

Banerjee, Abhijit; Chassang, Sylvain; Montero, Sergio; Snowberg, Erik

Working Paper

A Theory of Experimenters

CESifo Working Paper, No. 6678

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Banerjee, Abhijit; Chassang, Sylvain; Montero, Sergio; Snowberg, Erik (2017) : A Theory of Experimenters, CESifo Working Paper, No. 6678, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/171142>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

A Theory of Experimenters

Abhijit Banerjee, Sylvain Chassang, Sergio Montero, Erik Snowberg

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

A Theory of Experimenters

Abstract

This paper proposes a decision-theoretic framework for experiment design. We model experimenters as ambiguity-averse decision-makers, who make trade-offs between subjective expected performance and robustness. This framework accounts for experimenters' preference for randomization, and clarifies the circumstances in which randomization is optimal: when the available sample size is large enough or robustness is an important concern. We illustrate the practical value of such a framework by studying the issue of rerandomization. Rerandomization creates a trade-off between subjective performance and robustness. However, robustness loss grows very slowly with the number of times one randomizes. This argues for rerandomizing in most environments.

JEL-Codes: C900, D780, D810.

Keywords: experiment design, robustness, ambiguity aversion, randomization, rerandomization.

Abhijit Banerjee
Massachusetts Institute of Technology
USA - Cambridge, MA 02139
banerjee@mit.edu
economics.mit.edu/faculty/banerjee/

Sylvain Chassang
New York University
USA - New York, NY 10012
chassang@nyu.edu
nyu.edu/projects/chassang/

Sergio Montero
University of Rochester
USA - Rochester, NY 14627
smontero@rochester.edu
sas.rochester.edu/psc/smontero/

Erik Snowberg
UBC, Caltech
USA - 91125 Pasadena CA
snowberg@caltech.edu
hss.caltech.edu/~snowberg/

September 14, 2017

We are grateful to Angus Deaton, Pascaline Dupas, Jeff Ely, Guido Imbens, Charles Manski, Pablo Montagnes, Marco Ottaviani, Bruno Strulovici, Alexey Tetenov, Duncan Thomas, Chris Udry, as well as audience members at Bocconi, the Cowles Econometric Conference (2016), the ESSET Meeting at Gerzensee (2017), Emory, INSEAD, MIT, the NBER Development Economics Summer Meeting (2016), the North-American Econometric Society Meeting (2016), the SISL Conference at Caltech (2017), UBC, and 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

The proliferation of experiments in academia, business, and public policy has been accompanied by spirited debates about best practices for experiment design and the statistical analysis of experimental results. Topics of debate include pre-registration of experiment designs, pre-analysis plans, the pros and cons of rerandomization, clustering, stratification, and statistical significance testing (Bruhn and McKenzie, 2009; Deaton, 2010; Duflo et al., 2008; Humphreys et al., 2013; Imbens, 2010; Olken, 2015; Athey and Imbens, 2016; Benjamin et al., 2017). At the heart of these debates are differing—usually implicit—models of knowledge generation and how people interpret experimental evidence. This paper seeks to provide a decision-theoretic framework that may be able to guide, and help resolve, these debates. We show that our framework is consistent with important elements of observed experimental practice. We then use our framework to shed light on one of the less contentious open questions of experimental design, rerandomization. Finally, we discuss how our framework might be used to address other questions about experiment design.

Unfortunately, models of experimentation and information acquisition fail to explain why researchers (almost always) run randomized controlled trials (RCTs; see Kasy, 2016).¹ This is due to modeling decision-makers as subjective expected utility maximizers (Savage, 1954). As RCTs are mixed strategies over experimental assignments, they can never be strictly optimal for such a decision-maker.

To overcome this limitation, we model a Bayesian experimenter facing an adversarial audience. This is equivalent to assuming an ambiguity-averse decision-maker with standard maxmin preferences (Gilboa and Schmeidler, 1989). In our formulation, the decision-maker maximizes a mixture of her own subjective expected utility term and the welfare of an adversarial audience with non-common priors. Examples of such audiences abound: the

¹For seminal contributions to the literature on experimentation and information acquisition Rothschild (1974); Grossman and Stiglitz (1980); Aghion et al. (1991); Bergemann and Välimäki (1996); Persico (2000); Bergemann and Välimäki (2002, 2006).

Food and Drug Administration for drug trials, seminar audiences and journal referees for research papers, and governments or NGOs for development proposals.²

The paper reports two main sets of results. The first set of results shows that RCTs can be optimal for a decision-maker facing an adversarial audience, and clarifies the circumstances in which this 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. Indeed, as the sample size N grows large, RCTs permit robust prior-free inference, and achieve assignment losses of order $1/\sqrt{N}$. On the other hand, deterministic experiments are generically strictly optimal when the sample size is small and the decision-maker puts sufficiently high weight on her subjective expected utility.

This set of results accords with the observed heterogeneity of experimental practice. 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, optimal experiments are deterministic as such a design optimizes the informational value of each acquired data point (new product launches in select markets, preliminary medical research on diseases that would otherwise result in quick and certain death).

The second set of results applies our framework to the question of whether experimenters should rerandomize to improve covariate balance between treatment and control groups. Rerandomization draws multiple treatment assignments, then chooses the one that maximizes balance. 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 (Morgan and Rubin, 2012). Despite the ease of using rerandomization to improve balance, researchers are concerned that it may adversely affect the

²The audience could also be another side of the decision-maker herself, as ambiguity aversion is a natural way to model self-doubt.

reliability of findings (Bruhn and McKenzie, 2009).

The trade-offs at the heart of rerandomization are succinctly 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 is always bounded away from first-best. In contrast, robustness losses are vanishing in sample size N when the number of rerandomizations grows linearly or sublinearly in N .

Our framework builds on a long line of work in statistics, starting with Wald (1950). Variants of Wald’s framework have more recently been used in economics and econometrics to study questions of identification and model uncertainty (Manski, 2004, 2009; Kitagawa and Tetenov, 2015; Marinacci, 2015). However, prior work has not sought to provide a positive and normative model of experimenters.³ The prior literature falls into two broad groups: one focuses on ambiguity aversion, the other on regret minimization. While our main results are stated in an ambiguity-aversion framework, in Section 5 we show they extend to a regret-minimization framework. This is important as Tetenov (2012) shows that asymmetric regret minimization is sufficient for explaining the standard practice of null-hypothesis statistical testing (NHST) using t -statistics. Our basic insights therefore apply directly to classical statistical inference as well as to our more conventional decision-theoretic model.

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. Section 5 shows that our results extend to reference-dependent preferences, better suited to explain the use of t -statistics in decision-making. Section 6 discusses other applications for our framework,

³This paper is also related to the dormant literature in multi-Bayesian statistical decision theory (Weerahandi and Zidek, 1981, 1983). In these models, Bayesians with conflicting preferences adopt random decision rules, rather than randomized experiments.

including subgroup analysis. Appendix A explores the use of reference-dependent preferences as a foundation for t -statistics. Proofs are contained in Appendix B. Appendix C presents simulations.

2 A Framework for Studying 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 unit mass of individuals. For simplicity, we assume that the final policy choice $a \in \{0, 1\}$ is all-or-nothing and sets $\tau = a$ for all individuals.⁴ Potential outcomes for an individual with treatment status τ are random variables $Y^\tau \in \{0, 1\}$; $Y = 1$ is referred to as a success. Each individual has observable covariates $x \in X \subset \mathbb{R}^m$ that affect the distribution of outcomes Y . X is finite, and the distribution $q \in \Delta(X)$ of covariates in the population is known and has full support. The probability of success given covariate x is denoted by $p_x^\tau \equiv \text{prob}(Y^\tau = 1|x)$, and, conditional on x , outcomes are i.i.d.

The state of the world is described by the finite-dimensional vector $p = (p_x^0, p_x^1)_{x \in X} \in [0, 1]^{2X} \equiv P$ of success probabilities p_x^τ conditional on treatment status $\tau \in \{0, 1\}$ and covariates x . 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

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

Although covariates x are observable, our framework implicitly allows for unobservable characteristics. Denoting unobserved characteristics by z , we would have $p_x^\tau = \int p_{x,z}^\tau dF(z|x)$. Large shifts in correlations between x and z are captured by allowing the mapping $x \mapsto p_x$ to be discontinuous in x .

⁴See Section 6 for a discussion of sub-group analysis and targeting.

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_i s independent realizations of $Y_i^{\tau_i}$ given (x_i, τ_i) . Throughout, we assume that $N < 2|X|$, even as we increase N . A natural special case is unique individuals: $|X| = N$ and $q(X) = \frac{1}{|X|}$. This corresponds to settings where subjects are fixed, and the decision-maker must choose a treatment assignment.

