

A Theory of Experimenters: Robustness, Randomization, and Balance[†]

By ABHIJIT V. BANERJEE, SYLVAIN CHASSANG, SERGIO MONTERO,
AND ERIK SNOWBERG*

This paper studies the problem of experiment design by an ambiguity-averse decision-maker who trades off subjective expected performance against robust performance guarantees. This framework accounts for real-world experimenters' preference for randomization. It also clarifies the circumstances in which randomization is optimal: when the available sample size is large and robustness is an important concern. We apply our model to shed light on the practice of rerandomization, used to improve balance across treatment and control groups. We show that rerandomization creates a trade-off between subjective performance and robust performance guarantees. However, robust performance guarantees diminish very slowly with the number of rerandomizations. This suggests that moderate levels of rerandomization usefully expand the set of acceptable compromises between subjective performance and robustness. Targeting a fixed quantile of balance is safer than targeting an absolute balance objective. (JEL C90, D81)

The proliferation of experiments in academia, business, and public policy has been accompanied by spirited debates about best practices for experiment design and the 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 (Duflo, Glennerster, and Kremer 2008; Bruhn and McKenzie 2009; Deaton 2010; Imbens 2010; Humphreys, de la Sierra, and van der Windt 2013; Olken 2015; Athey and Imbens 2017; Benjamin et al. 2018). At the heart of these debates are different, usually implicit, models of both

*Banerjee: Department of Economics, Massachusetts Institute of Technology (email: banerjee@mit.edu); Chassang: Department of Economics, New York University (email: chassang@nyu.edu); Montero: Department of Political Science, University of Rochester (email: smontero@rochester.edu); Snowberg: Vancouver School of Economics, University of British Columbia (email: snowberg@mail.ubc.ca). Jeffrey Ely was the coeditor for this article. We are grateful to David Ahn, Angus Deaton, Pascaline Dupas, Guido Imbens, Charles Manski, Pablo Montagnes, Marco Ottaviani, Bruno Strulovici, Alexey Tetenov, Duncan Thomas, Chris Udry, several helpful referees, 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), Penn State, the SISL Conference at Caltech (2017), UBC, UC Berkeley, the UCL/Cemmap workshop on Econometrics for Public Policy (2016), and UPenn for several helpful discussions. Chassang and Snowberg gratefully acknowledge the support of NSF grant SES-1156154.

[†]Go to <https://doi.org/10.1257/aer.20171634> to visit the article page for additional materials and author disclosure statements.

knowledge generation and how people interpret experimental evidence. Moreover, classical models of experimentation, in which decision-makers are subjective expected utility-maximizers, cannot explain the strong preference for randomized controlled trials (RCTs) expressed by experimenters (Savage 1954, Kasy 2016).¹

This paper seeks to help clarify these debates. We propose an instrumental model of experimentation in which a decision-maker collects experimental data in order to improve decisions under uncertainty. A key step, which allows us to account for randomization, is to model the decision-maker as ambiguity-averse. Specifically, we use the maxmin framework of Gilboa and Schmeidler (1989), reformulated as the problem of a Bayesian decision-maker seeking to satisfy an adversarial audience. Examples of adversarial audiences abound: the Food and Drug Administration for drug trials, seminar audiences and journal referees for research papers, and governments or NGOs for public policy proposals.² This intuitive reinterpretation of Gilboa and Schmeidler (1989) (also known as the ε -contamination model; Huber 1964) facilitates meaningful comparative statics.

The paper reports two main sets of results. The first set shows that randomized experiments can be strictly optimal for a decision-maker facing an adversarial audience, and clarifies the circumstances in which this is the case. If the decision-maker cares sufficiently about the adversarial audience, and the sample size is sufficiently large, then it is strictly optimal for the decision-maker to use a randomized experiment. Specifically, all deterministic experiments are bounded away from the first-best, while a standard RCT approaches first-best up to a loss that is vanishing in the sample size. 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 own subjective expected utility.

This set of results accords with the observed heterogeneity in experimental practice. Randomized experiments tend to be used by decision-makers who put a high value on convincing an adversarial audience (for example, scientists and pharmaceutical companies) or when the decision-maker can afford large samples (for example, A/B testing in online marketing). When data points are few and the decision-maker puts little weight on satisfying an adversarial audience (for example, the CEO of a privately held firm deciding whether to adopt a new production technology) optimal experiments are deterministic and optimize the subjective informational value of each data point.

Our second set of results applies our model to an open issue in experiment design: whether to rerandomize to improve covariate balance between treatment and control groups. Rerandomization draws multiple treatment assignments, then chooses the one that maximizes a prespecified balance objective. 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

¹RCTs are mixed strategies over experimental assignments. As a result they can never be strictly optimal for a subjective expected utility-maximizer. For seminal contributions to the economic literature on experimentation and information acquisition, see Rothschild (1974), Grossman and Stiglitz (1980), Aghion et al. (1991), Bergemann and Välimäki (1996, 2002, 2006), and Persico (2000).

²The audience may also be seen as a stand-in for the decision-maker's self-doubt.

improve balance, researchers are concerned that it may adversely affect the reliability of findings (Bruhn and McKenzie 2009).

We show that the trade-offs at the heart of rerandomization are concisely captured in our framework. Successive rerandomizations improve balance, which can be expressed through the subjective expected utility component of preferences. However, rerandomization reduces the robust performance guarantees offered by RCTs. In the extreme case where the allocation is rerandomized until perfect balance is achieved, the allocation is effectively deterministic and worst-case performance guarantees are bounded away from first-best, even for large samples. In contrast, losses against the first-best are vanishing in sample size N provided the number of rerandomizations is polynomial in N . We clarify the potential robustness costs of alternative procedures used to achieve balance: setting a balance target and rerandomizing until it is met (Morgan and Rubin 2012); or equivalently, selecting an assignment uniformly from a constrained set of assignments achieving a prespecified balance objective.

Our framework builds on a long line of work in statistics, starting with Wald (1950). Variants of Wald's framework have been used in economics and econometrics to study questions of identification and model uncertainty (Gilboa and Schmeidler 1989; Hansen and Sargent 2001; Manski 2004, 2009; Marinacci 2015; Kitagawa and Tetenov 2018).³ Different approaches to model uncertainty, specifically ambiguity aversion or regret aversion, remain an ongoing subject of debate in this literature. Ambiguity aversion has more attractive normative foundations (see, for example, Marinacci 2015), while regret aversion is more directly prescriptive. We state our main results under ambiguity aversion, but show that they also hold when using regret aversion. The latter framework allows us to dispense with certain technical assumptions, as well as rationalize the practice of null-hypothesis statistical testing (NHST, Tetenov 2012). In a related paper (Banerjee, Chassang, and Snowberg 2017), three of the authors provide a nontechnical discussion of our results on randomization, and a limited discussion of our results on rerandomization.

The paper is structured as follows. Section I introduces our framework. Section II delineates the forces that determine whether running a randomized or deterministic experiment is optimal. Section III studies the trade-offs involved in rerandomization. Section IV shows that our results extend to reference-dependent preferences, better suited to explain the use of NHST in decision-making. Section V contains several discussions: of other possible rationalizations of randomization, of the positive implications of our theory, of practical implications for rerandomization, and of possible directions for future research. The Appendix extends our analysis to dynamic settings where participants arrive over time, and experimenters must choose treatment assignments in real time, without knowing what the ultimate sample of covariates will be.⁴

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.

⁴Online Appendix Section A clarifies that a traditional balance objective, minimizing the Mahalanobis distance between mean covariates across treatment and control samples, coincides with subjective expected utility maximization under specific priors. Online Appendix Section B clarifies how regret aversion can rationalize the use of NHST. Proofs are contained in online Appendix Section C. Online Appendix Section D presents simulations.

I. Model

In this section, we first lay out the problem of experiment design along lines similar to Banerjee, Chassang, and Snowberg (2017). We then specify the preferences and beliefs of the decision-maker.

A. A Framework for Studying Experiment Design

Decisions and Payoffs.—A decision-maker chooses whether to provide a treatment $\tau \in \{0, 1\}$ to a population of heterogeneous 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. Potential outcomes for individual i with treatment status τ are random variables $y_i^\tau \in \{0, 1\}$; $y = 1$ is referred to as a success. Each individual has observable covariates $x_i \in X$ that affect the distribution of outcomes y_i . Covariates x_i are uniformly distributed in X .

