Causal Discovery, Randomization and Individualized Treatment

Konstantin Genin
Group Leader: "Epistemology and Ethics of ML"

In collaboration with
Conor Mayo-Wilson, University of Washington (Seattle)

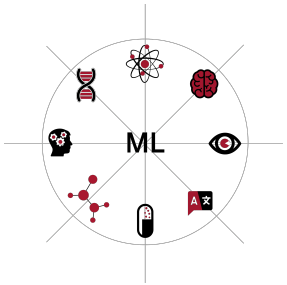

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

The Randomized, Controlled Trial (RCT)

“The RCT is the introduction of scientific method into the process of comparing treatments”

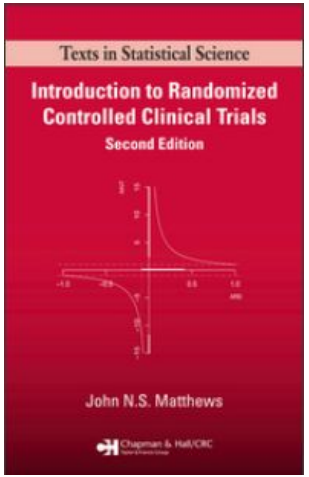

The Randomized, Controlled Trial (RCT)

Attempts to discover the relative effectiveness of a new intervention over standard treatment or placebo. Patients are assigned to the different “arms” of the trial by a randomization device.

- Widely considered the “gold standard” research design;
- Typically necessary for FDA approval;
- Raises a number of tricky ethical issues.

A call for RCTs for ML models in Clinical Settings

Perspective

Evaluating artificial intelligence in medicine: phases of clinical research

Yoonyoung Park[1], Gretchen Purcell Jackson[2,3], Morgan A. Foreman[1], Daniel Gruen[1], Jianying Hu[4] and Amar K. Das[1]

[1]Center for Computational Health, IBM Research Cambridge, Cambridge, Massachusetts, USA, [2]Center for AI, Research, and Evaluation, IBM Watson Health, Cambridge, MA, USA, [3]Department of Pediatric Surgery, Vanderbilt University Medical Center, Nashville, Tennessee, USA and [4]Center for Computational Health, IBM T.J. Watson Research Center, Yorktown Heights, New York, USA

Corresponding Author: Yoonyoung Park, IBM Research, 75 Binney Street, Cambridge, MA 02142, USA (yoonyoung.park@ibm.com)

Received 11 September 2019; Revised 15 May 2020; Editorial Decision 25 June 2020; Accepted 1 July 2020

ABSTRACT

Increased scrutiny of artificial intelligence (AI) applications in healthcare highlights the need for real-world evaluations for effectiveness and unintended consequences. The complexity of healthcare, compounded by the user- and context-dependent nature of AI applications, calls for a multifaceted approach beyond traditional in silico evaluation of AI. We propose an interdisciplinary, phased research framework for evaluation of AI imple-

