.

Indeed, assume that $[e_0] \neq [e_1]$ both belong to $\operatorname{argmax}_{[e]\in[E]} v(h_0, [e])$. For $\theta > 0$, consider the polynomial $M_\theta(p)$ in $p \in P$ defined by

$$M_\theta(p) = v\left((1-\theta)h_0 + \theta p, [e_0]\right) - v\left((1-\theta)h_0 + \theta p, [e_1]\right),$$

where $(1-\theta)h_0 + \theta p$ denotes the mixture probability measure that places mass $1-\theta$ on h, and mass θ on the Dirac mass at p. Since $[E]$ is finite, for all $\theta > 0$ small enough, it must be that

$$\operatorname*{argmax}_{[e]\in[E]} v((1-\theta)h_0 + \theta p, [e]) \subset \operatorname*{argmax}_{[e]\in[E]} v(h_0, [e]).$$

Consider such a $\theta > 0$. The fact that $[e_0] \neq [e_1]$ implies that $M_\theta(p)$ is not identically equal to 0. Hence, there exists p such that $v\left((1-\theta)h_0 + \theta p, [e_0]\right) \neq v\left((1-\theta)h_0 + \theta p, [e_1]\right)$. This implies that the inductive step holds at prior $\tilde{h} = (1-\theta)h_0 + \theta p$. Using the fact that $[E]$ is finite and $v(h, [e])$ is continuous in h, this implies that the inductive step holds at a prior that admits a density against the Lebesgue measure. Thus, when $\lambda = 1$, deterministic experiments are generically strictly optimal.

Now, given any λ, h, and $[e]$, since the decision-maker's utility only takes values in $[0,1]$, letting $\alpha_0 \in \operatorname{argmax}_{\alpha\in\mathcal{A}} \mathbb{E}_{h_0,e}[u(p, \alpha(e,y))]$ we have

$$\begin{aligned} v(\lambda h_0 + (1-\lambda)h, [e]) &\leq \lambda v(h_0, [e]) + (1-\lambda)v(h, [e]) \leq v(h_0, [e]) + (1-\lambda) \quad \text{and} \\ v(\lambda h_0 + (1-\lambda)h, [e]) &\geq \lambda v(h_0, [e]) + (1-\lambda)\mathbb{E}_{h,e}[u(p, \alpha_0(e,y))] \geq v(h_0, [e]) - (1-\lambda). \end{aligned}$$

As there are finitely many experiments, if $[e_0]$ is the unique maximizer of $v(h_0, [e])$, there exists $\eta > 0$ such that, for all $[e] \neq [e_0]$, $v(h_0, [e_0]) > v(h_0, [e]) + \eta$. Together, this implies that there exists $\underline{\lambda} \in (0,1)$ such that, for all $\lambda > \underline{\lambda}$, objective (3) is maximized at $\mathcal{E}$ if and

only if supp $\mathcal{E} \subset [e_0]$. ■

Proof of Proposition 3: To establish point (i) we use the strategy $(\mathcal{E}_{\text{RCT}}, \alpha_{\text{RCT}})$ such that

- $\mathcal{E}_{\text{RCT}}$ consists of sampling N participants with covariates independently drawn according to q and assigning treatments $\tau_i = \mathbf{1}_{i \leq N/2}$;
- $\alpha_{\text{RCT}}(e, y) \equiv \mathbf{1}_{\widehat{y}^1 > \widehat{y}^0}$, where $\widehat{y}^\tau$ is the sample average of outcomes among participants with treatment status τ.

Losses $L(p)$ from first best, given state of the world p, are defined as

$$L(p) \equiv \max_{a \in \{0,1\}} p^a - \mathbb{E}_{p, \mathcal{E}_{\text{RCT}}} \left[p^{\mathbf{1}_{\overline{y}^1 - \overline{y}^0 > 0}} \right].$$

By symmetry, it suffices to treat the case where $p_1 - p_0 > 0$. In this case, we have $L(p) = (p_1 - p_0)\text{prob}_{p, \mathcal{E}_{\text{RCT}}}(\overline{y}^1 - \overline{y}^0 \leq 0)$. The probability of choosing the suboptimal policy can be bounded using McDiarmid's inequality.[1] By applying McDiarmid's inequality to $f(y) \equiv \frac{2}{N} \sum_{i=1}^{N/2} y^0_{i+N/2} - y^1_i$, we obtain

$$\begin{aligned} \text{prob}_{p, \mathcal{E}_{\text{RCT}}}(\overline{y}^1 - \overline{y}^0 \leq 0) &= \text{prob}_{p, \mathcal{E}_{\text{RCT}}} \left(\widehat{y}^0 - \widehat{y}^1 - (p^0 - p^1) \geq (p^1 - p^0) \right) \\ &\leq \exp\left(-(N/2)(p^1 - p^0)^2 \right). \end{aligned}$$