The probability of success, given treatment τ and covariate x , is denoted by $p_x^\tau \equiv \Pr(y^\tau = 1|x)$. 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 covariate x . Outcomes y_i^τ for different individuals are drawn independently given state of the world 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

$$u(p, a) \equiv \mathbb{E}_p y^a = \frac{1}{|X|} \sum_{x \in X} p_x^a.$$

Although covariates x are observable, our framework is consistent with unobservable characteristics. Denoting unobserved characteristics by z , and given a joint distribution F over (x, z) we have $p_x^\tau = \int p_{x,z}^\tau dF(z|x)$. When x and z are embedded in \mathbb{R}^k , our framework captures traditional concerns over omitted variables by allowing the mapping $x \mapsto p_x^\tau$ from observed covariates to outcomes to be discontinuous. Even if $p_{x,z}^\tau$ is continuous in (x, z) , omitted variables cause $x \mapsto p_x^\tau$ to become discontinuous when the distribution of z conditional on x changes rapidly. That is, when individuals with similar values of x may have very different values of z .⁵

Experiments and Strategies.—To maximize her odds of making the correct policy choice, the decision-maker can run an experiment on N participants, which are a representative sample of a broader population. The decision-maker chooses a treatment status τ_i for each participant i , observes outcomes $y_i^{\tau_i}$, and then makes a policy decision a that applies to the entire broader population.

We simplify away the problem of representative sampling by taking as given the covariates $(x_i)_{i \in \{1, \dots, N\}}$ of experimental participants and assuming that they are *exactly* representative of the underlying set of types. That is, we assume that $N = |X|$ and $\{x_i\}_{i \in \{1, \dots, N\}} = X$. This means that random sampling is not needed to ensure that experimental participants are representative of the

⁵That is, $F(z|x)$ is not continuous in x .

overall population. This assumption is consistent with the Neyman-Rubin “potential outcomes” framework (Neyman 1923, Rubin 1974). Each person i is unique and cannot be observed in both the treated and untreated state. This assumption simplifies exposition and focuses on randomness in treatment assignment as the only random component of experiment design. Under this exact sampling assumption, the broader population can be thought of as identical copies of the experimental sample.⁶

Given covariates $(x_i)_{i \in \{1, \dots, N\}}$, an experimental assignment is a tuple $e = (\tau_i)_{i \in \{1, \dots, N\}} \in \{0, 1\}^N \equiv E$. 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) . We say that a covariate and treatment pair $(x, \tau) \in X \times \{0, 1\}$ is *sampled* by experiment e , denoted by $x, \tau \in e$, if and only if there exists $i \in \{1, \dots, N\}$ such that $(x_i, \tau_i) = (x, \tau)$.

The decision-maker’s strategy consists of both an experimental design $\mathcal{E} \in \Delta(E)$, which is a mixed strategy over experimental assignments e , 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 α .⁷

The Standard RCT.—For simplicity, we assume that $N/2$ is an integer throughout the paper. An experiment design of interest is the standard RCT, assigning 50 percent of participants to treatment $\tau = 1$.⁸ It corresponds to the strategy $(\mathcal{E}_{RCT}, \alpha_{RCT})$:

- \mathcal{E}_{RCT} draws an exchangeable profile $(\tau_i)_{i \in \{1, \dots, N\}} \in \{0, 1\}^N$ of treatment assignments such that $\sum_{i=1}^N \tau_i = N/2$.
- Policy a is chosen according to the *empirical success* rule: $\alpha_{RCT}(e, y) \equiv \mathbf{1}_{\bar{y}^1 \geq \bar{y}^0}$, where $\bar{y}^\tau \equiv (2/N) \sum_{i=1}^N y_i \mathbf{1}_{\tau_i = \tau}$ is the mean outcome for participants with treatment status τ .

Note that policy choice α_{RCT} is deterministic conditional on experimental outcomes: performance guarantees for RCTs established here stem from random assignment, not random policymaking.

B. Preferences and Beliefs

The experiment designer’s preferences and beliefs are at the heart of our analysis. As noted in the introduction, subjective expected utility maximization does not yield strict preferences for randomization. Ambiguity aversion does, under conditions explored below.

⁶Our results extend when $N < |X|$ and experimental participants are randomly sampled from X according to the underlying distribution of covariates.

⁷Following standard notation, for any $\sigma \in \Delta(\{0, 1\})$, $u(p, \sigma) \equiv \mathbb{E}_{a \sim \sigma}[u(p, a)]$.

⁸The analysis extends to experiments assigning any fixed share $\pi \in (0, 1)$ to treatment. Bounds use the reduced sample $N' = 2\min\{\pi, 1 - \pi\}N$ obtained by dropping excess data points in the larger subsample.

Preferences.—We consider an ambiguity-averse decision-maker that chooses strategy (\mathcal{E}, α) to maximize

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

where H is a set of priors $h \in \Delta(P)$ over states $p \in P$. These preferences coincide with the standard maxmin model of Gilboa and Schmeidler (1989): the decision-maker’s objective can be expressed as $\min_{h \in H_0} \mathbb{E}_{h, \mathcal{E}}[u(p, \alpha(e, y))]$, with $H_0 \equiv \{\lambda h_0 + (1 - \lambda)h' | h' \in H\}$. This parameterization, known as the ε -contamination model of Huber (1964, 2010), is popular in the robust statistics literature, with λ serving as a tuning parameter. It lets us emphasize subjective expected utility maximization as a special case ($\lambda = 1$), and facilitates comparative statics.

Representation (DP) of the standard maxmin model admits an intuitive interpretation: a subjective-expected-utility-maximizing decision-maker, with prior h_0 , faces an adversarial audience of Bayesian stakeholders with noncommon priors $h \in H$. Note that heterogeneous priors within the audience are essential: if the audience entertained a common prior h , then the decision problem would reduce to subjective expected utility maximization for the mixed prior $\lambda h_0 + (1 - \lambda)h$. If $\lambda = 1$, the decision-maker does not care about the audience, and we recover the standard subjective expected utility, or *Bayesian* model. Throughout the paper, we interpret results in terms of an adversarial audience, but it is equally appropriate to view decision problem (DP) as a model of the experimenter’s *internal* doubt.

Beliefs.—Our main assumption concerns the set of priors H entertained by audience members. We assume that for any given experiment e , there exists a welfare-minimizing prior h such that knowing, without error, the success rates $(p_x^\tau)_{x, \tau \in e}$ for pairs (x, τ) sampled by experiment e is not sufficient for efficient decision-making. Denoting the expected probability of success given policy $a \in \{0, 1\}$ by $p^a \equiv (1/|X|) \sum_{x \in X} p_x^a$, the assumption is as follows.

ASSUMPTION 1 (Limited Extrapolation): *There exists $\xi > 0$ such that, for all $e \in E$, there exists a prior $h \in \operatorname{argmin}_{h \in H} \mathbb{E}_h(\max_{a \in \{0, 1\}} p^a)$, such that*

$$\mathbb{E}_h \left[\max_{a \in \{0, 1\}} p^a - \max_{a \in \{0, 1\}} \mathbb{E}_h [p^a | (p_x^\tau)_{x, \tau \in e}] \right] \geq \xi.$$

Note that, for any experiment e , $(p_x^\tau)_{x, \tau \in e}$ is an upper bound, in the sense of Blackwell (1953), of the information generated by experiment e . The term $\mathbb{E}_h [p^a | (p_x^\tau)_{x, \tau \in e}]$ is the expected payoff of action a from the perspective of a decision-maker with initial prior h after observing the true success rate p_x^τ for each pair (x, τ) sampled by the experiment. For any deterministic experiment e , the limited extrapolation assumption implies there exists a pernicious prior h such that data from sampled pairs $(x, \tau) \in e$ do not permit extrapolation of the optimal policy for the entire population.

A key implication of Assumption 1 is that audience members cannot be arbitrarily pessimistic. They cannot be sure that both treatment and control would result in certain failure. If this were the case, nature would minimize the decision-maker’s welfare using the prior $p^1 = p^0 = 0$, which makes experimentation worthless.

It is useful to point out two specific features of our modeling exercise. The first is that any comparative static exercise in which the sample size N grows, while Assumption 1 is maintained for $\xi > 0$ fixed, also requires type space X to grow with N . If a finite type space X was held fixed while N grew, then all pairs (x, τ) could be sampled by a deterministic experiment when N is sufficiently large. The second feature is that Assumption 1 becomes unnecessary in an alternative model in which the decision-maker exhibits regret aversion, rather than ambiguity aversion, as described in Section IV. A key advantage of regret aversion is that the priors that minimize the decision-maker's welfare must still allow for the possibility of learning. Otherwise there is nothing to regret. The main disadvantage of regret aversion is that it may lead to intransitivity for some choice problems (Marinacci 2015).

The following lemma establishes that the set of priors satisfying Assumption 1 is nonempty. The proof is by construction: it contains an example that is useful for understanding the content of the limited extrapolation assumption.

LEMMA 1 (Existence and Example): *There exists a set of priors H that satisfies limited extrapolation for parameter $\xi = 1/8$.*

PROOF:

Consider all pairs (X', τ) , with $X' \subset X$ satisfying $|X'| = |X|/2$, and $\tau \in \{0, 1\}$. Such a pair specifies a set of covariates and a single treatment status. For each pair (X', τ) , consider the prior $h_{X', \tau}$ such that if $x \notin X'$ or $\tilde{\tau} \neq \tau$ then $p_x^{\tilde{\tau}} = 1/2$. If $x \in X'$ and $\tilde{\tau} = \tau$, then the prior over this pair is given by $p_x^{\tilde{\tau}} = (1 + \sigma)/2$ where σ takes values 1 and -1 with equal probability. Note that σ takes a single value for all such pairs $(x, \tilde{\tau})$. Set H is the set of such priors $h_{X', \tau}$.