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 $\alpha : E \times \mathcal{Y} \rightarrow \Delta(\{0, 1\})$, which maps experimental data (e, y) to policy decisions $a \in \{0, 1\}$. We denote by \mathcal{A} the set of such mappings.

Preferences. We consider an ambiguity-averse decision-maker with standard maxmin preferences (Gilboa and Schmeidler, 1989). We reinterpret these preferences in a way that is convenient for comparative static exercises. Under this interpretation, the decision-maker has her own prior h_0 and faces an adversarial audience of Bayesian stakeholders with non-common priors.⁵ She chooses a strategy (\mathcal{E}, α) to solve

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

where H is a convex set of priors $h \in \Delta(P)$ over states $p \in P$, and represents the set of priors belonging to members of the audience. The decision-maker places weight λ on her own subjective utility, and weight $1 - \lambda$ on the utility of her audience.⁶ When $\lambda = 1$, these

⁵If the audience entertained a common prior h , then the decision problem would boil down to subjective expected utility maximization for the mixed prior $\lambda h_0 + (1 - \lambda)h$.

⁶Whenever $h_0 \in H$, increasing $1 - \lambda$ also corresponds to increasing the set of priors that the decision-maker entertains as plausible. If audience members have veto power and enjoy a common outside option, then $\frac{1-\lambda}{\lambda}$ is the Lagrange multiplier placed on the audience's individual rationality constraint.

preferences coincide with standard subjective expected utility maximization. We sometimes refer to this case as *Bayesian*.

Equivalent experiments. As successes are independent conditional on covariates, experiments that differ only by a permutation of participants with identical covariates are equivalent from a decision-making perspective.

Definition 1 (equivalent experiments). *Two experiments $e = (x_i, \tau_i)_{i \in \{1, \dots, N\}}$ and $e' = (x'_i, \tau'_i)_{i \in \{1, \dots, N\}}$ are equivalent, denoted by $e \sim e'$, if there exists a permutation $\sigma : \{1, \dots, N\} \rightarrow \{1, \dots, N\}$ 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 $[e]$.⁷ We denote by $[E]$ the partition of possible experiments in 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 $\mathcal{E} \sim \mathcal{E}'$, $\max_{\alpha \in \mathcal{A}} U(\mathcal{E}, \alpha) = \max_{\alpha \in \mathcal{A}} U(\mathcal{E}', \alpha)$.*

All proofs can be found in Appendix B.

Standard RCTs. As many of our results deal with randomized controlled trials (RCTs), it is useful to define these explicitly. A standard RCT that assigns 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}_{RCT} samples N exchangeable participants labelled by $i \in \{1, \dots, N\}$, with covariates $(x_i)_{i \in \{1, \dots, N\}}$ drawn according to the distribution of observable covariates q ;
- \mathcal{E}_{RCT} assigns treatment $\tau_i = \mathbf{1}_{i \leq \pi N}$;
- $\alpha_{\text{RCT}}(e, y) \equiv \mathbf{1}_{\hat{y}^1 \geq \hat{y}^0}$, where $\hat{y}^\tau \equiv \sum_{i=1}^N y_i \mathbf{1}_{\tau_i = \tau} / \sum_{i=1}^N \mathbf{1}_{\tau_i = \tau}$ is the mean outcome for participants with treatment status τ .

⁷It is convenient to include distributions \mathcal{E} with support in $[e]$ in the equivalence class of e .

2.1 Assumption about the Audience

We impose the following *limited extrapolation* condition on X , N , and H : Denote by $p^a \equiv \sum_{x \in X} q(x)p_x^a$ the expected probability of success given policy $a \in \{0, 1\}$. Given an experiment $e = (\tau_i, x_i)_{i \in \{1, \dots, N\}}$, denote by $\bar{p}_e \equiv (p_{x_i}^{\tau_i})_{i \in \{1, \dots, N\}}$ the subset of success rates for participants in the experiment. Vector \bar{p}_e is an upper bound to the information generated by experiment e in the sense of Blackwell (1953).

Assumption 1 (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 \bar{p}_e ,*

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

This imposes two important constraints on the members of the adversarial audience. First, the information from any experiment is insufficient to fully convince every single audience member (represented by some prior $h \in H$) of the correct policy $a \in \{0, 1\}$. Second, audience members cannot be arbitrarily pessimistic; they cannot be certain that $p^1 = p^0 = 0$.

Note that the requirement that $N < 2|X|$ ensures that Assumption 1 can be satisfied. We show in Section 5 that Assumption 1 is dispensable when the decision-maker exhibits regret aversion.

3 Optimal Design and Randomization

3.1 Bayesian Experimentation

When $\lambda = 1$, the decision-maker is a standard subjective expected utility maximizer. In this case, it is known that deterministic experiments are *weakly* optimal. In fact, we show that

for generically every prior—that is, for an open and dense set of priors under an appropriate distance—deterministic experiments are *strictly* optimal when λ is close to 1.⁸

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

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 (DP) if and only if $\text{support } \mathcal{E} \subset [e^*]$.*

In recent work, Kasy (2016) uses a similar result to argue that RCTs may be suboptimal. However, his key point is that deterministic assignments are much more likely to achieve covariate balance between treatment and control groups. Instead, we suggest that Proposition 2 shows the limits of subjective expected utility maximization as a positive model of experimenters. We argue that the decision problem defined by (DP) is more successful at explaining the range of information acquisition strategies observed in practice. Later, in Section 4, we study rerandomization as a non-deterministic alternative to improving balance.

3.2 Adversarial Experimentation

We now examine the case where the experimenter cares about her audience’s preferences.

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*

$$\max_{\alpha \in \mathcal{A}} \min_{h \in H} \mathbb{E}_{h, \mathcal{E}^*} [u(p, \alpha(e, y))] \geq \min_{h \in H} \mathbb{E}_h \left[\max_{a \in \{0, 1\}} u(p, a) \right] - \sqrt{\frac{\ln 2}{N}}.$$

⁸We use the uniform norm on distributions: $d(h, h') \equiv \sup_{A \subset P} |h(A) - h'(A)|$.

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

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

where ξ is defined in Assumption 1.

Point (i) shows that the efficiency loss of the optimal experiment compared to the first-best decision is bounded above by a term of order $1/\sqrt{N}$. Point (ii) 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; therefore, it must be randomized.

The rationale for randomization can be understood by observing that the decision-maker is playing a zero-sum game against nature. The decision-maker first chooses an experimental design and a decision rule. Nature then picks the prior which maximizes the chance of the experimenter choosing the wrong policy, given the decision-maker's experiment design. If there is a known pattern in the choice of experimental assignments, nature can exploit it to lower the decision-maker's payoff. Randomization eliminates patterns that nature can exploit. This is related to the fact that ambiguity-averse agents may have preferences for randomization even if they exhibit risk-aversion over known lotteries (Saito, 2015).⁹

3.3 RCTs as a Rule of Thumb

A corollary of the proof of Proposition 3 is that the standard RCT $(\mathcal{E}_{\text{RCT}}, \alpha_{\text{RCT}})$ (defined in Section 2) provides a near optimal solution to decision problem (DP).

⁹A key modeling choice here is that nature cannot observe the outcome of the decision-maker's randomization before picking a prior. Kasy (2016) assumes that nature observes the outcome of the experimenter's randomization and then picks a prior, which renders randomization useless. We believe our assumption is more consistent with research practice: Referees typically complain about research design, not about why a particular observation ended up in the treatment or control group. The exception is when there is an extreme imbalance between treatment and control on a covariate the referee considers important.

Corollary 1. *Experimentation policy $(\mathcal{E}_{\text{RCT}}, \alpha_{\text{RCT}})$ is such that for all priors $h \in \Delta([0, 1]^X)$,*

$$\mathbb{E}_{h, \mathcal{E}_{\text{RCT}}} [u(p, \alpha_{\text{RCT}}(e, y))] \geq \mathbb{E}_h \left[\max_{a \in \{0, 1\}} u(p, a) \right] - \sqrt{\frac{\ln 2}{2\underline{\pi}N}},$$

where $\underline{\pi} \equiv \min\{\pi, 1 - \pi\}$.

As this result holds for every prior h , standard RCTs allow approximately optimal decision-making for both Bayesian and ambiguity-averse decision-makers. Thus, an RCT can be interpreted as a rule-of-thumb that avoids the complexity of specifying the space of priors in decision problem (DP) and deriving an optimal solution.


Note that Corollary 1 holds even though there may be arbitrarily many covariates, some of which have a significant impact on treatment effects, and exact balance cannot be ensured. As outcomes are bounded, it is not possible for rare—and therefore hard to balance—covariate realizations to have a substantial impact on payoffs. Balance in expectation is sufficient to guarantee approximately optimal decision-making.