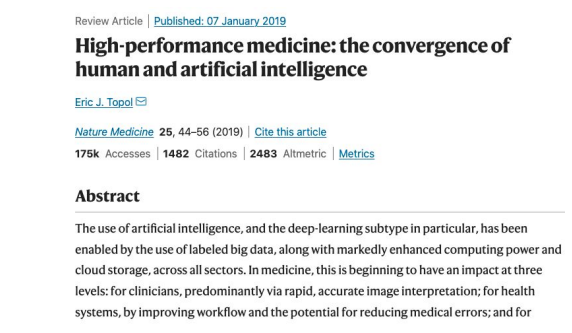

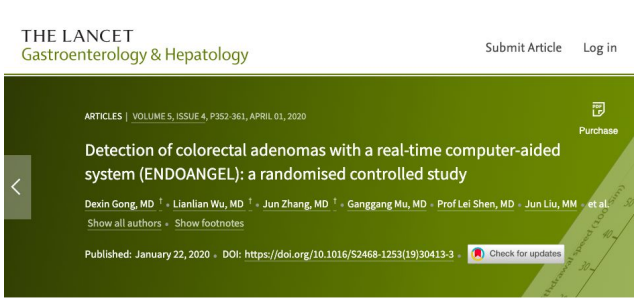

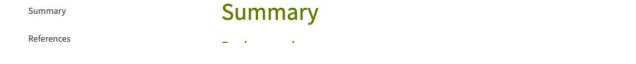

A Proliferation of ML methods for Causal Discovery

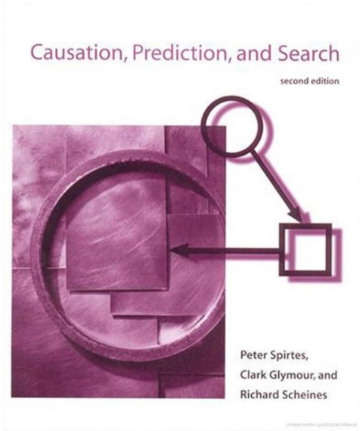

1993/2001

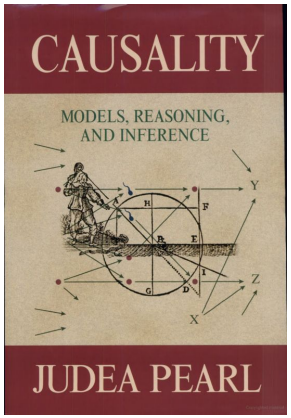

2000

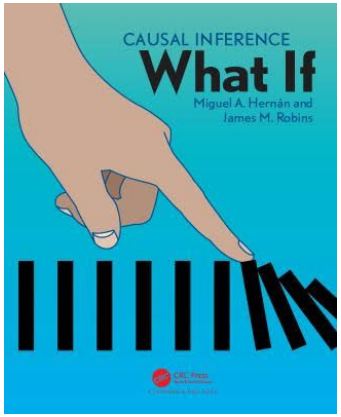

2010

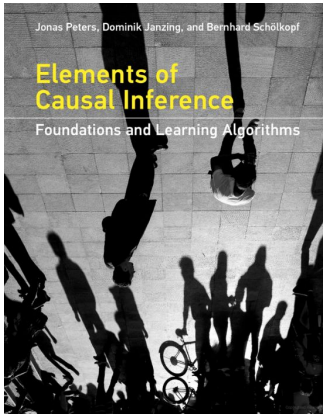

2017

A Proliferation of ML methods for Causal Discovery

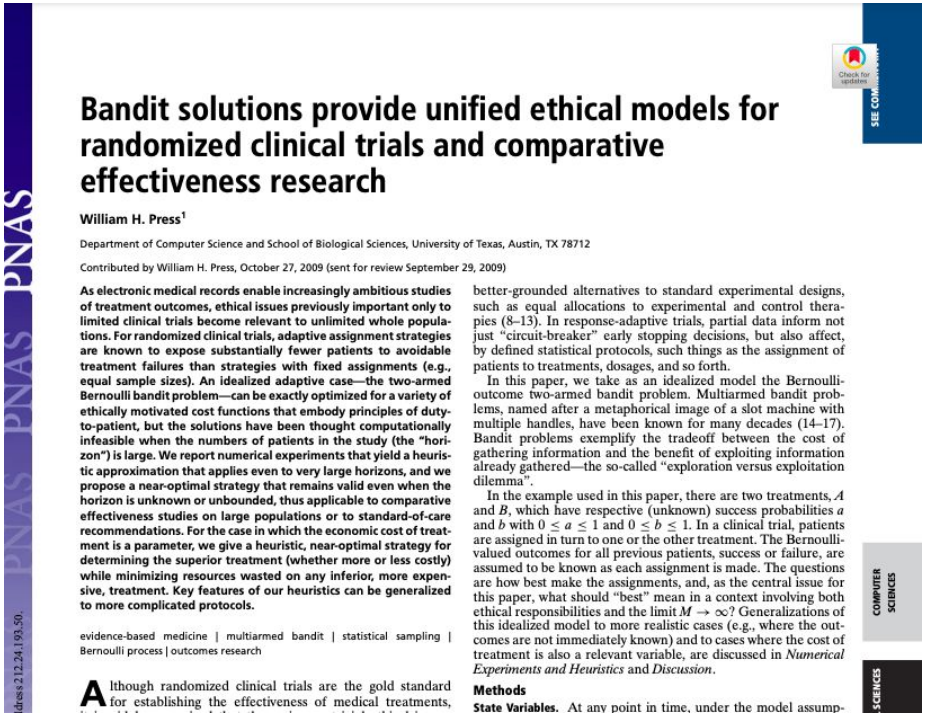

PNAS PNAS PNAS PNAS

SEE COMMENTARY

Bandit solutions provide unified ethical models for randomized clinical trials and comparative effectiveness research

William H. Press[1]

Department of Computer Science and School of Biological Sciences, University of Texas, Austin, TX 78712

Contributed by William H. Press, October 27, 2009 (sent for review September 29, 2009)

As electronic medical records enable increasingly ambitious studies of treatment outcomes, ethical issues previously important only to limited clinical trials become relevant to unlimited whole populations. For randomized clinical trials, adaptive assignment strategies are known to expose substantially fewer patients to avoidable treatment failures than strategies with fixed assignments (e.g., equal sample sizes). An idealized adaptive case—the two-armed Bernoulli bandit problem—can be exactly optimized for a variety of ethically motivated cost functions that embody principles of duty-to-patient, but the solutions have been thought computationally infeasible when the numbers of patients in the study (the "horizon") is large. We report numerical experiments that yield a heuristic approximation that applies even to very large horizons, and we propose a near-optimal strategy that remains valid even when the horizon is unknown or unbounded, thus applicable to comparative effectiveness studies on large populations or to standard-of-care recommendations. For the case in which the economic cost of treatment is a parameter, we give a heuristic, near-optimal strategy for determining the superior treatment (whether more or less costly) while minimizing resources wasted on any inferior, more expensive, treatment. Key features of our heuristics can be generalized to more complicated protocols.