Set H satisfies limited extrapolation. First, the expected value $\mathbb{E}_h(\max_{a \in \{0, 1\}} p^a)$ is constant and equal to $5/8$ for all $h \in H$. This implies that $\operatorname{argmin}_{h \in H} \mathbb{E}_h(\max_{a \in \{0, 1\}} p^a) = H$. Next, for any deterministic experiment e there exists a prior $h \in H$ such that decision-making is bounded away from first-best. As $N = |X|$, there must exist X' and $\tau \in \{0, 1\}$ such that no pair in $X' \times \{\tau\}$ is sampled by deterministic experiment e . Thus, under the corresponding belief $h_{X', \tau}$, data generated by experiment e carries no information about the sign of σ . As a result, for all $\tau \in \{0, 1\}$, $\max_{a \in \{0, 1\}} p^a - p^\tau = 1/4$ with probability $1/2$. Thus, limited extrapolation holds for $\xi = 1/8$.⁹ ■

Two points raised by the proof are worth highlighting. First, in order to ensure that Assumption 1 is nonempty, $|X|$ needs to be larger than $N/2$: we assume $|X| = N$. Second, pernicious priors are such that nonsampled pairs (x, τ) create significant residual uncertainty about the correct treatment. As deterministic experiment e provides no information about these pairs, this uncertainty is not resolved by the experiment.

⁹Note that H in this example will continue to satisfy Assumption 1 even if it is expanded by adding other priors as long as $\mathbb{E}_{h'}(\max_{a \in \{0, 1\}} p^a) \geq 5/8$ for each added prior h' .

II. Optimal Design and Randomization

This section shows that both deterministic and randomized experiments can be strictly optimal for a decision-maker solving (DP). Which type of experiment is optimal depends on the sample size N and the weight λ given to subjective performance $\mathbb{E}_{h_0} u(a, p)$, versus worst-case performance guarantees $\min_{h \in H} \mathbb{E}_h u(a, p)$.

A. 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 h_0 under which all data points are valuable, deterministic experiments are *strictly* optimal. We use the topological version of genericity, that is, open, dense sets, under the total variation distance over distributions, $d(h, h') \equiv \sup_{A \text{ meas.}} |h(A) - h'(A)|$, sometimes described as *the* statistical distance.

We assume that the decision-maker values every data point. Formally, for any deterministic experiment $e \in E$, let e_{-i} denote the experiment in which sample point (x_i, τ_i) has been removed. Let $E_{-1} = \{e_{-i} | i \in \{1, \dots, N\}, e \in E\}$ denote the set of experiments with samples of length $N - 1$. We denote by \mathcal{A}_{-1} the set of policy rules mapping an experiment $e \in E_{-1}$ and outcomes $y \in \{0, 1\}^{N-1}$ to $\Delta(\{0, 1\})$.

DEFINITION 1: *We say that all data points are valuable under prior h_0 whenever*

$$\max_{\alpha \in \mathcal{A}} \max_{e \in E} \mathbb{E}_{h_0} [u(p, \alpha(e, y)) | e] > \max_{\alpha \in \mathcal{A}_{-1}} \max_{e \in E_{-1}} \mathbb{E}_{h_0} [u(p, \alpha(e, y)) | e].$$

That is, when all data points are valuable under prior h_0 , reducing the sample size reduces the quality of decision-making. This holds automatically if there is some cost of acquiring data points and the sample size N is chosen optimally by the decision-maker. However, as N is exogenously given here, we assume all data points are valuable. This rules out trivial indifferences between experiments.

PROPOSITION 1 (Near-Bayesians Do Not Randomize): *For every prior h_0 , if $\lambda = 1$, then there exists a deterministic experiment e^* solving (DP).*

For generically every prior h_0 such that all data points are valuable, there exists $\underline{\lambda} \in (0, 1)$ and a deterministic experiment e^ such that, for all $\lambda > \underline{\lambda}$, e^* is uniquely optimal.*

That is, when all data points are valuable, a deterministic experiment is generically strictly optimal when $\lambda = 1$, and also when the decision-maker puts a small, but nonzero, weight $1 - \lambda > 0$ on the preferences of her audience.¹⁰ In recent work, Kasy (2016) uses a similar result to argue that RCTs may be suboptimal. Instead, we

¹⁰Here genericity applies relative to the set of priors under which all data points are valuable. Note, however, that this is an open set under the statistical distance. Hence, the set of priors satisfying the second half of Proposition 1 is in fact open within the set of all priors.

believe that Proposition 1 shows the limits of subjective expected utility maximization as a positive model of experimenters.

B. Adversarial Experimentation

We now examine the case where the experimenter cares about her audience's preferences. We reiterate that the type space X must grow with N in order to maintain Assumption 1 for a fixed value $\xi > 0$, and that $H_0 \equiv \{\lambda h_0 + (1 - \lambda)h \mid h \in H\}$.

PROPOSITION 2: *Take weight $\lambda \in (0, 1)$ and parameter $\xi > 0$ in Assumption 1 as given. There exists \underline{N} such that for all $N \geq \underline{N}$, and all decision problems (DP) satisfying Assumption 1 with parameter ξ , optimal experiments are randomized. More precisely, the following hold:*

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

$$\max_{\alpha \in \mathcal{A}} \min_{h \in H_0} \mathbb{E}_{h, \mathcal{E}^*} [u(p, \alpha(e, y))] \geq \min_{h \in H_0} \mathbb{E}_h \left[\max_{a \in \{0, 1\}} u(p, a) \right] - \sqrt{\frac{1}{N}}.$$

(ii) *All deterministic experiments $e \in E$ are bounded away from first-best:*

$$\forall e \in E, \quad \max_{\alpha \in \mathcal{A}} \min_{h \in H_0} \mathbb{E}_{h, e} [u(p, \alpha(e, y))] < \min_{h \in H_0} \mathbb{E}_h \left[\max_{a \in \{0, 1\}} u(p, a) \right] - (1 - \lambda)\xi.$$

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 $1/\sqrt{N}$. This performance bound is potentially attained by several experiment designs, including the standard RCT: $(\mathcal{E}_{RCT}, \alpha_{RCT})$. Point (ii) shows that the loss from a deterministic experiment is bounded below by $(1 - \lambda)\xi$, where ξ is defined by Assumption 1. Hence, taking ξ as given, there exists \underline{N} such that for all $N > \underline{N}$ the optimal experiment is not deterministic; it is randomized. Note that this is not an asymptotic result: threshold \underline{N} is the lowest value such that $\sqrt{1/\underline{N}} < (1 - \lambda)\xi$.

The traditional interpretation of an ambiguity-averse decision-maker's problem as a zero-sum game played against nature can be used to understand the value of randomization. The decision-maker selects an experimental design \mathcal{E} and a decision rule α , while nature picks the prior h to minimize the decision-maker's payoff. If there is a known pattern in the choice of experimental assignments, nature can exploit the pattern 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).¹¹ In our context, for large sample size N , the Bayesian component of preferences is close to the first-best for any prior. Thus, the subjective benefits of

¹¹ A 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.

running a deterministic experiment are small. In contrast, under Assumption 1, the impact of randomization on robust payoff guarantees is bounded away from 0.

Randomization can also be seen as a way to get parties with divergent priors to agree on a design. Such stakeholders need not be satisfied by any given deterministic experiment: there may always exist a prior under which the experiment’s design is badly flawed. Running a randomized experiment guarantees each stakeholder that, in expectation, the final decision will be close to optimal from her perspective.

Regardless of the interpretation, our results emphasize the importance of actual randomization. Assignments that are only “nearly” random, such as assignment based on time of day (see Green and Tuscisny 2012 for a critique), or alphabetical order (as in the case of Miguel and Kremer 2004; see Deaton 2010 for a critique), remain problematic under some adversarial priors.¹² Randomization provides ex ante performance guarantees even under the most skeptical priors. Approximate randomization does not.

C. Standard RCTs as a Rule-of-Thumb

A well-known corollary of the proof of Proposition 2 (see, for example, Manski 2004) is that the standard RCT $(\mathcal{E}_{RCT}, \alpha_{RCT})$ leads to near-optimal decisions under every prior.

COROLLARY 1: *Experimentation policy $(\mathcal{E}_{RCT}, \alpha_{RCT})$ is such that for all priors $h \in \Delta(P)$,*

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

This result implies that a standard RCT guarantees approximately optimal decision-making for both Bayesian and ambiguity-averse decision-makers.¹³ Importantly, the decision-maker does not even need to specify her own preferences. When sample N is large enough, a RCT is near optimal for all parameters λ and sets of priors H . Note also that the corollary holds using the deterministic empirical success rule $\alpha_{RCT}(e, y)$. This shows that the robustness of RCTs stems from random assignment, not random policymaking.