3.4 Positive Implications

Figure 1 maps out some positive implications of Propositions 2 and 3 for experiment design. When sample points are scarce, or when the decision-maker does not put much weight on satisfying her audience (λ close to 1), the optimal experiment 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 at a few production sites: The firm will focus on the places where it can learn the most. This may, for example, involve assigning treatment to sites where success is the least likely (Banerjee et al., 2017). Similarly, early stage medical research often does not feature randomization—especially when treating conditions known to result in severe disability or death.¹⁰ In this case, there is no

¹⁰For recent examples, see Harmon (2016) and Yuhas (2016).

Figure 1: When to randomize? Theory matches broad patterns.



adversarial audience. Scientists are trying to learn whether a particular research direction is worth exploring.

When the decision-maker must instead satisfy a sufficiently adversarial audience, or she has a sufficiently large sample, she will randomize. The former is the case in scientific research. The latter is the case for firms doing A/B testing online. 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 effectively address internal concerns about robustness. This is also the case for later stage medical research seeking regulatory approval: government regulators and investors form the adversarial audience for pharmaceutical and medical device companies.

4 Rerandomization as a Refined Rule-of-Thumb

Having established that our framework makes reasonable positive predictions about experimental practice, we now use it to shed light on an open question of experimental design:

rerandomization. As noted in Section 3.3, a standard RCT is only a near-optimal solution to (DP). In particular, it may sometimes result in an unbalanced sample that permits little real opportunity for learning.

This could be avoided by computing an optimal solution to (DP). For appropriate problems, the solution is a distribution over possible assignments that puts greater weight on balanced assignments, and may put zero weight on excessively unbalanced assignments. This distribution maintains sufficient randomization to ensure robust decision-making. Unfortunately, it is implausible that one could carefully and reliably elicit full preferences from experimenters and their audience at the onset of an experiment. Computing an optimal solution to (DP) is a tricky exercise, and a well designed rule-of-thumb may be more valuable.

In practice, experimenters resolve this difficulty through rerandomization: they repeatedly randomize until they obtain an assignment that satisfies their balance objectives. As Bruhn and McKenzie (2009) highlight, this practice is widespread, poorly theorized and—in principle—a substantial deviation from one-shot randomization. Our framework can be used to clarify the trade-offs involved in rerandomization. We argue that, used in moderation, rerandomization is a valuable rule-of-thumb that provides a flexible way to improve balance—and/or reflect ancillary design objectives—without sacrificing significant decision-making robustness. Compared to other common approaches to ensuring balance, it has the added benefit of carefully controlling losses in robustness.

4.1 *K*-rerandomized Experiments

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

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

Decision problem (DP) sets $B(e) \equiv \mathbb{E}_{h_0}[u(p, \alpha(e, y))|e]$, capturing “balance”, and $R(\mathcal{E}) \equiv \min_{h \in H} \mathbb{E}_{h, \mathcal{E}}[u(p, \alpha(e, y))]$, capturing robustness. We re-write (DP) in this form to emphasize

that our results hold for *any* objective function $B(e)$ the experimenter, or other stakeholders, might have.

With this mapping, we define a K -rerandomized experiment \mathcal{E}_K as follows:

1. Independently draw a set of K assignments $\{e_1, \dots, e_K\}$ 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 $e_K^* \in \operatorname{argmax}_{e \in \{e_1, \dots, e_K\}} B(e)$ that maximizes balance, breaking ties with a uniform draw.
3. Run experiment e_K^* , generating outcomes y_K .
4. Policy is chosen according to $\alpha_{\text{RCT}}(e_K^*, y_K) \equiv \mathbf{1}_{\hat{y}_K^1 \geq \hat{y}_K^0}$, where \hat{y}_K^τ is the mean outcome for participants with treatment status τ .

Rerandomization protocols 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. However, our results show open-ended stopping times may lead to significant robustness losses.

4.2 The Trade-off of Rerandomization

For any balance function, the gains from rerandomization are immediately apparent.

Remark 1. $B(e_K^*)$ *first-order stochastically dominates* $B(e_{K-1}^*)$.

The value of balance is formally and intuitively clear. In the context of statistical inference, Morgan and Rubin (2012) study the value of rerandomization when $B(e) = -M(e)$, in which $M(e)$ is the Mahalanobis distance between the treatment and control samples. They show that rerandomization leads to significantly more precise estimates of treatment effects when outcomes come from a linear Gaussian model.¹¹ Bungi et al. (2016) show, more

¹¹The Mahalanobis distance—defined as $M(e) \equiv (\hat{x}^1 - \hat{x}^0)'[\operatorname{cov}(\hat{x}^1 - \hat{x}^0)]^{-1}(\hat{x}^1 - \hat{x}^0)$ —is commonly used in multivariate matching (Rubin, 1980; Cochran and Rubin, 1973; Rubin, 1979). Outcomes follow a Gaussian linear model when they are defined by $Y_i^{\tau_i} = \langle x_i, \beta \rangle + \delta\tau_i + \sigma\varepsilon_i$, with $\varepsilon_i \sim \mathcal{N}(0, 1)$.

generally, that balance from symmetric stratification procedures also substantially increases precision.

The Cost of Rerandomization. Although rerandomization provides clear benefits in terms of improved balance, there are common, but vague, concerns about its potential costs. As Bruhn and McKenzie (2009) document, this leads many researchers to use rerandomization but not report it. Our framework clarifies the issue by showing that rerandomization can reduce robustness, but that this cost is relevant only when the number of rerandomizations is very large:

Proposition 4. *There exists $\underline{\rho} > 0$ such that for all 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 becomes essentially deterministic. Proposition 3 implies that this precludes first-best robustness.¹² However, the number of rerandomizations K needed to cause fixed losses in robustness is exponential in the sample size.

A moderate number of rerandomizations has little impact on robustness:

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 for all $h \in H$,*

$$\mathbb{E}_{h, \mathcal{E}_K} [u(p, \alpha_{RCT}(e, y))] \geq \mathbb{E}_h \left[\max_{a \in \{0,1\}} u(p, a) \right] - \sqrt{\frac{\ln K}{\underline{\pi} N}},$$

where $\underline{\pi} \equiv \min\{\pi, 1 - \pi\}$.

¹²The process by which a realized experiment e is reached is irrelevant for a Bayesian: $h_0(p | e, y, e \sim \delta_e) = h_0(p | e, y, e \sim \mathcal{E}_{RCT}) = h_0(p | e, y, e \sim \mathcal{E}_K)$, in which δ_e denotes the deterministic experiment e .

As the additional loss from rerandomization is of order $\sqrt{\ln K}$, which is less than 5 for $K \leq 10^{10}$, K linear in N does not change the order of magnitude of robustness losses. This suggests the following guideline:

Suggested Guideline. *Set $K \leq N$.*¹³

Remark 1 and Propositions 4 and 5 clarify the pros and cons of rerandomization, but do not provide a hard and fast rule. Our suggested guideline defines a range of acceptable behavior that limits potential robustness-losses from rerandomization. Within this range, different decision-makers may choose differently. A Bayesian decision-maker using the linear Gaussian model studied by Morgan and Rubin (2012) should choose the largest possible number of rerandomizations. An ambiguity-averse decision-maker entertaining a symmetric set of possible priors may prefer to randomize only once (see Appendix C.1 for a stylized numerical example).

4.3 Other Approaches for Increasing Balance

Rerandomization can be seen as an algorithm to bias the distribution of experimental assignments towards high-balance ones. It is simpler than many matching algorithms, especially when one wants to establish balance on multiple continuous covariates. Moreover, as our results hold for any balance function $B(e)$, they are also true if this function is specified after seeing the K possible rerandomizations. As we highlight in Section 6, this degree of freedom may be very useful in practice as a way to respond to the preferences of stakeholders and implementation partners.

It is instructive to relate rerandomization to an experiment design that ensures balance in a more direct way: by randomizing conditional on strata. Formally, the experimenter first defines a set of acceptable assignments $E_{\dagger} \subset E$. Then, an assignment is drawn from

¹³The probability that the rerandomized assignment will be in the top 5% most balanced is $1 - 0.95^K > 99\%$ when $K = 100$. Thus, experimenters may wish to limit the number of rerandomizations to 100, even if N is much larger (Banerjee et al., 2017).