evidence-based medicine | multiarmed bandit | statistical sampling | Bernoulli process | outcomes research

Although randomized clinical trials are the gold standard for establishing the effectiveness of medical treatments, better-grounded alternatives to standard experimental designs, such as equal allocations to experimental and control therapies (8–13). In response-adaptive trials, partial data inform not just "circuit-breaker" early stopping decisions, but also affect, by defined statistical protocols, such things as the assignment of patients to treatments, dosages, and so forth.

In this paper, we take as an idealized model the Bernoulli-outcome two-armed bandit problem. Multiarmed bandit problems, named after a metaphorical image of a slot machine with multiple handles, have been known for many decades (14–17). Bandit problems exemplify the tradeoff between the cost of gathering information and the benefit of exploiting information already gathered—the so-called "exploration versus exploitation dilemma".

In the example used in this paper, there are two treatments, A and B, which have respective (unknown) success probabilities a and b with $0 \le a \le 1$ and $0 \le b \le 1$. In a clinical trial, patients are assigned in turn to one or the other treatment. The Bernoulli-valued outcomes for all previous patients, success or failure, are assumed to be known as each assignment is made. The questions are how best make the assignments, and, as the central issue for this paper, what should "best" mean in a context involving both ethical responsibilities and the limit $M \to \infty$? Generalizations of this idealized model to more realistic cases (e.g., where the outcomes are not immediately known) and to cases where the cost of treatment is also a relevant variable, are discussed in *Numerical Experiments and Heuristics* and *Discussion*.