Together with Proposition 2, Corollary 1 makes the following predictions about experimenter preferences when the sample size is large and robustness is a concern. First, experimenters strongly prefer standard RCTs over any deterministic experiment.¹⁴ Second, while the standard RCT need not be exactly optimal, the marginal

¹²Specifically, Miguel and Kremer (2004) stratifies by administrative unit before assigning schools to three treatment groups according to alphabetical order. A later paper argues that these precautions allowed for valid inference (Baird et al. 2016). Still, the Deaton (2010) critique provides an example of the challenge an adversarial prior can pose for nonrandomized designs.

¹³Note that Corollary 1 holds even in models where $|X|$ is arbitrarily large compared to N , provided covariates are sampled in a representative way. As outcomes are bounded, it is not possible for rare, and therefore hard to sample, covariate realizations to have a substantial impact on payoffs.

¹⁴Cardinal measures of the intensity of experimenter preferences could be obtained by letting the experimenter trade-off the choice of her preferred experiment design against making policy decision a at random with some probability.

improvement from using the exactly optimal random design is small compared to the improvement from using the standard RCT over a deterministic design. Note that while standard RCTs are also near optimal under a subjective expected utility model, such a model does not rationalize the strong preferences for randomization that experimenters exhibit in practice. In economics, randomization frequently requires costly negotiations with implementation partners. Furthermore, randomization comes at the cost of balance, which experimenters clearly care about, as revealed by covariate collection efforts.

III. Rerandomization

As noted above, a standard RCT is only a near-optimal solution to (DP). With small probability, it may result in unbalanced assignments that provide little real opportunity for learning under reasonable priors. To improve balance, treatment assignment needs to be non-exchangeable: if someone with particular covariates is assigned to treatment, then a participant with similar covariates should be assigned to the control group.

In practice, experimenters sometimes resolve this difficulty through rerandomization: they repeatedly randomize until they obtain an assignment that satisfies their balance objective. As Bruhn and McKenzie (2009) highlights, this practice is widespread, poorly theorized, and, in principle, a substantial deviation from one-shot randomization.

Our framework is well suited to clarify the trade-offs involved in rerandomization: it improves the subjective value of experiments at some cost to robustness. We argue that, used in moderation, rerandomization provides a simple way to trade off the subjective value of an experiment design for tolerable losses in robustness. Rerandomization captures what may be an important aspect of optimal solutions to (DP): it correlates assignment across participants. It provides an expanded class of rules-of-thumb that lets experimenters trade off subjective performance for robust performance guarantees, while ensuring near-optimal performance along both objectives.

We discuss two approaches to rerandomization: K -rerandomization, which allows for a fixed number K of rerandomizations, and balance-targeting rerandomization, which sets a specific balance target for the assignment, and draws independent assignments until that target is reached.

A. K -Rerandomized Experiments

Now, K -rerandomized experiments extend the standard RCT. As in our previous analysis, we fix the policy rule α to be the empirical success rule: $\alpha_{RCT}(e, y) = \mathbf{1}_{\bar{y}^1 \geq \bar{y}^0}$.

Given a profile of covariates $(x_i)_{i \in \{1, \dots, N\}}$, K -rerandomized experiment design \mathcal{E}_K takes the following form:

- (i) Independently draw K assignments $\{e_1, \dots, e_K\}$ with each $e_k = (\tau_{i,k})_{i \in \{1, \dots, N\}}$ an exchangeable assignment such that 50 percent of participants receive treatment $\tau = 1$.

- (ii) Select an assignment $e_K^* \in \operatorname{argmax}_{e \in \{e_1, \dots, e_K\}} \mathbb{E}_{h_0} [u(p, \alpha_{RCT}(e, y))]$ that maximizes the decision-making value of experiment e under prior h_0 .¹⁵
- (iii) Run experiment e_K^* , generating outcomes y_K .
- (iv) Choose a policy $a \in \{0, 1\}$ according to α_{RCT} .

More generally, one could replace the objective $\mathbb{E}_{h_0} [u(p, \alpha_{RCT}(e, y))]$ in stage (ii) by any utility function $B: e \in E \mapsto B(e) \in \mathbb{R}$ defined directly over experimental assignments. Propositions 3 and 4 hold for this more general class of experiment designs.

We discuss two practical aspects of rerandomization in the Appendix. First, the Appendix extends our analysis of rerandomization to sequential settings where participants arrive one after the other, and treatment must be chosen in real time. This makes balance more difficult to achieve. Second, online Appendix Section A describes priors h_0 under which maximizing subjective expected utility term $\mathbb{E}_{h_0} [u(p, \alpha_{RCT}(e, y))]$ coincides with a standard balance objective: minimizing the Mahalanobis distance between mean covariates across the treated and control groups (Rubin 1979, 1980; Cochran and Rubin 1973; Morgan and Rubin 2012).¹⁶

The remainder of this section uses our framework to clarify the potential costs and benefits of rerandomization.

B. The Trade-Off of Rerandomization

The Benefit of Rerandomization.—The benefit of rerandomization is immediately apparent: it increases the subjective expected performance of the implemented design.¹⁷

Remark 1: Subjective performance $\mathbb{E}_{h_0} [u(p, \alpha_{RCT}(e_K^*, y))]$, in which the expectation is taken given a realization for e_K^* , first-order stochastically dominates $\mathbb{E}_{h_0} [u(p, \alpha_{RCT}(e_{K-1}^*, y))]$.

In the special case where prior h_0 rationalizes minimizing the Mahalanobis distance between mean covariates, this observation coincides with the analysis of Morgan and Rubin (2012). In that work, rerandomization leads to significantly more precise estimates of treatment effects when outcomes come from a linear Gaussian model. More generally, Bugni, Canay, and Shaikh (2018) shows that balance from symmetric stratification procedures also substantially increases precision.

¹⁵ An assignment is selected with uniform probability if multiple assignments maximize subjective performance.

¹⁶ The Mahalanobis distance defined by covariates $\mathbf{x} = (x_i)_{i \in \{1, \dots, N\}}$ (with $x_i \in \mathbb{R}^k$) between two arbitrary vectors $v_1, v_2 \in \mathbb{R}^k$ is given by $(v_1 - v_2)^T \operatorname{cov}(\mathbf{x})^{-1} (v_1 - v_2)$. This is equivalent to the standard Euclidean distance in the rotated and rescaled coordinate system in which the sample covariance matrix of covariates \mathbf{x} becomes the identity matrix: that is, the Euclidean distance in the basis defined by the principal components of \mathbf{x} .

¹⁷ This holds for *any* preference over deterministic assignments used to select e_K^* among K assignments.

The Cost of Rerandomization.—Although rerandomization provides clear benefits in terms of improving the subjective value of experiments, there are common, but vague, concerns about its potential costs. As Bruhn and McKenzie (2009) documents, this leads many researchers to rerandomize, but not report the fact that they did. Our framework clarifies the issue by showing that rerandomization can indeed reduce robustness. However, this cost grows very slowly with the number of rerandomizations.

We assume that the number of experimental assignments e maximizing $\mathbb{E}_{h_0}[u(p, \alpha_{RCT}(e, y))]$ is bounded above by a finite number, independently of N .¹⁸

PROPOSITION 3: *There exists $\underline{\rho} > 0$ such that, for all N and H satisfying Assumption 1 with parameter $\xi > 0$, 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 experimental assignment is essentially deterministic. Proposition 2 implies that this precludes first-best performance guarantees. As a result, Proposition 3 encourages caution toward balance-targeting rerandomization schemes that set a fixed balance target, and rerandomize until that target is reached. If very few assignments satisfy this target, then such rerandomization algorithms cause nonvanishing robustness losses.

Still, the number of rerandomizations K necessary to cause nonvanishing robustness losses is exponential in the sample size. This suggests that a moderate number of rerandomizations may have little impact on robustness. This is indeed the case.

PROPOSITION 4: *Given $K \geq 2$, consider K -rerandomized experiment \mathcal{E}_K . For all $h \in \Delta(P)$,*

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

Proposition 4 clarifies that additional losses from rerandomization are of order $\sqrt{\ln K}$.¹⁹ When $K = O(N)$, this implies that the K -rerandomized experiment still approaches the first-best as N grows large. Still, whether this potential loss in robustness is worthwhile depends on experimenter preferences. We propose what we believe are reasonable guidelines for practice in Section V, but emphasize that they are ultimately subjective.

Remarkably, Proposition 4 holds for *any* preference over deterministic assignments used to select e_K^* among K assignments. As we discuss in Section V, this suggests that K -rerandomized experiment designs could be used to trade off robustness for objectives other than standard statistical balance, including the preferences of stakeholders.

¹⁸In the case where the balance objective coincides with the Mahalanobis distance, there will generically be two optimal assignments, as treatment and control status can be switched without changing balance.

¹⁹This bound is non-asymptotic and conservative.