For any $a > 0$, $x \mapsto x \exp(-ax^2)$ is log-concave and maximized at $x = (2a)^{-1/2}$. This implies that

$$\max_{a \in \{0,1\}} p^a - \mathbb{E}_{p, \mathcal{E}_{\text{RCT}}} \left[p^{\mathbf{1}_{\overline{y}^1 - \overline{y}^0 > 0}} \right] \leq \frac{1}{\sqrt{N}}. \tag{8}$$

An analogous argument delivers (8) also for the case where $p^1 - p^0 \leq 0$. Hence, given any $h \in H$,

$$\mathbb{E}_h \left(\max_{a \in \{0,1\}} u(p, a) \right) - \mathbb{E}_{h, \mathcal{E}_{\text{RCT}}} \left[u(p, \alpha_{\text{RCT}}(e, y)) \right] \leq \frac{1}{\sqrt{N}}.$$

[1] McDiarmid's (1989) inequality can be stated as follows. Let $X_1, \ldots, X_n$ be independent random variables, with X_k taking values in a set A_k for each k. Suppose that the (measurable) function $f : \times_k A_k \to \mathbb{R}$ satisfies $|f(x) - f(x')| \leq c_k$ whenever x and x' differ only in the kth coordinate. Then, for any $t > 0$, $\text{prob}\left(f(X_1, \ldots, X_n) - \mathbb{E}[f(X_1, \ldots, X_n)] \geq t \right) \leq \exp\left(-2t^2 / \sum_k c_k^2 \right)$.

To establish point (ii), fix a deterministic experiment $e \in E$. From Assumption 2, there exists $h \in H$ such that for almost every $\overline{p}_e$,

$$\min\left\{\mathbb{E}_h\left[\max_{a\in\{0,1\}} p^a - p^0 \middle| \overline{p}_e\right], \mathbb{E}_h\left[\max_{a\in\{0,1\}} p^a - p^1 \middle| \overline{p}_e\right]\right\} > \xi. \text{ Hence,}$$

$$\begin{aligned}\max_{\alpha} \mathbb{E}_{h,e}\left[u(p,\alpha(e,y))\right] &\leq \mathbb{E}_{h,e}\left[\max_{a\in\{0,1\}} \mathbb{E}_{h,e}\left[u(p,a)|\overline{p}_e\right]\right] \\ &\leq \mathbb{E}_{h,e}\left[\max_{a\in\{0,1\}} u(p,a)\right] - \xi.\end{aligned}$$

■

Proof of Proposition 4: Fix any $e^\dagger \in \operatorname{argmax}_{e\in\text{supp }\mathcal{E}_K} B(e)$. The kth rerandomized trial, $k \in \{1,\ldots,K\}$, selects each experiment in its support with probability at least $r \equiv (\underline{q}\underline{\pi})^N$, where $\underline{q} \equiv \min_{x\in X} q(x) \leq 1/|X|$ and $\underline{\pi} \equiv \min\{\pi, 1-\pi\} \leq 1/2$. Therefore, the odds of rerandomization picking experiment $e^\dagger$ are at least $\rho \equiv 1-(1-r)^K$. For $K \geq (2|X|)^N$,

$$\rho = 1 - \exp(K\ln(1-r)) \sim 1 - \exp(-Kr) \geq 1 - 1/\exp > 0.$$

Hence, there exists $\underline{\rho} > 0$ such that, for all N, rerandomized experiment $\mathcal{E}_K$ selects deterministic experiment $e^\dagger$ with probability at least $\underline{\rho}$.