E_{\dagger} with uniform probability. This procedure involves two difficulties. First, if E_{\dagger} is defined implicitly—for example, the set of all assignments whose balance is above a certain threshold—computing it may be quite difficult. Second, this procedure obfuscates potential robustness costs. Indeed, if the set of acceptable assignments E_{\dagger} turns out to be small, the intuition underlying Proposition 4 suggests that robustness will be bounded away from first-best.¹⁴

Our framework can be used to explicitly evaluate the potential robustness costs of such approaches. Note that the procedure described above sets $B(e) \equiv \mathbf{1}_{e \in E_{\dagger}}$. Define $\mathcal{E}_{E_{\dagger}}$ as the randomized experiment that picks an assignment $e \in E_{\dagger}$ with uniform probability. Finally, let $p_{E_{\dagger}} > 0$ denote the probability that a uniformly chosen assignment $e \in E$ belongs to E_{\dagger} . Note that computing p_{\dagger} may be difficult, especially for complex or opaque matching algorithms. In that case, Monte Carlo simulations may provide an approximate value for $p_{E_{\dagger}}$.

Corollary 2. *For all $h \in H$,*

$$\mathbb{E}_{h, \mathcal{E}_{E_{\dagger}}} [u(p, \alpha_{RCT}(e, y))] \geq \mathbb{E}_h \left[\max_{a \in \{0,1\}} u(p, a) \right] - \min_{K_{\dagger} \in \mathbb{N}} \frac{1}{1 - (1 - p_{\dagger})^{K_{\dagger}}} \sqrt{\frac{\ln K_{\dagger}}{\pi N}}$$

This implies that a stratification or matching procedure will come at a limited cost of robustness if an acceptable assignment can be reached with high probability within a small number of rerandomizations.

5 Extension to Reference Dependent Preferences

¹⁴This may occur when an experimenter uses a matching algorithm to achieve balance across several continuous covariates. Note that stratification and matching can considerably improve the precision of estimates (Athey and Imbens, 2016; Bungi et al., 2016). However, it is important to assess whether a proposed design limits the set of possible assignments so severely that it causes significant losses in robustness.

It has been shown that decision problem (DP) does not rationalize the use of null-hypothesis statistical testing (NHST) using t -statistics. For this, we need reference-dependent preferences with the status quo as the reference point (Tetenov, 2012).¹⁵ As most researchers use NHST, it is important to extend our main results to such reference-dependent preferences.

Interestingly, this extension also allows us to dispense with the limited extrapolation assumption. One of the key features of this assumption is to ensure that learning is still possible even for the worst-case prior for any experimental design. Without this constraint, the worst prior would always put unit mass on the probability of success being zero, regardless of the policy choice. However, when considering regret minimization, nature favors environments in which the decision-maker can make an incorrect decision—otherwise there is no possibility of regret. As a result, under regret aversion, there is always a value for information even at the unconstrained worst-case prior.

Let $\Delta^a \equiv p^a - p^{1-a}$ denote true expected outcome differences between action $a \in \{0, 1\}$ and the alternative. We consider a decision-maker who seeks to solve

$$\max_{\mathcal{E}, \alpha} \lambda \mathbb{E}_{h_0}[w_\alpha(\Delta^a)] + (1 - \lambda) \min_{h \in H} \mathbb{E}_h[w_\alpha(\Delta^a)]. \quad (\text{DP}')$$

in which, for all $a \in \{0, 1\}$,

$$w_a(\Delta^a) = \Delta^a + \kappa_a \times \Delta^a \mathbf{1}_{\Delta^a < 0}, \quad \text{with } \kappa_a > 0.$$

In words, the decision-maker cares about improvements rather than absolute levels, and is $\kappa_a > 0$ implies she is particularly averse to making wrong choices. Note that H in (DP') is now the set of all priors over P . For simplicity, we assume $N = |X|$, so that each subject is unique.¹⁶

¹⁵See Appendix A as well as Tetenov (2012) for micro-foundations of the standard approach to NHST.

¹⁶Our results extend to the case when $N < 2|X|$, as before. Moreover, our analysis applies more generally to functions w_a that are strictly increasing and concave in Δ^a . We restrict w_a to the form above to clarify the relationship with the prior literature.

We show in Appendix A that this class of preferences rationalizes the use of t -statistics whenever $\kappa_1 > \kappa_0$, that is, when a preference for success is not symmetric with respect to the policy choice a . In addition it coincides with pure regret-averse preferences when κ_0 and κ_1 grow large.¹⁷

The conditions under which deterministic and randomized experiments are optimal are qualitatively unchanged:

Proposition 6. *Consider a decision-maker solving Problem (DP'):*

- (i) *For generically every prior h_0 , there exists $\underline{\lambda} \in (0, 1)$ such that, for all $\lambda > \underline{\lambda}$, it is optimal to use a deterministic experiment.*
- (ii) *For every h_0 and every $\lambda \in (0, 1)$, as N becomes arbitrarily large, deterministic experiments remain bounded away from efficiency, and randomized experiments are strictly optimal.*

Additionally, the near-optimality of K -randomized trials continues to hold:

Proposition 7. *There exists $M > 0$ such that for every prior $h \in \overline{H}$,*

$$\mathbb{E}_h [w_a(\Delta^a) \mid e \sim \mathcal{E}_K, a \sim \alpha_{RCT}] \geq \mathbb{E}_h \left[\max_{a \in \{0,1\}} w_a(\Delta^a) \right] - M \sqrt{\frac{\ln K}{\underline{\pi} N}}.$$

6 Discussion

The ambition of this paper is to propose a decision-theoretic framework that takes seriously the concerns of experimenters, and can help clarify current experimental practices as well as resolve open debates. In this final section, we discuss further implications of our framework, and preview other possible applications.

¹⁷Regret-averse preferences have received extensive attention from statisticians and econometricians (Wald, 1950; Manski, 2004, 2009; Kitagawa and Tetenov, 2015). However, failure of transitivity can reduce their appeal (see, for example, Marinacci, 2015).

Why randomize? Our analysis clarifies the circumstances in which randomization is optimal: when robustness matters and the sample size is large. When robustness considerations are motivated by an adversarial audience with non-common priors, randomization can be interpreted as a way to let parties with diverging priors agree on a process. Indeed, stakeholders with divergent priors need not be satisfied by any given deterministic experiment: there may always exist a prior under which the experiment’s design is flawed. In contrast, running a randomized experiment guarantees each stakeholder that, in expectation, the final decision will be optimal from her perspective. However, there is a tension here: when the assignment is revealed, some stakeholders may wish to change it *ex post*.

Whether robustness is motivated by an external audience, or internal doubts, our analysis 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 Tuscisny, 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 her audience. Randomization provides a defense against the most skeptical priors; near-randomization does not.

Rerandomization as an additional degree of freedom. Our framework also clarifies the trade-offs of rerandomization: it improves subjective balance, but comes at the expense of robustness. However, the loss in robustness from rerandomization grows slowly in the number of times we rerandomize. As a result, applied in moderation, rerandomization provides a useful degree of freedom for experimenters who have subjective preferences over balance.

While pre-specifying the objective function used in rerandomization has the added benefit of permitting randomization inference tests (Fisher, 1935; Morgan and Rubin, 2012; Young, 2016; see Appendix C.2.3 for examples), we emphasize that it is not a requirement for our results to hold. Indeed, Proposition 5 remains true if the balance function $B(e)$ is specified *after* experimental assignments $\{e_1, \dots, e_K\}$ are drawn.

This suggests a novel use of rerandomization: as a way to let stakeholders, such as implementation partners, express preferences over assignments. The experimenter may draw K treatment assignments in advance and let her implementation partners pick their preferred assignment. This process could be used with any stakeholder, including regulators, funders, or the communities from which experimental participants are drawn. These stakeholders often have equity concerns and distributional preferences. They may further care about targeting treatment to those they believe will benefit the most, or may simply dislike randomization and wish to exert some control over assignment. The ability to at least *partially* accommodate the preferences of stakeholders by letting them pick their preferred assignment among K options may help build good will and ensure cooperation.¹⁸

Extending the framework: Subgroup analysis. While we believe that our framework is well suited to extensions designed to answer other important questions about experimental practice. We illustrate this by providing an outline of how our framework may be extended to address another important methodological question: sub-group analysis.

Consider a partition of the population into subgroups $G \in \mathcal{G}$ where \mathcal{G} is a partition of X . Treatment and policy α may now be targeted by subgroup. Assignment α becomes a mapping from experimental outcomes (e, y) to a distribution over targeted assignments $\mathbf{a} = (a_G)_{G \in \mathcal{G}} \in \{0, 1\}^{\mathcal{G}}$. The decision-maker's problem is now:

$$\max_{\mathcal{E}, \alpha} \lambda \mathbb{E}_{h_0, \mathcal{E}, \alpha}[u(p, \mathbf{a})] + (1 - \lambda) \min_{h \in H} \mathbb{E}_{h, \mathcal{E}, \alpha}[u(p, \mathbf{a})] \quad (\text{DP}'')$$

where $u(p, \mathbf{a}) = \sum_{G \in \mathcal{G}} \sum_{x \in G} q(x) p_x^{a_G}$.