C. Direct Approaches to Balance Involve Similar Trade-Offs

We conclude this section by showing that understanding K -rerandomization helps us understand other, more direct, approaches to balance.

Consider the following experiment design, denoted by $\mathcal{E}_{E_{\dagger}}$. First, the experimenter defines a set of acceptable assignments $E_{\dagger} \subset E$. For example, the set of assignments that is perfectly balanced on a few discrete-valued dimensions of covariates $x \in \mathbb{R}^k$, and tolerably balanced on the remaining dimensions. Second, an assignment e is drawn from E_{\dagger} with uniform probability and implemented. This overall design could be implemented by stratifying on these few dimensions, achieving perfect balance on those dimensions, and then drawing stratified samples until they satisfy the overall balance objective.²⁰

The potential issue here is that set E_{\dagger} may be quite small. If this is the case, then the experiment is near deterministic, and the argument of Proposition 3 applies: balance constraints reduce performance guarantees by an amount that does not vanish as N gets large. This is especially problematic if set E_{\dagger} is difficult to compute explicitly, so that the decision-maker is not naturally aware of how demanding a balance target is. The bound K in K -rerandomization mechanically ensures that the experimenter does not seek to achieve excessively difficult balance objectives. Intuitively, if it takes many draws of random samples to reach the acceptable set E_{\dagger} , then it is likely that the set E_{\dagger} is excessively small.

This intuition can be used to obtain a lower bound on the performance of experiment design $\mathcal{E}_{E_{\dagger}}$. Let $p_{E_{\dagger}}$ denote the probability that a uniformly chosen assignment $e \in E$ belongs to E_{\dagger} .²¹ Proposition 4 can be used to establish the following efficiency bound.

PROPOSITION 5: For all $h \in \Delta(P)$,

$$\mathbb{E}_{h, \mathcal{E}_{E_{\dagger}}}[u(p, \alpha_{RCT}(e, y))] \geq \mathbb{E}_h\left[\max_{a \in \{0,1\}} u(p, a)\right] - \min_{K \geq 2} \left[2\sqrt{\frac{\ln K}{N}} + (1 - p_{E_{\dagger}})^K\right].$$

This implies that a procedure seeking to achieve a given balance objective will come at a limited cost to robustness if an acceptable assignment can be reached with high probability within a small number of rerandomizations.

Quantile-Targeting Rerandomization.—This analysis helps qualify the balance-targeting rerandomization procedure endorsed by Morgan and Rubin (2012). Proposition 5 suggests specifying balance targets as a quantile of balance among all possible assignments. This makes the size of set E_{\dagger} explicit and ensures that the balance objective picked by the experimenter does not cause significant losses in robustness. For instance, if the experimenter targets the

²⁰Stratification is also sometimes referred to as “blocking.” It usually involves dividing a sample along some particular dimension, say gender, and then assigning treatment through draws without replacement within each gender. In this way, the treatment and control groups will be exactly balanced on this dimension. In our framework, stratification on all dimensions is not possible, as individuals are unique. However, stratification is possible along a few discrete-valued dimensions of covariates $x \in \mathbb{R}^k$, for instance gender.

²¹Computing $p_{E_{\dagger}}$ may be difficult for complex balance criteria. In that case, Monte Carlo simulations can be used to compute an estimate of $p_{E_{\dagger}}$.

ninety-fifth percentile of balance, then $p_{E_{\dagger}} = 5\%$ and robustness losses due to rerandomization vanish as sample size N grows large.

IV. Reference Dependence and Regret Aversion

The use of null-hypothesis statistical testing (NHST) is another common practice that standard models fail to predict. Recent work, however, shows that NHST can be rationalized if the decision-maker exhibits reference-dependence (Tetenov 2012). We establish that reference-dependent preferences are compatible with our framework.

The version of reference-dependence we use nests regret aversion as a special case. It has the additional benefit of allowing us to dispense with Assumption 1. In particular, consider payoffs

$$w(p, a) = \Delta_p^a \times (1 + \kappa_a \mathbf{1}_{\Delta_p^a < 0}), \quad \text{with } 0 < \kappa_0 \leq \kappa_1,$$

where $\Delta_p^a \equiv p^a - p^{1-a}$. These payoffs depend on both the policy choice a made by the decision-maker, and the performance difference Δ_p^a between this choice and the alternative.²² Crucially, the parameters κ_0, κ_1 , which only affect preferences when the decision-maker makes a mistake ($\Delta_p^a < 0$), imply that the possibility of a success of size Δ_p^a cannot offset an equal possibility of a mistake of the same size.²³ When $\kappa_0 = \kappa_1$, regret about a mistake of a given magnitude is the same whether a equals 0 or 1. This is the standard formulation of regret aversion in the statistical decision theory literature. On the other hand, when $\kappa_0 < \kappa_1$, mistakes made when the decision-maker chooses $a = 1$ are given more weight than those made when the decision-maker chooses $a = 0$. Thus, $a = 1$ will only be chosen if the decision-maker is sufficiently certain that $p^1 > p^0$.

We consider the variant of (DP) in which the decision-maker seeks to maximize

$$(DP') \quad \lambda \mathbb{E}_{h_0, \varepsilon} [w(p, \alpha)] + (1 - \lambda) \min_{h \in H} \mathbb{E}_{h, \varepsilon} [w(p, \alpha)],$$

where $H = \Delta(P)$ is now the set of all possible priors over states $p \in P$.²⁴

Note that we no longer impose Assumption 1 (limited extrapolation). The set of priors H is simply the set of all priors over P . As noted earlier, Assumption 1 is needed to rule out the doctrinaire prior in which $p^1 = p^0 = 0$. This would be the unconstrained worst-case prior chosen by nature under problem (DP), and no learning is possible at these priors. Nature does not select such a prior under regret aversion. Under the worst-case prior for problem (DP'), one policy must be better than the other. Otherwise, the decision-maker cannot make a mistake, and there is nothing to regret. As a result, under regret aversion, information is valuable even under the unconstrained worst-case prior.

In this environment, analogs of our earlier propositions hold.

²² Online Appendix Section B shows that these preferences rationalize NHST as an optimal decision rule.

²³ The decision-maker exhibits loss aversion, consistent with Prospect Theory (Kahneman and Tversky 1979).

²⁴ Alternatively, for any $\underline{p} < 1/2 < \bar{p}$, we could consider the set of priors $H = \Delta([\underline{p}, \bar{p}]^{2X})$.

PROPOSITION 6 (Randomization): *Consider a decision-maker solving problem (DP^{*}):*

- (i) *Whenever $\lambda = 1$, running a deterministic experiment is weakly optimal.*
- (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.*

PROPOSITION 7 (Rerandomization): *There exists $M > 0$ such that for every prior $h \in H$ and $K \geq 1$,*

$$\mathbb{E}_{h, e \sim \mathcal{E}_K, a \sim \alpha_{RCT}}[w(p, a)] \geq \mathbb{E}_h \left[\max_{a \in \{0, 1\}} w(p, a) \right] - M \sqrt{\frac{\ln(K + 1)}{N}}.$$

V. Discussion

The ambition of this paper is to provide a decision-theoretic framework for experimental design that can clarify debates about experimental practice. Our primary requirement is that any such framework should rationalize revealed preferences for randomization. In this last section, we discuss alternative rationales for randomization and formulate suggestions for practical experiment design.

A. Alternative Rationales for Randomization

An alternative rationale for randomization is that experimenters are indifferent over experimental assignment, and may as well randomize. This indifference could be ascribed to a lack of opinion over the variation in treatment effects associated with different covariates. Alternatively, experimenters may have such opinions, but find the collection of covariate information too costly.

We believe that this rationale is not consistent with the preferences revealed by the behavior of real-life experimenters. First, experimenters indicate strong preferences against deterministic assignments and in favor of randomized designs, including, but not limited to, standard RCTs. Ensuring proper randomization is often costly, and experimenters are concerned when they engage in any behavior that strays from this gold standard. This is especially true in economics, where experimenters often invest considerable time and energy trying to convince implementation partners to randomize. Second, experimenters are not indifferent about balance, and reveal mild preferences for rerandomized assignments over standard RCTs. Indeed, experimenters do obtain covariates, and frequently use them to determine experimental assignment by rerandomizing or stratifying. Concerns over balance typically vanish as the sample size grows large. Our model captures all these facts: deterministic assignments are unattractive even when the sample size is large; standard RCTs are a sizable improvement, although they may be improved by enhancing balance; the value of this last improvement vanishes as N gets large.

Another potential avenue to rationalize randomization is to view it as a delegation rule in a principal-agent setting: the audience wants data to be collected, but

wants to limit the ways in which an interested experimenter could bias findings. This concern about moral hazard is not present in our framework: we believe that experimenters and their audience find randomization valuable even in the absence of incentive issues. In other words, we believe that a team of researchers and policy designers sharing the same objectives, but not necessarily the same beliefs, would find randomization a valuable tool. While we certainly believe that some experimenters are interested in obtaining specific outcomes, for example, pharmaceutical companies, it is still difficult to explain randomization without introducing heterogeneous priors. A Bayesian regulator would prefer mandating specific experimental assignments, for example, assignments maximizing balance.