The proof of Proposition 3 implies that there exists $h^\dagger \in H$ such that

$$\begin{aligned}\forall e \in E,\ \max_{\alpha\in\mathcal{A}} \mathbb{E}_{h^\dagger,e}[u(p,\alpha(e,y))] &\leq \min_{h\in H} \mathbb{E}_h\left(\max_{a\in\{0,1\}} u(p,a)\right), \\ \text{and } \max_{\alpha\in\mathcal{A}} \mathbb{E}_{h^\dagger,e^\dagger}[u(p,\alpha(e^\dagger,y))] &\leq \min_{h\in H} \mathbb{E}_h\left(\max_{a\in\{0,1\}} u(p,a)\right) - \xi. \\ \text{Hence, } \max_{\alpha\in\mathcal{A}} \min_{h\in H} \mathbb{E}_{h,\mathcal{E}_K}[u(p,\alpha(e,y))] &\leq \min_{h\in H} \mathbb{E}_h\left(\max_{a\in\{0,1\}} u(p,a)\right) - \underline{\rho}\xi.\end{aligned}$$

■

Proof of Proposition 5: Denote by $(\overline{y}_{0,k}, \overline{y}_{1,k})$ the sample average of outcomes by treatment

group for experiment e_k, and by $(\overline{y}_0^*, \overline{y}_1^*)$ the sample average of outcomes by treatment group for the experimental design e_K^* selected by rerandomized experiment $\mathcal{E}_K$.

Losses $L(p)$ from first best given state of the world p are defined as $L(p) \equiv \max_{a\in\{0,1\}} p^a - \mathbb{E}_{p,\mathcal{E}_K}\left[p^{\mathbf{1}_{\overline{y}_1^* - \overline{y}_0^* > 0}}\right]$. By symmetry, it suffices to treat the case where $p^1 - p^0 > 0$. In this case, we have

$$\begin{aligned}
L(p) &= (p^1 - p^0)\text{prob}_{p,\mathcal{E}_K}(\overline{y}_1^* - \overline{y}_0^* \leq 0) \\
&\leq (p^1 - p^0)\text{prob}_{p,\mathcal{E}_K}\left(\min_{k\in\{1,\ldots,K\}} \overline{y}_{1,k} - \overline{y}_{0,k} \leq 0\right) \\
&\leq (p^1 - p^0)\min\left\{1, \sum_{k=1}^{K} \text{prob}_{p,\mathcal{E}_K}(\overline{y}_{1,k} - \overline{y}_{0,k} \leq 0)\right\} \\
&\leq (p^1 - p^0)\min\left\{1, K\exp\left(-2\pi(1-\pi)(p^1 - p^0)^2 N\right)\right\} \\
&\leq (p^1 - p^0)\min\left\{1, K\exp\left(-\kappa(p^1 - p^0)^2 N\right)\right\},
\end{aligned}$$

where the second-to-last step used McDiarmid's inequality (McDiarmid, 1989), already invoked in the proof of Proposition 3, applied to $f(y) \equiv \overline{y}_{0,k} - \overline{y}_{1,k}$, and the last step follows from $2\pi(1-\pi) \geq \kappa \equiv \min\{\pi, 1-\pi\}$.

We have that $K\exp(-\kappa(p^1 - p^0)^2 N) \leq 1 \iff p^1 - p^0 \geq \sqrt{\frac{\ln(K)}{\kappa N}}$, which implies that

$$L(p) \leq \begin{cases} p^1 - p^0 & \text{if} \quad p^1 - p^0 < \sqrt{\frac{\ln(K)}{\kappa N}}, \\ K(p^1 - p^0)\exp(-\kappa(p^1 - p^0)^2 N) & \text{if} \quad p^1 - p^0 \geq \sqrt{\frac{\ln(K)}{\kappa N}}. \end{cases} \tag{9}$$

The mapping $x \mapsto x\exp(-\kappa N x^2)$ is ln-concave and maximized at $x = \sqrt{\frac{1}{2\kappa N}}$. Since $K \geq 2$, we have that $\sqrt{\frac{\ln(K)}{\kappa N}} \geq \sqrt{\frac{1}{2\kappa N}}$, which implies that both terms on the right-hand side of (9) are maximized at $p^1 - p^0 = \sqrt{\frac{\ln(K)}{\kappa N}}$. This implies that indeed $L(p) \leq \sqrt{\frac{\ln(K)}{\kappa N}}$. Identical reasoning applies in the case where $p^1 - p^0 < 0$. ■

References

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien, "Optimal Learning by Experimentation," *The Review of Economic Studies*, 1991, *58* (4), 621–654.