Methods

State Variables. At any point in time, under the model assump-

COMPUTER SCIENCES

SCIENCES

Check for updates

Article

Rethinking the Gold Standard With Multi-armed Bandits: Machine Learning Allocation Algorithms for Experiments

Organizational Research Methods
2021, Vol. 24(1) 78-103

DOI: 10.1177/1094428119854153
journals.sagepub.com/home/orm
SAGE

Chris Kaibel[1] and Torsten Biemann[1]

Abstract

In experiments, researchers commonly allocate subjects randomly and equally to the different treatment conditions before the experiment starts. While this approach is intuitive, it means that new information gathered during the experiment is not utilized until after the experiment has ended. Based on methodological approaches from other scientific disciplines such as computer science and medicine, we suggest machine learning algorithms for subject allocation in experiments. Specifically, we discuss a Bayesian multi-armed bandit algorithm for randomized controlled trials and use Monte Carlo simulations to compare its efficiency with randomized controlled trials that have a fixed and balanced subject allocation. Our findings indicate that a randomized allocation based on Bayesian multi-armed bandits is more efficient and ethical in most settings. We develop recommendations for researchers and discuss the limitations of our approach.

Keywords

experiments, randomized controlled trial, multi-armed bandit, exploration versus exploitation, machine learning, ethics in research

Between Morals and Methodology

If the ethical costs of RCTs are justified, it must be in virtue of the **epistemic superiority** of the randomized, controlled design.

- Exactly what is the epistemic good which RCTs secure?

- Can these goods not be secured with some **other** methodology? Preferably without the same ethical costs?

An Early *Controlled* Trial

1747: James Lind, surgeon aboard the HMS Salisbury treats 12 sailors, receiving the same rations, suffering from scurvy.

2 with cider; 2 with seawater; 2 with horseradish;

2 with vinegar; 2 with sulfuric acid

and 2 with **lemons and oranges.**

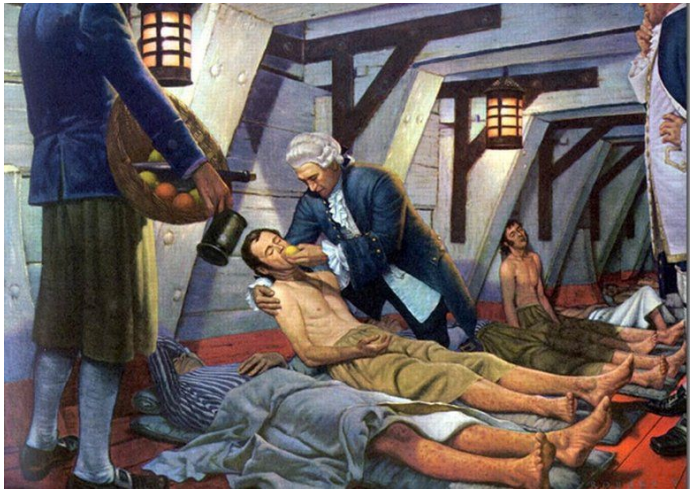

The First *Randomized* Controlled Trial

1948: A. Bradford Hill, facing a shortage of streptomycin, runs the first randomly allocated trial of streptomycin for tuberculosis.

1965: Bradford Hill proposed a set of nine criteria for epidemiological evidence of a causal relationship. Henceforth widely known as the **Bradford Hill criteria.**

Early Randomizers

1935: Fisher's publishes *The Design of Experiments.*

1925: Fisher publishes *Statistical Methods for Research Workers.*

1921: Fisher publishes *Studies in Crop Variation.*

1883: C.S. Peirce and Jastrow perform randomized experiments in psychophysics.

1780: Charles Deslon proposes a randomized trial to test Mesmer's claims.

The Trouble with Randomization

Randomization comes into prima facie conflict with **therapeutic obligation:**

“A physician should not recommend for a patient therapy such that, given present medical knowledge, the hypothesis that the particular therapy is inferior to some other therapy is more probable than the opposite hypothesis” (Marquis, 1983).

The Trouble with Randomization

Randomization comes into prima facie conflict with **individualized treatment:**

“Although a patient who has been enrolled as a research subject in a RCT may benefit from the therapeutic effects of the treatment being tested, the fact that the treatment cannot be entirely tailored to that patient’s special needs seems to violate the physician’s obligation of unqualified fidelity to his patient’s health” (Schafer, 1983).

Clinical Equipoise

Since theoretical equipoise is very fragile, Freedman proposes **clinical equipoise** instead, which obtains when

“[t]here exists (or, … may soon exit) an honest, professional disagreement among expert clinicians about the preferred treatment” (1987, 144).