B. *Deterministic versus Randomized Experimentation*

As shown in Section II, our framework predicts either deterministic experimentation or randomized experimentation depending on the sample size N and the weight the decision-maker places on satisfying her audience λ . Here we describe how these findings relate to stylized facts about experimentation, as summarized in Figure 1. Throughout, we maintain Assumption 1 for a fixed parameter $\xi > 0$ (alternatively, we could assume the reference-dependent preferences of Section IV). Under this maintained assumption, changes in sample size N are accompanied with changes in the set of type X and priors H .

When sample points are scarce (N is small), or when the decision-maker does not put much weight on satisfying her audience (λ is close to 1), the optimal experiment is deterministic, driven by prior h_0 . The experimenter assigns treatment and control to maximize the decision-making value of the data collected. For example, in the context of process improvement, a firm testing out new machinery may assign treatment to its best-performing teams so that the evaluation is not muddled by operational errors. When refining a new product, firms test their prototypes with selected alpha users.

Similarly, early-stage medical research often does not feature randomization. For example, Phase I trials, which examine the toxicity of a new compound, are not randomized. These trials recruit healthy volunteers, which allows adverse reactions to be cleanly ascribed to treatment rather than to any underlying disease condition. Sample sizes in Phase I trials also tend to be relatively small. Deterministic experimentation also occurs when studying patients who have conditions known to rapidly result in severe disability or death.²⁵ Any improvement can be reliably assigned to the treatment. In our framework, the reasoning underlying these two cases is the same: relatively little disagreement between the experimenter and the adversarial audience leads to deterministic designs maximizing the experimenter's subjective expected utility.

When the decision-maker cares sufficiently about satisfying her adversarial audience, or when she has a sufficiently large sample, she will randomize. The former is the case in scientific research, especially when the experiment has significant policy implications, for example, in development economics. The latter is the case for firms

²⁵For recent examples, see Amy Harmon, "After Long Fight, Drug Gives Sudden Reprieve," *The New York Times*, February 22, 2010 (<http://www.nytimes.com/2010/02/23/health/research/23trial.html>), and Alan Yuhas, "Cancer Researchers Claim 'Extraordinary Results' Using T-Cell Therapy," *Guardian*, February 15, 2016 (<https://www.theguardian.com/science/2016/feb/15/cancer-extraordinary-results-t-cell-therapy-research-clinical-trials>).

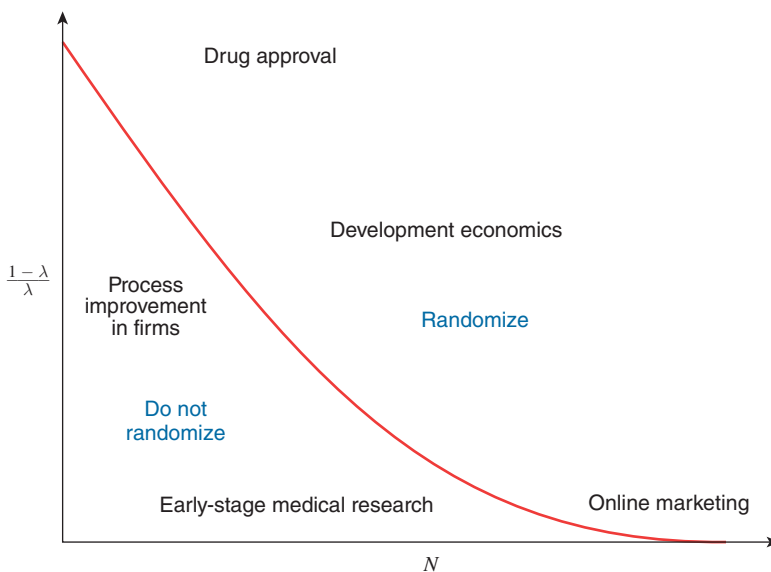


FIGURE 1. WHEN TO RANDOMIZE? THEORY MATCHES BROAD PATTERNS

doing A/B testing online. Although the firm only needs to convince itself of the effectiveness of a particular ad, there are so many observations available 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, doctors, patients, and investors form the adversarial audience for pharmaceutical and medical device companies.

C. The Practice of Rerandomization

We now discuss how our results may be used to inform the practice of rerandomization. We emphasize that our results apply to environments where the set of covariates is known at the time of treatment assignment. We extend our results to the case of sequential treatment assignment in the Appendix.

Proposition 2 and Corollary 1 establish that: (i) randomization is a key aspect of any compromise between subjective expected utility and robust payoff guarantees; and (ii) the standard RCT is an acceptable compromise, providing approximately optimal performance for each objective. We interpret Propositions 3 and 4 as expanding the range of such compromises. We emphasize two insights from our analysis that seem most useful for practice: first, how to set balance targets so that they do not inadvertently cause severe losses in robustness, and second, how to use rerandomization to allow stakeholders to express preferences.

Setting Balance Targets Properly.—Our analysis identifies the following issue with the practice of rerandomization: excessively ambitious absolute balance objectives may lead to severe losses in robustness. Propositions 4 and 5 suggest two ways to address the issue by changing the metric used to set balance targets.

The first approach is to use *K*-rerandomization, with *K* less than sample size *N*. This ensures losses on the order of $\sqrt{\ln(N)}/N$. A rule-of-thumb that strikes

us as a (subjectively) attractive compromise is to set $K \leq \min\{100, N\}$ (Banerjee, Chassang, and Snowberg 2017). Indeed, note that with probability $1 - 0.95^K$ the K -rerandomized assignment is in the ninety-fifth quantile of most balanced assignments. For any $K \geq 100$ this event has probability greater than 99 percent.

The second approach, described in Section III C, is to set a quantile of balance as an objective, for instance the ninety-fifth quantile of balance among all assignments $e \in E$, and rerandomize until that target is achieved. For any such fixed quantile, Proposition 5 ensures that losses against the first-best vanish as sample size N grows large. Concretely, an approximate way to generate such an assignment would be to draw a large set of independent, exchangeable assignments, and choose one with uniform probability among the 5 percent most balanced.

In both cases, the idea is to avoid setting absolute balance goals, so that rerandomization does not inadvertently lead to an extremely reduced set of possible assignments. Expressing rerandomization objectives either as a bound on the number of rerandomizations or as targeting a quantile of balance clarifies how unusual possible assignments are. While the optimal degree of rerandomization is ultimately subjective, we think these guidelines, setting $K = \min\{N, 100\}$ or targeting the ninety-fifth quantile of balance, offer a reasonable compromise between subjective performance and robust performance guarantees.

Expressing Stakeholder Preferences.—As noted in the text surrounding Remark 1 and Proposition 4, these results do not rely on the fact that selected experimental assignment e_K^* maximizes a specific subjective performance objective. They hold regardless of the way the final assignment e_K^* is selected, provided it is selected among K uniformly drawn random assignments.

This means that one can use a K -rerandomized design as a way to let stakeholders and implementation partners express preferences, albeit in a constrained way. Regulators, funders, or the communities from which experimental participants are drawn often have equity concerns and distributional preferences. They may care about targeting treatment to those they believe will benefit the most, or may simply dislike the lack of control inherent in randomization and wish to exert some control over assignment. The ability to at least *partially* accommodate the preferences of stakeholders, by using their preferences to select an assignment among K options, may help build goodwill and ensure cooperation.

Concretely, such a protocol would take the following form: (i) let stakeholders express preferences over assignments; (ii) draw K random assignments; (iii) pick the assignment e_K^* from that group of K randomizations that maximizes the preferences of stakeholders; (iv) run the experiment; (v) take policy decision $a = \mathbf{1}_{\bar{y}_1 \geq \bar{y}_0}$.

We note that in this description, the preferences of stakeholders are expressed *ex ante*, although Proposition 4 would continue to hold even if stakeholders picked the assignment *ex post* without prespecifying their preferences. Stating preferences *ex ante* has two benefits. First, formally specifying the preferences of stakeholders permits randomization inference tests (Fisher 1935, Morgan and Rubin 2012, Young 2016). That is, given the process for assignment, a statistician can simulate the distribution of treatment effects that would be observed if $p_x^1 = p_x^0$ for all $x \in X$. This procedure can be used to calculate exact p -values, and infer standard errors. Second, randomization helps parties with differing priors

agree on a process *ex ante*, but not *ex post*: if audience members start looking into the details of the realized assignment, someone may well find issue with it.

D. Future Directions

We believe that debates over the proper way to do empirical research are an opportunity to both improve and put to gainful use economists' models of decision-making. Other possible uses of our framework include understanding subgroup analysis and multiple-hypothesis testing. We believe these factors can be introduced into our framework by considering more complex policies that tailor treatment to the covariates of participants. Pre-analysis plans, that is, the practice of prespecifying which statistics of the data will be reported, also form a challenge to traditional models of experimentation. This practice is difficult to rationalize using a Bayesian framework as long as all collected data are made available. A related, but more fundamental, difficulty is explaining the use of low-dimensional statistics, rather than just publishing data. A possible path toward understanding these behaviors may be to embed selective attention into our framework.