Banerjee, Abhijit, Sylvain Chassang, and Erik Snowberg, "Decision Theoretic Approaches to Experiment Design and External Validity," in Esther Duflo and Abhijit Banerjee, eds., *Handbook of Field Experiments*, Elsevier, forthcoming.

Bergemann, Dirk and Juuso Välimäki, "Learning and Strategic Pricing," *Econometrica*, September 1996, *64* (5), 1125–1149.

_ **and** _ , "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 2002, *70* (3), 1007–1033.

_ **and** _ , "Bandit Problems," 2006. Cowles Foundation discussion paper.

Bruhn, Miriam and David McKenzie, "In Pursuit of Balance: Randomization in Practice in Development Field Experiments," *American Economic Journal: Applied Economics*, 2009, *1* (4), 200–232.

Cochrane, William G. and Donald B. Rubin, "Controlling Bias in Observational Studies," *Sankhyā: The Indian Journal of Statistics, Series A*, 1973, *35* (4), 417–446.

Deaton, Angus, "Instruments, Randomization, and Learning about Development," *Journal of Economic Literature*, June 2010, *48* (2), 424–455.

Duflo, Esther, Rachel Glennerster, and Michael Kremer, "Using Randomization in Development Economics Research: A Tool Kit," in T. Paul Schultz and John Strauss, eds., *Handbook of Development Economics, Vol. 4*, Amsterdam: Elsevier, 2008, pp. 3895–3962.

Gerber, Alan S and Donald P Green, "The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment," *American Political Science Review*, 2000, *94* (3), 653–663.

Gilboa, Itzhak and David Schmeidler, "Maxmin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics*, 1989, *18* (2), 141–153.

Green, Donald P and Andrej Tusicisny, "Statistical Analysis of Results from Laboratory Studies in Experimental Economics: A Critique of Current Practice," 2012. Columbia University, *mimeo.*

Grossman, Sanford J and Joseph E Stiglitz, "On the Impossibility of Informationally Efficient Markets," *The American Economic Review*, June 1980, *70* (3), 393–408.

Humphreys, Macartan, Raul Sanchez de la Sierra, and Peter Van der Windt, "Fishing, Commitment, and Communication: A Proposal for Comprehensive Nonbinding Research Registration," *Political Analysis*, 2013, *21* (1), 1–20.

Imai, Kosuke, "Do Get-out-the-vote Calls Reduce Turnout? The Importance of Statistical Methods for Field Experiments," *American Political Science Review*, 2005, *99* (2), 283–300.

Imbens, Guido W, "Better LATE than Nothing: Some Comments on Deaton (2009) and Heckman and Urzua (2009)," *Journal of Economic Literature*, June 2010, *48* (2), 399–423.

Kasy, Maximilian, "Why Experimenters Should not Randomize, and What they Should do Instead," 2013. Harvard University, *mimeo.*

Kitagawa, Toru and Aleksey Tetenov, "Who should be treated? Empirical welfare maximization methods for treatment choice," Technical Report, Collegio Carlo Alberto 2015.

Machina, Mark J, "Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty," *Journal of Economic Literature*, 1989, *27* (4), 1622–1668.

McDiarmid, Colin, "On the Method of Bounded Differences," *Surveys in Combinatorics*, 1989, *141* (1), 148–188.

Miguel, Edward and Michael Kremer, "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica*, January 2004, *72* (1), 159–217.

Morgan, Kari Lock and Donald B. Rubin, "Rerandomization to Improve Covariate Balance in Experiments," *The Annals of Statistics*, 2012, *40* (2), 1263–1282.

Olken, Benjamin A, "Promises and Perils of Pre-Analysis Plans," *The Journal of Economic Perspectives*, 2015, *29* (3), 61–80.

Persico, Nicola, "Information Acquisition in Auctions," *Econometrica*, 2000, *68* (1), 135–148.

Rothschild, Michael, "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 1974, *9* (2), 185–202.

Rubin, Donald B., "Using Multivariate Matched Sampling and Regression Adjustment to Control Bias in Observational Studies," *Journal of the American Statistical Association*, 1979, *74* (366a), 318–328.

_ , "Bias Reduction using Mahalanobis-metric Matching," *Biometrics*, 1980, *36* (2), 293–298.

Saito, Kota, "Preferences for flexibility and randomization under uncertainty," *The American Economic Review*, 2015, *105* (3), 1246–1271.

Savage, Leonard J., *The Foundations of Statistics*, Courier Corporation, 1954.