The Tragic View of Clinical Research

The discussion around clinical equipoise presupposes

- There is some valuable **epistemic good** secured by randomization;
- **Any** trial methodology which secures this good must **inevitably** come into conflict with the requirements of individual treatment**.**

The Tragic View of Clinical Research

The job of clinical ethics is to reconcile clinicians to this tragic situation:

“These clinical instincts, while understandable and laudable, have the potential to obscure the true nature of clinical research, as investigators and participants alike try to convince themselves that clinical research involves nothing more than the provision of clinical care. One way to try to address this collective and often willful confusion would be to identify a justification for exposing research participants to net risks for the benefit of others.” (Wendler, 2021).

The Tragic View of Clinical Research

The job of clinical ethics is to reconcile clinicians to this tragic situation:

“These clinical instincts, while understandable and laudable, have the potential to obscure the true nature of clinical research, as investigators and participants alike try to convince themselves that clinical research involves nothing more than the provision of clinical care. One way to try to address this collective and often willful confusion would be to identify a justification for exposing research participants to net risks for the benefit of others.” (Wendler, 2021).

David S. Wendler, M.A., Ph.D.

Senior Investigator

Department of Bioethics

NIH Clinical Center

VIEW SITE

The Tragic View of Clinical Research

But is the tragic view **right?**

- What is the valuable **epistemic good** secured by randomization?

- Is there really **no** methodology that reconciles this good with the ethical requirements of individualized treatment?

Critics of Randomization

Randomization has come in for criticism on purely epistemic grounds.

- Bayesians have a hard time rationally reconstructing randomization (Savage 1961,1962; Kadane & Seidenfeld, 1999; Kasy 2016).

- The theory of optimal design does not endorse randomization (Kiefer 1959; Harville 1975).

- Philosophers of science have criticized the coherence of randomization (Urbach 1985; Worrall 2002).

Randomization On its Own Terms

What is the best frequentist justification for randomization?

The Causal Situation

T := treatment (binary);
E := effect (binary);
M := measured covariates;
U := unmeasured covariates;
I := randomizer.

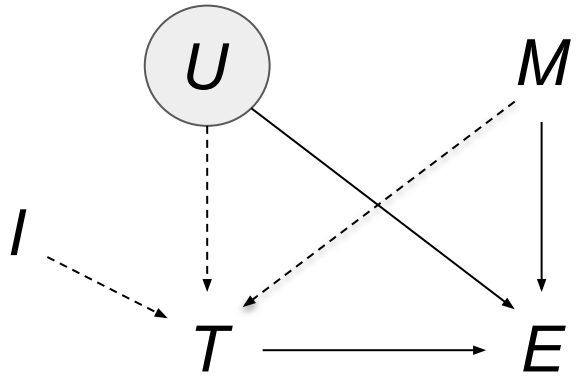

Average Treatment Effect

The goal is to estimate the
average treatment effect (ATE):

$$P(E = 1|do(T = 1)) - P(E = 1|do(T = 0))$$

Or, in the notation of the potential outcomes framework:

$$\frac{1}{n}\sum_{i \leq n} P(E_i^{t=1} = 1) - P(E_i^{t=0} = 1)$$

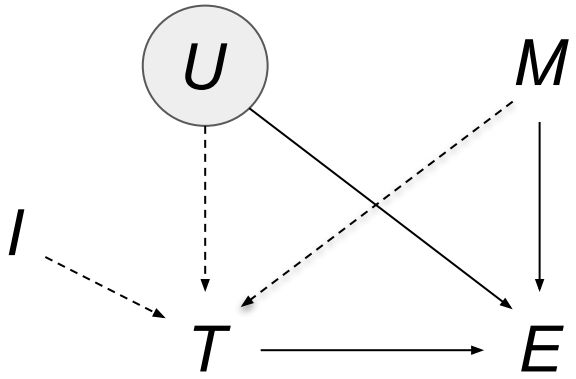

Trouble with Observational Studies

If there is an unobserved common cause of *T, E*
it is easy to come up with examples in which the ATE
is **not identified.**