APPENDIX

A. Sequential Treatment Assignment

In many settings, subjects arrive over time and treatment assignment has to be made sequentially. This is particularly true in settings where recruitment takes time (for example, medical trials), or in settings where the cost of interacting with subjects is large (for example, economic development experiments occurring in hard to reach rural areas). In such settings, it may not be practical to first learn the set of covariates X in the full sample and assign treatments afterward. In this Appendix, we show that much of our analysis holds in such environments. We propose a variant of K -rerandomization appropriate for sequential assignment.

Sequential treatment assignment has received attention from statisticians. For example, Efron (1971) advocates biased coin designs to ensure that treatment and control are sampled at similar rates, without generating significant bias in treatment assignment. More recently, Atkinson (1982, 2002) extends the approach to the design of sequentially balanced experiments when participants have covariates. Guiteras, Levine, and Polley (2016) provides simulations and details from field implementations highlighting the value of such designs in practice. To our knowledge, this literature does not address the impact of such designs on robust decision-making.

B. Framework

Label N experimental participants by their arrival period i . Each participant $i \in \{1, \dots, N\}$ exhibits an observable covariate $d_i \in D$, where D is finite. We denote by $x_i = (i, d_i)$ covariates augmented with the time of arrival of participants. In principle, outcomes $y_i \in \{0, 1\}$ may depend on the time of arrival. Let $X = \{x_i\}_{i \in \{1, \dots, N\}}$ denote the final set of covariates. Note that X is a random variable. We denote by $\mathcal{X} = \text{supp}X$ the set of potential sets of covariates X . A state of the world is a pair

(X, p) , where p takes the form $(p_x^\tau)_{\tau \in \{0,1\}, x \in X}$. A prior h now denotes a prior over the pair (X, p) . As in Section I, the payoff of a decision-maker taking decision a at state (X, p) is $u(X, p, a) = (1/|X|) \sum_{x \in X} p_x^a$.

We continue to denote by $\tau_i \in \{0, 1\}$ the treatment assignment of the i th participant. A design history s_n in period $n \leq N$ is a tuple $(x_i, \tau_i, x_n)_{i \in \{1, \dots, n-1\}}$. An experimental assignment function $e \in E$ is a mapping from design histories s_n to treatment assignments $\tau_n \in \{0, 1\}$. A randomized experiment \mathcal{E} is a distribution over experimental designs $e \in E$. The final realized experimental assignment is described by history $s_{N+1} = (x_i, \tau_i)_{i \in \{1, \dots, N\}}$.

A policy rule α is a mapping from the final experiment design s_{N+1} and realized outcome data $y = (y_1, \dots, y_N)$ to policy choices $\Delta(\{0, 1\})$. The decision-maker evaluates choices over experiment designs \mathcal{E} and policy rules α according to criterion

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

where h_0 is a prior over (X, p) , and H is a set of priors h over (X, p) . For simplicity, we assume that H takes the form $H = \{\delta_X \times h_p \mid X \in \mathcal{X}, h_p \in H_p\}$, where δ_X is the Dirac mass at X , and $H_p \subset \Delta([0, 1]^{2N})$ is a set of priors over p satisfying Assumption 1 for some fixed $\xi > 0$.

The benchmark sequential RCT design \mathcal{E}_{sRCT} takes the following form:

- For each history s_n , \mathcal{E}_{sRCT} assigns treatment $\tau = 1$ with probability 0.5.
- Policy a is chosen according to the empirical success rule: $\alpha_{RCT}(e, y) \equiv \mathbf{1}_{\bar{y}^1 \geq \bar{y}^0}$ defined in Section I.

C. Benchmark Results

The following results hold:

- (i) If $\lambda = 1$, there exists a deterministic experiment design e that maximizes objective (DP^\dagger) .

The decision-maker's value at any history s_n satisfies the Bellman equation

$$V(s_n) = \max_{\tau \in \{0,1\}} \mathbb{E}_{h_0} V(s_{n+1})$$

with $V(s_{N+1}) = \mathbb{E}_{h_0, s_{N+1}} \left[\max_{a \in \{0,1\}} \mathbb{E}_{h_0} [u(X, p, a) \mid s_{N+1}, y] \right]$.

At any history s_n the decision-maker picks τ_n , solving the first equation above. As she is an expected utility-maximizer, there is no advantage to random treatment choice. A deterministic contingent plan is optimal.

- (ii) If $\lambda > 0$ and Assumption 1 holds for $\xi > 0$ fixed, then any deterministic experiment is bounded away from the first-best by a term greater than $(1 - \lambda)\xi$.

The result follows from the fact deterministic sequential experiment designs form a subset of the deterministic experiment designs stud-

ied in Section I. As a result, payoffs under deterministic sequential experiment designs are bounded above by payoffs under deterministic experiments.

- (iii) Sequential trial design \mathcal{E}_{sRCT} yields payoffs within $O(\sqrt{1/N})$ of first-best. The proof is identical to that of Corollary 1.

D. Sequential K-Rerandomization

It is possible to provide a sequential version of the K -rerandomization process described in Section III. The design, denoted by \mathcal{E}_{sK} , takes as given a value function W defined over states s_n and period n assignments τ_n . We provide a few possibilities below. Assignments are defined as follows:

- (i) Draw K sequential assignment mappings $(e_1, \dots, e_K) \in E \times \dots \times E$ according to K sequential RCT designs $(\mathcal{E}_{sRCT}^1, \dots, \mathcal{E}_{sRCT}^K)$. Note that each individual design \mathcal{E}_{sRCT}^k assigns treatment independently across covariates, but the realizations of the assignment mappings themselves may be correlated across values $k \in \{1, \dots, K\}$. We provide examples below (in particular, see Figure A1).
- (ii) For each history s_n , we define the set of feasible treatments $T(s_n)$ at s_n as

$$T(s_n) = \{e(x_n) \mid e \in \{e_1, \dots, e_K\} \text{ such that } \forall n' < n, e(x_{n'}) = \tau_{n'}\}.$$

Note that $T(s_n)$ is the set of assignments $e_1(x_n), \dots, e_K(x_n)$, for mappings e_k whose assignments for past covariates $x_{n'}$ are consistent with realized past assignments $\tau_{n'}$ (with $n' < n$).

Treatment τ_n is chosen from $\operatorname{argmax}_{\tau \in T(s_n)} W(s_n, \tau)$. Indifferences are resolved with a uniform draw.

- (iii) Policy rule α is the empirical success rule $\alpha_{RCT} \equiv \mathbf{1}_{\bar{y}^1 \geq \bar{y}^0}$.

The proof of Proposition 4 implies that this experiment design guarantees payoffs close to the first-best, up to losses of order $\sqrt{(\ln K)/N}$. Indeed, the proof of Proposition 4 only relies on the fact that each assignment \mathcal{E}_{sRCT}^k for $k \in \{1, \dots, K\}$ provides robust payoff guarantees. Proposition 4 holds regardless of the correlation between assignment mappings.

There are several ways to define value function W . Given a prior μ_0 over X , and a balance function over final assignments $B(s_{N+1})$, one can define W via the Bellman equation

$$\begin{aligned} V(s_{N+1}) &= B(s_{N+1}), \\ V(s_n) &= \max_{\tau_n \in T(s_n)} \mathbb{E}_{\mu_0} [V(s_{n+1}) \mid s_n, \tau_n], \\ W(s_n, \tau_n) &= \mathbb{E}_{\mu_0} [V(s_{n+1}) \mid s_n, \tau_n]. \end{aligned}$$




FIGURE A1. CORRELATED ASSIGNMENTS

In practice, it may be convenient to use the sample empirical prior estimated over past covariates to predict future covariates at each history. Concretely, for any n , $\hat{\mu}_n$ is the belief over future covariates $((n + 1, d_{n+1}), \dots, (N, d_N))$ obtained by resampling covariates (d_1, \dots, d_n) with replacement. At every n we define V_n by

$$V_n(s_{N+1}) = B(s_{N+1}),$$

$$\forall n' \geq n, \quad V_n(s_{n'}) = \max_{\tau_n \in T(s_{n'})} \mathbb{E}_{\hat{\mu}_n} [V_n(s_{n'+1}) | s_{n'}, \tau_n],$$

and define $W(s_n, \tau) = \mathbb{E}_{\hat{\mu}_n} [V_n(s_{n+1}) | s_n, \tau]$.