This framework can provide guidelines on how finely to target treatment as a function of the data and sample size. An ambiguity-averse decision-maker would not want to tailor assignments based on within-group treatment effects when the number of groups $G \in \mathcal{G}$ is

¹⁸Constraints from implementation partners led Miguel and Kremer (2004) to assign treatment alphabetically, eliciting concerns from Deaton (2010).

of order N . If treatment effects are constant across groups, targeting leads to inefficient treatment choices for a non-vanishing share of the population. A Bayesian decision-maker putting sufficient weight on treatment effects being constant or correlated across groups would come to the same conclusion. On the other hand, if the number of groups is small compared to N , and the decision-maker's subjective prior entertains the possibility that the optimal treatment is different across groups, then the solution to (DP'') will target by group (and approach first-best efficiency).

More generally, we believe that related decision-theoretic frameworks are needed to address the trickier aspects of experiment design. For example, if a random, uninformed statement is included in a pre-analysis plan, should this affect inference? It seems logical that the answer would be no, but this calls into question the entire exercise. Or, why should formulating an additional hypothesis affect our confidence in the analysis of others? Yet, this is the impact of corrections for multiple hypotheses testing. These questions, as well as others, await a rigorous treatment with solid decision-theoretic foundations.

References

- Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien, “Optimal Learning by Experimentation,” *The Review of Economic Studies*, 1991, 58 (4), 621–654.
- Athey, Susan and Guido Imbens, “The Econometrics of Randomized Experiments,” *arXiv preprint arXiv:1607.00698*, 2016.
- 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, 2017, pp. 141–174.
- Benjamin, Daniel, James Berger, Magnus Johannesson, Brian Nosek, Eric-Jan Wagenmakers, Richard Berk, Kenneth Bollen, Bjorn Brembs, Lawrence Brown, and Colin Camerer, et al., “Redefine Statistical Significance,” *Nature Human Behaviour*, 2017, *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.
- Blackwell, David, “Equivalent Comparisons of Experiments,” *The Annals of Mathematical Statistics*, 1953, 24 (2), 265–272.
- 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.
- Bungi, Federico A., Ivan A. Canay, and Azeem M. Shaikh, “Inference under Covariate-Adaptive Randomization,” 2016. *mimeo*. Duke University.
- 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.

- Fisher, Ronald Aylmer**, *The Design of Experiments*, Edinburgh and London: Oliver & Boyd, 1935.
- 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.
- Harmon, Amy**, “After Long Fight, Drug Gives Sudden Reprieve,” *The New York Times*, February 22, 2016. Published online at: <http://www.nytimes.com/2010/02/23/health/research/23trial.html>.
- 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.
- 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 Might not Always Want to Randomize, and what they Could Do Instead,” *Political Analysis*, 2016, 24 (3), 324–338.
- Kitagawa, Toru and Aleksey Tetenov**, “Who should be Treated? Empirical Welfare Maximization Methods for Treatment Choice,” 2015. Collegio Carlo Alberto, *mimeo*.
- Luenberger, David G.**, *Optimization by Vector Space Methods*, New York: John Wiley & Sons, 1969.
- Manski, Charles F.**, “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 2004, 72 (4), 1221–1246.
- , “The 2009 Lawrence R. Klein Lecture: Diversified Treatment under Ambiguity,” *International Economic Review*, 2009, 50 (4), 1013–1041.
- Marinacci, Massimo**, “Model Uncertainty,” *Journal of the European Economic Association*, 2015, 13 (6), 1022–1100.
- 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.
- Tetenov, Aleksey**, “Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria,” *Journal of Econometrics*, 2012, 166 (1), 157–165.
- Wald, Abraham**, *Statistical Decision Functions*, New York: Wiley, 1950.
- Weerahandi, Samaradasa and James V. Zidek**, “Multi-Bayesian Statistical Decision Theory,” *Journal of the Royal Statistical Society. Series A (General)*, 1981, 144 (1), 85–93.
- and —, “Elements of multi-Bayesian decision theory,” *The Annals of Statistics*, 1983, 11 (4), 1032–1046.
- Young, Alwyn**, “Channelling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results,” 2016. *mimeo.*, London School of Economics.
- Yuhas, Alan**, “Cancer Researchers Claim ‘Extraordinary Results’ using T-cell Therapy,” *The Guardian*, February 15, 2016. Published online at: <https://www.theguardian.com/science/2016/feb/15/cancer-extraordinary-results-t-cell-therapy-research-clinical-trials>.

Online Appendix—Not Intended for Publication

A Rationalizing t -statistics

A primary goal of this paper is to provide a suitable positive model of experimental practice that can be used to examine and guide innovations. However, decision problem (DP) is incompatible with hypothesis testing using t -statistics, a mainstay of experimental practice. The raw treatment effect—that is, the difference in average outcomes $\Delta^1 \equiv \bar{p}^1 - \bar{p}^0$ —is sufficient for near-optimal decision-making.

This appendix describes a class of experimenter preferences that rationalize the use of t -statistics. Hypothesis testing favors implementation of the status quo (or null) treatment $a = 0$. We first clarify that standard preferences (including risk-aversion) do not rationalize t -statistics. We then turn to reference-dependent preferences.

Ambiguity aversion does not play a role in this argument. We consider a decision-maker with independent Gaussian posteriors $\mathcal{N}(\hat{p}^a, \frac{\sigma_a^2}{N})$ over the mean outcome \bar{p}^a of action $a \in \{0, 1\}$.¹ A risk-neutral Bayesian decision-maker solving $\max_{a \in \{0, 1\}} \mathbb{E}[\bar{p}^a]$ will simply take action $a = 1$ if and only if $\hat{p}^1 - \hat{p}^0 > 0$. However, the t -statistic for a given treatment effect is given by

$$t \equiv \sqrt{N} \frac{\hat{p}^1 - \hat{p}^0}{\sqrt{\sigma_0^2 + \sigma_1^2}}. \quad (1)$$

Thus, decision rules that choose $a = 1$ if and only if $t > \underline{t} > 0$ are suboptimal. To see this, note that there always exists σ_0 large enough to cause the decision-maker to stick with $a = 0$ regardless of the estimated treatment effect.

¹Parameters \hat{p}^a and $\frac{\sigma_a^2}{N}$ could be derived from a standard Gaussian learning model.

Risk-aversion over policy outcomes. A natural supposition is that risk aversion may drive the reliance on hypothesis testing using t -statistics. However, this is not the case. To show this, we assume (w.l.o.g.) that $\sigma_0 < \sigma_1$, and consider a decision-maker who wants to solve $\max_{a \in \{0,1\}} \mathbb{E}[\Gamma(\bar{p}^a)]$ where Γ is thrice continuously differentiable and concave. For N large, the second order development of $\mathbb{E}[\Gamma(\bar{p}^a)]$ is

$$\mathbb{E}[\Gamma(\bar{p}^a)] = \Gamma(\hat{p}^a) - \frac{1}{2} \Gamma''(\hat{p}^a) \frac{\sigma_a^2}{N} \text{ implying sufficient statistic } \theta = 2N \frac{\Gamma(\hat{p}^1) - \Gamma(\hat{p}^0)}{\Gamma''(\hat{p}^1)\sigma_1^2 - \Gamma''(\hat{p}^0)\sigma_0^2}.$$

As $\sigma_0 < \sigma_1$, the decision-maker will choose policy $a = 1$ whenever $\theta > \gamma$, where $\gamma = \frac{1}{2} \frac{\Gamma'}{\Gamma''}$. In the special case where Γ is quadratic, θ reduces to

$$\theta = N \frac{\hat{p}^1 - \hat{p}^0}{\sigma_1^2 - \sigma_0^2}.$$

This differs significantly from a t -statistic: mean treatment effect $\hat{p}^1 - \hat{p}^0$ is scaled by the difference of variances, rather than the sum of standard deviations. Indeed, risk-aversion means that the decision-maker values certainty (a small variance in outcomes) as well as a higher mean outcome. Variance—rather than standard deviation—plays a role as Γ is smooth, and can be approximated by a 2nd order polynomial.

Reference-dependent preferences. The preceding discussion suggests that hypothesis testing can only be motivated by reference-dependent preferences that treat $a = 0$ and $a = 1$ asymmetrically (see Tetenov, 2012). Let $\Delta^a \equiv \bar{p}^a - \bar{p}^{1-a}$. Consider a decision-maker who seeks to solve

$$\max_{a \in \{0,1\}} \mathbb{E}[w_a(\Delta^a)], \tag{2}$$

where

$$\forall a \in \{0,1\}, \quad w_a(\Delta^a) = \Delta^a + \kappa_a \times \Delta^a \mathbf{1}_{\Delta^a < 0}$$

with $\kappa_1 > \kappa_0 \geq 0$.