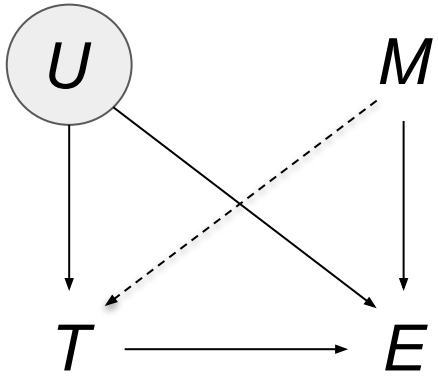

Trouble with Observational Studies

If there is an unobserved common cause of *T, E*
it is easy to come up with examples in which the ATE
is **not identified.**

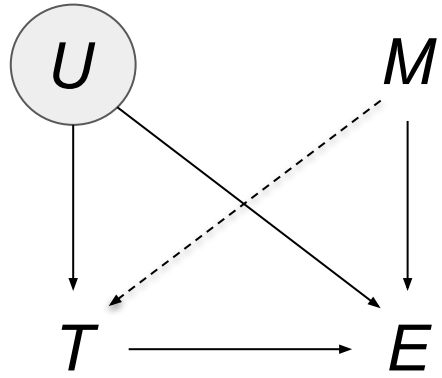

Trouble with Observational Studies

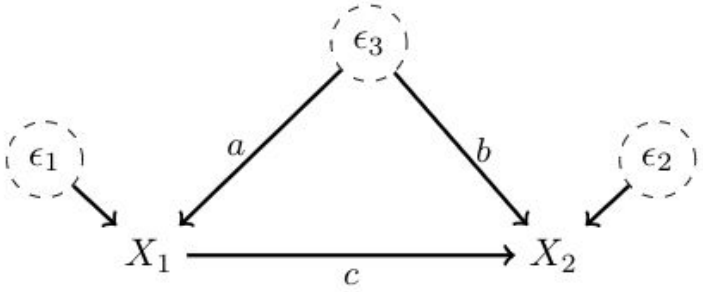

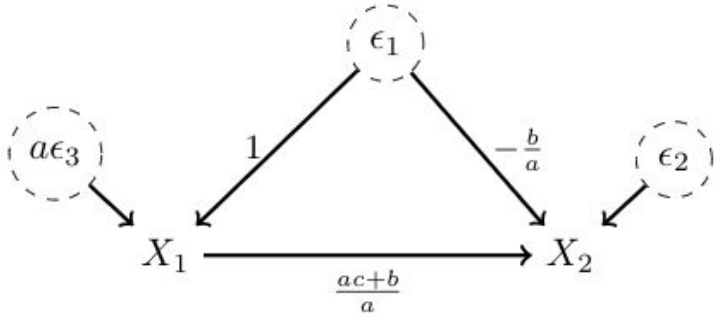

The Point of Randomization

Randomization “breaks edges” into treatment, so that any association between *T* and *E* is due to the causal effect of *T* on *E* and not shared common causes.

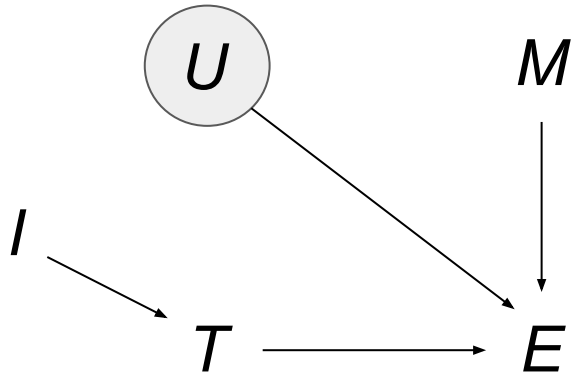

The Point of Randomization

It ensures that the ATE is **identified** and equal to

$$P(E = 1|T = 1) - P(E = 1|T = 0)$$

Moreover an **unbiased estimate** of the ATE is easily obtained.