The reason that one may want to correlate assignments e_1, \dots, e_K is to optimally delay selecting the ultimate assignment. For instance, if $K = 2$, it may be beneficial to draw two assignments e_1 and e_2 that are identical up to participant $N/2$, and independent after participant $N/2$. That is, for all $d \in D$, $e_1((i, d)) = e_2((i, d))$ if $i \leq N/2$, and $e_1((i, d))$ independent of $e_2((i, d))$ if $i > N/2$. This allows the experimenter to first learn the distribution of covariates z in the population before picking the continuation assignment that maximizes expected balance.

More generally, if the number of randomizations is $K = C^G$, with C and G two integers, one could select the final assignment e in G steps, with choices made at times n_1, \dots, n_G . At each period n_g , the decision-maker chooses between C independently drawn assignment mappings for covariates realized between time n_g and $n_{g+1} - 1$. In that case, action set $T(s_n)$ is a singleton for all periods $n \notin \{n_1, \dots, n_G\}$. Figure A1 illustrates such correlated assignments in a setting with $K = 4$, $G = 2$, $C = 2$, $n_1 = N/4$, and $n_2 = N/2$.

REFERENCES

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien. 1991. "Optimal Learning by Experimentation." *Review of Economic Studies* 58 (4): 621–54.

Athey, Susan, and Guido Imbens. 2017. "The Econometrics of Randomized Experiments." In *Handbook of Field Experiments*, Vol. 1, edited by Abhijit Banerjee and Esther Dufo, 73–140. New York: Elsevier.

Atkinson, Anthony C. 1982. "Optimum Biased Coin Designs for Sequential Clinical Trials with Prognostic Factors." *Biometrika* 69 (1): 61–67.

Atkinson, Anthony C. 2002. "The Comparison of Designs for Sequential Clinical Trials with Covariate Information." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 165 (2): 349–73.

- Baird, Sarah, Joan Hamory Hicks, Michael Kremer, and Edward Miguel.** 2016. "Worms at Work: Long-Run Impacts of a Child Health Investment." *Quarterly Journal of Economics* 131 (4): 1637–80.
- Banerjee, Abhijit, Sylvain Chassang, and Erik Snowberg.** 2017. "Decision Theoretic Approaches to Experiment Design and External Validity." In *Handbook of Field Experiments*, Vol. 1, edited by Abhijit Banerjee and Esther Duflo, 141–74. New York: Elsevier.
- Benjamin, Daniel J., James O. Berger, Magnus Johannesson, Brian A. Nosek, Eric-Jan Wagenmakers, Richard Berk, Kenneth A. Bollen, et al.** 2018. "Redefine Statistical Significance." *Nature Human Behaviour* 2 (1): 6–10.
- Bergemann, Dirk, and Juuso Välimäki.** 1996. "Learning and Strategic Pricing." *Econometrica* 64 (5): 1125–49.
- Bergemann, Dirk, and Juuso Välimäki.** 2002. "Information Acquisition and Efficient Mechanism Design." *Econometrica* 70 (3): 1007–33.
- Bergemann, Dirk, and Juuso Välimäki.** 2006. "Bandit Problems." Cowles Foundation Paper 1292.
- Blackwell, David.** 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics* 24 (2): 265–72.
- Bruhn, Miriam, and David McKenzie.** 2009. "In Pursuit of Balance: Randomization in Practice in Development Field Experiments." *American Economic Journal: Applied Economics* 1 (4): 200–32.
- Bugni, Federico A., Ivan A. Canay, and Azeem M. Shaikh.** 2018. "Inference under Covariate-Adaptive Randomization." *Journal of the American Statistical Association* 113 (524): 1784–96.
- Cochran, William G., and Donald B. Rubin.** 1973. "Controlling Bias in Observational Studies: A Review." *Sankhya: The Indian Journal of Statistics, Series A* 35 (4): 417–46.
- Deaton, Angus.** 2010. "Instruments, Randomization, and Learning about Development." *Journal of Economic Literature* 48 (2): 424–55.
- Duflo, Esther, Rachel Glennerster, and Michael Kremer.** 2008. "Using Randomization in Development Economics Research: A Toolkit." In *Handbook of Development Economics*, Vol. 4, edited by T. Paul Schultz and John A. Strauss, 3895–962. Amsterdam: Elsevier.
- Efron, Bradley.** 1971. "Forcing a Sequential Experiment to Be Balanced." *Biometrika* 58 (3): 403–17.
- Fisher, Ronald A.** 1935. *The Design of Experiments*. Edinburgh: Oliver and Boyd.
- Gilboa, Itzhak, and David Schmeidler.** 1989. "Maxmin Expected Utility with a Non-Unique Prior." *Journal of Mathematical Economics* 18 (2): 141–53.
- Green, Donald P., and Andrej Tuscisny.** 2012. "Statistical Analysis of Results from Laboratory Studies in Experimental Economics: A Critique of Current Practice." Unpublished.
- Grossman, Sanford J., and Joseph E. Stiglitz.** 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70 (3): 393–408.
- Guiteras, Raymond P., David I. Levine, and Thomas H. Polley.** 2016. "The Pursuit of Balance in Sequential Randomized Trials." *Development Engineering* 1: 12–25.
- Hansen, Lars Peter, and Thomas J. Sargent.** 2001. "Robust Control and Model Uncertainty." *American Economic Review* 91 (2): 60–66.
- Huber, Peter J.** 1964. "Robust Estimation of a Location Parameter." *Annals of Mathematical Statistics* 35 (1): 73–101.
- Huber, Peter J.** 2010. "Robust Statistics." In *International Encyclopedia of Statistical Science*, edited by Miodrag Lovric, 1248–51. Berlin: Springer.
- Humphreys, Macartan, Raul Sanchez de la Sierra, and Peter van der Windt.** 2013. "Fishing, Commitment, and Communication: A Proposal for Comprehensive Nonbinding Research Registration." *Political Analysis* 21 (1): 1–20.
- Imbens, Guido W.** 2010. "Better LATE than Nothing: Some Comments on Deaton (2009) and Heckman and Urzua (2009)." *Journal of Economic Literature* 48 (2): 399–423.
- Kahneman, Daniel, and Amos Tversky.** 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–92.
- Kasy, Maximilian.** 2016. "Why Experimenters Might Not Always Want to Randomize, and What They Could Do Instead." *Political Analysis* 24 (3): 324–38.
- Kitagawa, Toru, and Aleksey Tetenov.** 2018. "Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice." *Econometrica* 86 (2): 591–616.
- Manski, Charles F.** 2004. "Statistical Treatment Rules for Heterogeneous Populations." *Econometrica* 72 (4): 1221–46.
- Manski, Charles F.** 2009. "The 2009 Lawrence R. Klein Lecture: Diversified Treatment under Ambiguity." *International Economic Review* 50 (4): 1013–41.
- Marinacci, Massimo.** 2015. "Model Uncertainty." *Journal of the European Economic Association* 13 (6): 1022–1100.
- Miguel, Edward, and Michael Kremer.** 2004. "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities." *Econometrica* 72 (1): 159–217.

- Morgan, Kari Lock, and Donald B. Rubin.** 2012. "Rerandomization to Improve Covariate Balance in Experiments." *Annals of Statistics* 40 (2): 1263–82.
- Neyman, Jerzy.** 1923. "On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9." *Roczniki Nauk Rolniczych Tom X*: 1–51. Translated by D. M. Dabrowska and T. P. Speed (1990, *Statistical Science* 5 (4): 465–72.)
- Olken, Benjamin A.** 2015. "Promises and Perils of Pre-Analysis Plans." *Journal of Economic Perspectives* 29 (3): 61–80.
- Persico, Nicola.** 2000. "Information Acquisition in Auctions." *Econometrica* 68 (1): 135–48.
- Rothschild, Michael.** 1974. "A Two-Armed Bandit Theory of Market Pricing." *Journal of Economic Theory* 9 (2): 185–202.
- Rubin, Donald B.** 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66 (5): 688–701.
- Rubin, Donald B.** 1979. "Using Multivariate Matched Sampling and Regression Adjustment to Control Bias in Observational Studies." *Journal of the American Statistical Association* 74 (366): 318–28.
- Rubin, Donald B.** 1980. "Bias Reduction Using Mahalanobis-Metric Matching." *Biometrics* 36 (2): 293–98.
- Saito, Kota.** 2015. "Preferences for Flexibility and Randomization under Uncertainty." *American Economic Review* 105 (3): 1246–71.
- Savage, Leonard J.** 1954. *The Foundations of Statistics*. New York: Dover Publications.
- Tetenov, Aleksey.** 2012. "Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria." *Journal of Econometrics* 166 (1): 157–65.
- Wald, Abraham.** 1950. *Statistical Decision Functions*. New York: Wiley.
- Weerahandi, Samaradasa, and James V. Zidek.** 1981. "Multi-Bayesian Statistical Decision Theory." *Journal of the Royal Statistical Society, Series A (General)* 144 (1): 85–93.
- Weerahandi, Samaradasa, and James V. Zidek.** 1983. "Elements of Multi-Bayesian Decision Theory." *Annals of Statistics* 11 (4): 1032–46.
- Young, Alwyn.** 2016. "Channelling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results." Unpublished.