Lemma A.1. *Consider a reference-dependent agent solving (2). The optimal-decision rule takes the form $t > t^*$, with $t^* > 0$.*

Proof. Let $\bar{t} \equiv \frac{\bar{p}^1 - \bar{p}^0}{\sqrt{\sigma_0^2 + \sigma_1^2}}$. As $\bar{p}^1 - \bar{p}^0 \sim \mathcal{N}(\hat{p}_1 - \hat{p}_0, \sqrt{\sigma_0^2 + \sigma_1^2})$, it follows that $\bar{t} \sim \mathcal{N}(t, 1)$. As both w_0 and w_1 are positively homogeneous of degree 1, the decision-maker chooses $a = 1$ if and only if:

$$\begin{aligned} \mathbb{E}[w_1(\Delta^1) - w_0(\Delta^0)] > 0 &\iff \mathbb{E}_{\bar{t}} \left[w_1 \left(\bar{t} \sqrt{\sigma_0^2 + \sigma_1^2} \right) - w_0 \left(-\bar{t} \sqrt{\sigma_0^2 + \sigma_1^2} \right) \middle| t \right] > 0 \\ &\iff \mathbb{E}_{\bar{t}} [w_1(\bar{t}) - w_0(-\bar{t}) | t] > 0 \\ &\iff t > t^* \end{aligned}$$

for some value of \underline{t} . Note that $w_1(\bar{t}) - w_0(-\bar{t}) = (2 + \kappa_0)\bar{t} + (\kappa_1 - \kappa_0)\bar{t}\mathbf{1}_{\bar{t} < 0}$. As $\kappa_1 > \kappa_0$ it follows that $w_1(\bar{t}) - w_0(-\bar{t})$ is concave in \bar{t} . This implies that $\mathbb{E}_{\bar{t}} [w_1(\bar{t}) - w_0(-\bar{t}) | t = 0] < 0$, so that $t^* > 0$. \square

B Proofs

Proof of Lemma: By the Minimax Theorem (Luenberger, 1969), the decision-maker's indirect utility from running experiment \mathcal{E} can be written as

$$\begin{aligned} V(\mathcal{E}) &\equiv \max_{\alpha \in \mathcal{A}} U(\mathcal{E}, \alpha) = \max_{\alpha \in \mathcal{A}} \min_{h \in H'} \mathbb{E}_{h, \mathcal{E}}[u(p, \alpha(e, y))] \\ &= \min_{h \in H'} \max_{\alpha \in \mathcal{A}} \mathbb{E}_{h, \mathcal{E}}[u(p, \alpha(e, y))]. \end{aligned}$$

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

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

where $v(h, e) \equiv \sum_{y \in \mathcal{Y}} \max_{a \in \{0,1\}} \mathbb{E}_{p \sim h} [\text{prob}(y|p, e) u(p, a)]$. As $v(h, e') = v(h, e) \equiv v(h, [e])$ 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 begin by showing that deterministic experiments are always weakly optimal for a Bayesian decision-maker. The decision-maker's indirect utility from running experiment \mathcal{E} can be written as

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

where $v(h_0, e)$ is the indirect utility from decision-making given realized experiment e : $v(h_0, e) \equiv \sum_{y \in \mathcal{Y}} \text{prob}(y|e) \max_{a \in \{0,1\}} \mathbb{E}_{p \sim h_0} [u(p, a)|e, y]$. Any deterministic experiment e^* solving $\max_{e \in E} v(h_0, e)$ is optimal. More strongly, \mathcal{E} solves (DP) if and only if $\text{support } \mathcal{E} \subset \arg\max_{e \in E} v(h_0, e)$.

To show that deterministic experiments are generically strictly optimal, we begin by showing that $\arg\max_{[e] \in [E]} v(h_0, [e])$ is generically a singleton for $\lambda = 1$. 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 . As E is finite, there exists $\eta > 0$ such that $v(h_0, [e]) < v(h_0, [e_0]) - \eta$ for all $[e] \neq [e_0]$. As $v(h, e)$ 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

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

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 . As $[E]$ 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 $v((1 - \theta)h_0 + \theta p, [e_0]) \neq v((1 - \theta)h_0 + \theta p, [e_1])$. 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.

We now consider the case of $\lambda < 1$. Given λ , h , and $[e]$, as the decision-maker's utility

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 (DP) is maximized at \mathcal{E} if and only if support $\mathcal{E} \subset [e_0]$. ■

Proof of Proposition 3: To establish point (i) and Corollary 1, we use the standard RCT $(\mathcal{E}_{\text{RCT}}, \alpha_{\text{RCT}})$. Losses $L(p)$ from first best, given state of the world p , are defined as

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

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

$$\begin{aligned} \operatorname{prob}_{p, \mathcal{E}_{\text{RCT}}}(\bar{y}^1 - \bar{y}^0 \leq 0) &= \operatorname{prob}_{p, \mathcal{E}_{\text{RCT}}}(\bar{y}^0 - \bar{y}^1 - (p^0 - p^1) \geq (p^1 - p^0)) \\ &\leq \exp\left(-\frac{2(p^1 - p^0)^2}{\frac{1}{(1-\pi)N} + \frac{1}{\pi N}}\right) \\ &= \exp(-2\pi(1-\pi)N(p^1 - p^0)^2) \\ &\leq \exp(-\pi N(p^1 - p^0)^2), \end{aligned}$$

²McDiarmid (1989): Let X_1, \dots, X_n be independent random variables, with X_k taking values in a set A_k for each k . Suppose that the (measurable) function $f : \times_k A_k \rightarrow \mathbb{R}$ satisfies $|f(x) - f(x')| \leq c_k$ whenever x and x' differ only in the k th coordinate. Then, for any $t > 0$, $\operatorname{prob}(f(X_1, \dots, X_n) - \mathbb{E}[f(X_1, \dots, X_n)] \geq t) \leq \exp(-2t^2 / \sum_k c_k^2)$.

where the last inequality follows from $2\pi(1 - \pi) \geq \underline{\pi} \equiv \min\{\pi, 1 - \pi\}$. For any $a > 0$, the mapping $x \mapsto x \exp(-ax^2)$ is ln-concave and maximized at $x = (2a)^{-1/2}$. This implies that

$$\max_{a \in \{0,1\}} p^a - \mathbb{E}_{p, \mathcal{E}_{\text{RCT}}} [p^{1_{\bar{y}^1 - \bar{y}^0 > 0}}] \leq \sqrt{\frac{\exp(-1)}{2\underline{\pi}N}} \leq \sqrt{\frac{\ln 2}{2\underline{\pi}N}}. \quad (3)$$

An analogous argument delivers (3) also for the case where $p^1 - p^0 \leq 0$. Hence, given any $h \in H$,

$$\mathbb{E}_h \left(\max_{a \in \{0,1\}} u(p, a) \right) - \mathbb{E}_{h, \mathcal{E}_{\text{RCT}}} [u(p, \alpha_{\text{RCT}}(e, y))] \leq \sqrt{\frac{\ln 2}{2\underline{\pi}N}}.$$

Setting $\pi = 1/2$ yields point (i) and the tightest bound.

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

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

■

Proof of Proposition 4: Consider an experiment $e^\dagger \in \operatorname{argmax}_{e \in \operatorname{support} \mathcal{E}_K} B(e)$. As the total number of subsets of $\{1, \dots, N\}$ is equal to 2^N , the number of experiments that assign treatment to πN participants out of N is necessarily less than or equal to 2^N . Hence the probability that the k th rerandomized trial selects e^\dagger is at least $\rho \equiv 1 - (1 - 2^{-N})^K$. For $K \geq 2^N$,

$$\rho \geq 1 - \exp(2^N \ln(1 - 2^{-N})) \sim 1 - \exp(-1) > 0.$$

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 $(\bar{y}_{0,k}, \bar{y}_{1,k})$ the sample average of outcomes by treatment group for experiment e_k , and by $(\bar{y}_0^*, \bar{y}_1^*)$ 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 $L(p) \equiv \max_{a \in \{0,1\}} p^a - \mathbb{E}_{p, \mathcal{E}_K} [\text{prob}(\mathbf{1}_{\bar{y}_1^* - \bar{y}_0^* > 0})]$. By symmetry, it suffices to treat the case where $p^1 - p^0 > 0$. In this case, we have

$$\begin{aligned} L(p) &= (p^1 - p^0) \text{prob}_{p, \mathcal{E}_K} (\bar{y}_1^* - \bar{y}_0^* \leq 0) \\ &\leq (p^1 - p^0) \text{prob}_{p, \mathcal{E}_K} \left(\min_{k \in \{1, \dots, K\}} \bar{y}_{1,k} - \bar{y}_{0,k} \leq 0 \right) \\ &\leq (p^1 - p^0) \min \left\{ 1, \sum_{k=1}^K \text{prob}_{p, \mathcal{E}_K} (\bar{y}_{1,k} - \bar{y}_{0,k} \leq 0) \right\}. \end{aligned}$$

As in the proof of Proposition 3, by applying McDiarmid's inequality to $f_k(y) \equiv \bar{y}_{0,k} - \bar{y}_{1,k}$, we obtain $\text{prob}_{p, \mathcal{E}_K} (\bar{y}_{1,k} - \bar{y}_{0,k} \leq 0) \leq \exp(-\underline{\pi} N (p^1 - p^0)^2)$, where $\underline{\pi} \equiv \min\{\pi, 1 - \pi\}$.

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

$$L(p) \leq \begin{cases} p^1 - p^0 & \text{if } p^1 - p^0 < \sqrt{\frac{\ln K}{\pi N}}, \\ K(p^1 - p^0) \exp(-\pi N(p^1 - p^0)^2) & \text{if } p^1 - p^0 \geq \sqrt{\frac{\ln K}{\pi N}}. \end{cases} \quad (4)$$

The mapping $x \mapsto x \exp(-\pi N x^2)$ is ln-concave and maximized at $x = \sqrt{\frac{1}{2\pi N}}$. As $K \geq 2$, we have $\sqrt{\frac{\ln K}{\pi N}} > \sqrt{\frac{1}{2\pi N}}$, which implies that both terms on the right-hand side of (4) are maximized at $p^1 - p^0 = \sqrt{\frac{\ln K}{\pi N}}$. This implies that indeed $L(p) \leq \sqrt{\frac{\ln K}{\pi N}}$. Identical reasoning applies in the case where $p^1 - p^0 < 0$. ■

Proof of Corollary 2: Consider a balance function minimized by assignments $e \in E_{\dagger}$. Let L_K denote the loss in efficiency from a K -rerandomized trial, and $L_{E_{\dagger}}$ the loss in efficiency from experiment $\mathcal{E}_{E_{\dagger}}$. The likelihood that the assignment e drawn from \mathcal{E}_K belongs to E_{\dagger} is $1 - (1 - p_{E_{\dagger}})^K$. As $L_K \geq (1 - (1 - p_{E_{\dagger}})^K) L_{E_{\dagger}}$, it follows from Proposition 5 that

$$L_{E_{\dagger}} \leq \frac{1}{1 - (1 - p_{E_{\dagger}})^K} \sqrt{\frac{\ln K}{\pi N}}.$$
■

Proof of Proposition 6: Point (i) follows from a reasoning similar to that of Proposition 2. For $\lambda = 1$, given an experiment \mathcal{E} , the decision-maker's indirect utility is

$$\max_{\alpha, \mathcal{E}} \mathbb{E}_{h_0}[w_{\alpha}(\Delta^{\alpha})] = \sum_{e \in E} \mathcal{E}(e) w(h_0, e),$$

where $w(h_0, e) \equiv \sum_{y \in \mathcal{Y}} \text{prob}(y|e) \max_{a \in \{0,1\}} \mathbb{E}_{p \sim h_0}[w_a(\Delta^a)|e, y]$. Hence, an experiment \mathcal{E} is optimal if and only if support $\mathcal{E} \subset \arg \max_e w(h_0, e)$.

To conclude the proof of point (i), it is sufficient to show that $\arg \max_{[e] \in [E]} w(h_0, [e])$ is a singleton. A reasoning identical to that of Proposition 2, with $w(h_0, [e])$ in place of $v(h_0, [e])$, shows that this is the case for generically every prior h_0 . The existence of an appropriate threshold $\underline{\lambda} < 1$ follows from the fact that $w(h_0, [e])$ is Lipschitz continuous in h_0 , and there are finitely many possible experimental assignments.

We now turn to point (ii). We know from Proposition 7 that there exist randomized experiments leading to optimal decisions up to a penalty of order $O(1/\sqrt{N})$. This implies that the decision-maker can guarantee herself a payoff greater than $-O(1/\sqrt{N})$. Consider a deterministic experiment e . For $d \in \mathbb{R}$, let the state $p(d)$ such that

$$\begin{aligned} p_x^0 &= \frac{1}{2} + d, & p_x^1 &= \frac{1}{2} & \text{if } \tau_x &= 1; \\ p_x^0 &= \frac{1}{2}, & p_x^1 &= \frac{1}{2} - d & \text{if } \tau_x &= 0. \end{aligned}$$

Consider the prior h_e that puts probability 0.5 on $p(d = \nu)$ and 0.5 on $p(d = -\nu)$ for $\nu \in (0, 1/2)$. By construction the information generated by the experiment is independent of whether $d = \nu$ or $d = -\nu$. In addition, $\Delta^1 = \bar{p}^1 - \bar{p}^0 = -d$. Hence, regardless of the action a taken by the decision-maker, there is probability 0.5 that $\Delta^a = -\nu$ and probability 0.5 that $\Delta^a = +\nu$. As w_a is locally strictly concave around 0, it follows that expected payoffs from running a deterministic experiments are bounded away below 0.

This implies that for N large enough, randomized experiments are strictly optimal. ■

Proof of Proposition 7: The proof is closely related to that of Proposition 5. Consider first the case where $\Delta \bar{p} > 0$. Then the efficiency loss compared to first-best is equal to

$$L = \mathbb{E}_h[w_1(\Delta^1) - w_{\alpha_{RCT}}(\Delta^1) | \Delta^1 > 0].$$

As functions w_a are Lipschitz continuous and $w_a(0) = 0$, there exists $M > 0$ such that

$$L \leq M\mathbb{E}_h[(1 - \alpha_{RCT})\Delta^1 | \Delta^1 > 0].$$

In turn, if $\Delta\bar{p} < 0$, there exists M such that the efficiency loss satisfies $L \leq M\mathbb{E}_h[-\alpha_{RCT}\Delta^1 | \Delta^1 < 0]$.

The proof of Proposition 5 implies that $\mathbb{E}_h[(1 - \alpha_{RCT})\Delta^1 | \Delta^1 > 0]$ and $\mathbb{E}_h[-\alpha_{RCT}\Delta^1 | \Delta^1 < 0]$ are bounded above by $\sqrt{\frac{\log K}{\pi N}}$. This proves Proposition 7. \blacksquare

C Simulations

C.1 Optimal Experimentation

A simple example helps clarify the use of rerandomization as a way to pick a trade-off between subjectively optimal experiments and the robustness of RCTs. We consider an environment where $N = |X| = 2$, $X = \{0, 1\}$, and the experimenter is given a sample of participants with $(x_1, x_2) = (1, 0)$.

We parameterize the set H of admissible priors over the set P of states of the world as follows. First, for the expected probability of success given policy $a \in \{0, 1\}$, $p^a = qp_1^a + (1 - q)(1 - p_0^a)$, we consider priors such that, with equal probability, the pair (p^0, p^1) describes either a “high probability of success” world with $(\min_{a \in \{0, 1\}} p^a, \max_{a \in \{0, 1\}} p^a) = (1/2, 3/4)$ or a “low probability of success” world with $(\min_{a \in \{0, 1\}} p^a, \max_{a \in \{0, 1\}} p^a) = (1/4, 1/2)$.

Next, given p^a , to further constrain the support of priors in H , we construct a grid over the set of values of the conditional success probabilities p_0^a and p_1^a consistent with p^a . Note that as $p_0^a = (p^a - qp_1^a)/(1 - q)$, $p_0^a \geq 0$ implies that $p_1^a \leq h^a \equiv \min\{1, p^a/q\}$, while $p_0^a \leq 1$ implies that $p_1^a \geq l^a \equiv \max\{0, (q - 1 + p^a)/q\}$. Letting $m^a \equiv l^a/2 + h^a/2$, we restrict priors in

H to assign positive probability only to values of p_1^a defining an equally-spaced, 10-segment grid from $l^a/2 + m^a/2$ to $m^a/2 + h^a/2$.

With only two participants, we consider experiments with $\tau_1 = 1 - \tau_2$, and we solve numerically for the optimal probability with which participant 1 should be treated to maximize the adversarial component of the experimenter's payoff as defined in (DP) with policy chosen according to α_{RCT} . Figure C.1 shows how the optimal experiments vary with the value of $q \in [0, 1]$, the share of agents with $x = 1$ in the population. Relative to a standard RCT with probability $\pi = 1/2$ of treating each participant, the experimenter only experiences gains from running an optimal experiment for intermediate values of q . In this range, rerandomization offers a compromise between optimality and the robustness of the RCT. However, if q is close to zero or one, the advantage of rerandomization over the standard RCT disappears.




Figure C.1: Range of optimal experiments as a function of the environment.

C.2 The Trade-Offs of Rerandomization

We provide two numerical explorations of our results. We begin by considering a well-behaved case in which treatment effects are continuous with respect to a small number of

underlying characteristics, so that there is limited tension between balance and robustness. We then turn to a much more discontinuous setting designed to oppose the desire to balance and robust policy making. The message from both simulations is clear: rerandomization increases balance with very little to no increase in mistaken decisions.

C.2.1 Smooth Priors

We consider the following environment. 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 :

$$\text{prob}(Y_i = 1|x) = \frac{\exp(\mu \cdot x)}{1 + \exp(\mu \cdot x)}.$$

Parameter μ is drawn according to a five-dimensional truncated normal: $\mu \sim \prod_{k=1}^5 \mathcal{N}(0, 1)_{|[-2, 2]}$.

We denote by τ^* and α the Bayes optimal assignment of treatment and policy choice under this model.

We report balance—captured by the negative of the L_2 norm between mean characteristics across treatment and control—as well as several efficiency losses of interest (see Figure C.2):


- *Bayes Loss given Bayes Optimal Assignment*

$$\mathbb{E}_{\mu, x, \tau^*} \left[\max_{a \in \{0, 1\}} u(p, a) - u(p, \alpha) \right]; \quad (5)$$

- *Loss under worst prior given Bayes optimal assignment*

$$\max_{\mu} \mathbb{E}_{x, \tau^*} \left[\max_{a \in \{0, 1\}} u(p, a) - u(p, \alpha) \right]; \quad (6)$$

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




- *Loss under worst prior, and worst assignment τ*

$$\max_{\mu} \mathbb{E}_x \max_{\tau} \mathbb{E} \left[\max_{a \in \{0,1\}} u(p, a) - u(p, \alpha) \right]. \quad (7)$$

The ex-ante Bayes expected loss (5) is essentially identical under randomization and rerandomization. Loss measure (6) chooses the prior that maximizes the error rate given the experimental strategy \mathcal{E} of the experimenter. While this is substantially higher than the Bayes expected loss — as expected — it is not substantially different between randomization and rerandomization. Finally, loss measure (7) 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

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



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 C.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.


C.2.2 The Case of Non-Smooth Priors

We now consider an environment designed to create a tension between balance and robustness. Also, we pick assignment τ using balance objective $B(e) = -\|\bar{x}^1 - \bar{x}^0\|_2$. Policy is chosen according to $\alpha(e, y) \equiv \arg\max_{a \in \{0,1\}} \bar{y}^a - \bar{y}^{1-a}$.

The environment involves a single covariate $x \in X = \{1, 2, \dots, 10,000\}$. 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 $n \in \{1, 2, \dots, 5,000\}$,

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

Figure C.4: Rerandomization substantially increase balance with no cost to robustness.



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 numerous even observations, which are not an appropriate comparison group.³

Figure C.4 examines the error rates and balance of randomization and rerandomization. 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

³Indeed, using pairwise matching to assign treatment and control status increases inferential errors, but does so equally for randomization and rerandomization

though we only re-draw the experimental allocation 100 times.

C.2.3 (Re-)Randomization Inference

As noted in the text, if $B(e)$ is pre-specified, randomization inference can be used to calculate standard errors (Fisher, 1935; Young, 2016). Here, we illustrate this process using the simulations, and balance functions, from the prior subsections.

In particular, randomization inference uses the following procedure:

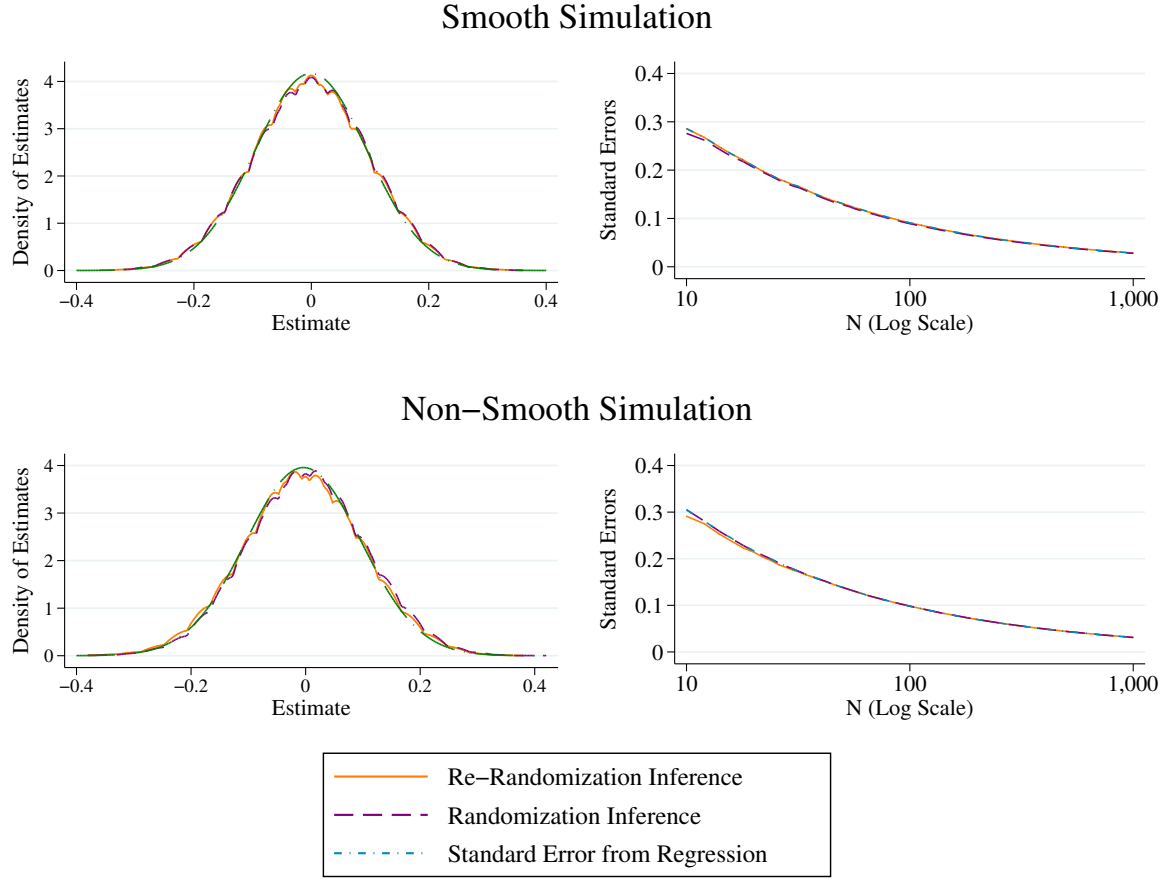
1. Randomize assignment $(\tau_i)_{i \in \{1, \dots, N\}}$
2. Generate data $(y_i^{\tau_i})_{i \in \{1, \dots, N\}}$
3. Calculate treatment effect $\Delta \hat{y} \equiv \hat{y}^1 - \hat{y}^0$
4. Re-draw experimental assignment using the same randomization procedure $(\tau_i^j)_{i \in \{1, \dots, N\}}$
5. Calculate treatment effect *as if* τ^j had been implemented $\Delta^j \hat{y} \equiv \hat{y}^{\tau_i^j=1} - \hat{y}^{\tau_i^j=0}$
6. Return to Step 4, repeat until $j = J$ (usually 10,000).

The simulated distribution of $\Delta^j \hat{y}$ is the distribution of treatment effects that would be observed if $p_x^1 = p_x^0, \forall x \in X$. This can be used to calculate p -values, or infer standard errors.

Clearly, knowing $B(e)$, and the number of rerandomizations, K , is critical to properly implementing Step 4. If this is the case, we say that the experimenter can conduct *rerandomization inference*.

In Figure C.5, we compare rerandomization inference and the standard errors that would come from a regression of outcomes on treatment status. Finally, we also include *naïve randomization inference*, where the experimenter does not know $B(e)$, and consequently draws a single randomization to calculate each $\Delta^j \hat{y}$.

Figure C.5: Comparing rerandomization inference with other inference strategies.



The right-hand-side panels of the figure show the distribution of $\Delta^j \hat{y}$ under each technique. The distribution shown for the regression comes from a simulated normal with standard deviation equal to the standard error from the regression coefficient. The left-hand-side panels show how standard errors from each source change as the sample grows larger. Note that we set $K = \min\{N, 100\}$ for both the experimental simulation and rerandomization inference.

The results under all schemes are substantially the same. This may be due, in part, to the fact that outcomes are bounded in our simulations. However, it is certainly the case, as

our theoretical results show, that rerandomization has a very small impact on the robustness of decision-making.