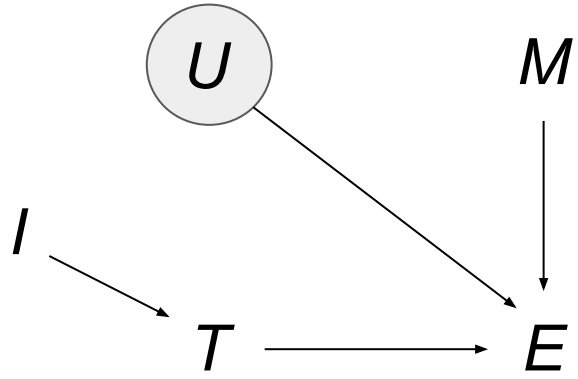

The Point of Randomization

“In ideal randomized experiments, association is causation”

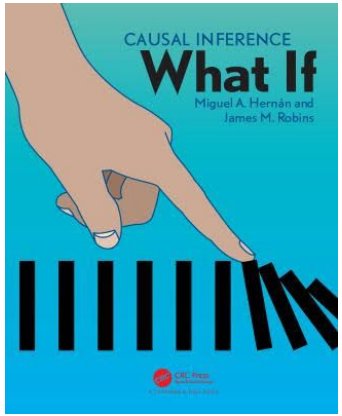

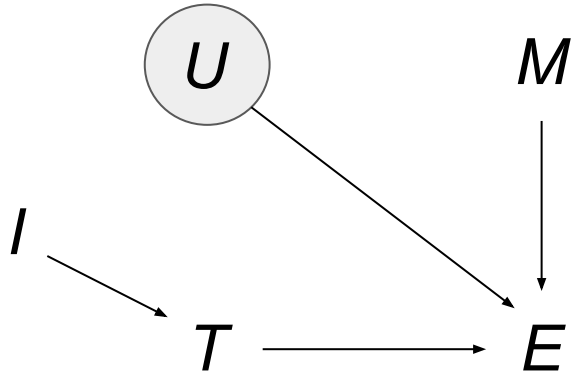

No Other Way?

But is breaking edges into *T* the only way to render the ATE identified and construct unbiased estimates?

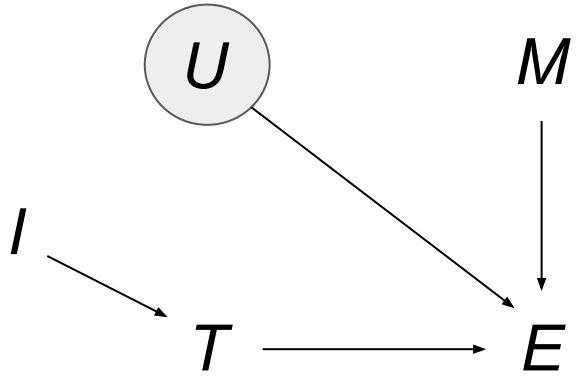

No Other Way?

But is breaking edges into *T* the only way to render the ATE identified and construct unbiased estimates?

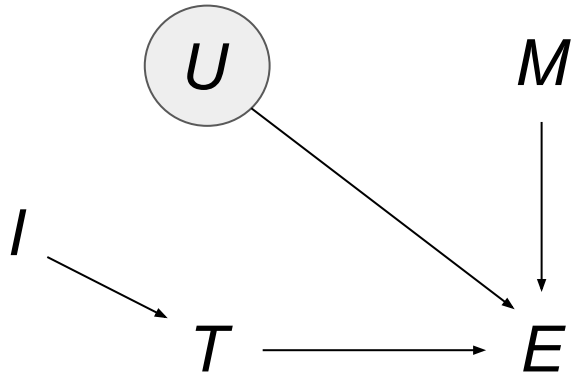

No!

Instrumental Variables

I is an **instrumental variable** if (roughly)

- *I* is statistically independent of *U,M;*
- the only unblocked path from *I* to *E* goes through *T*

  (a path is blocked if it contains a sequence like … → *T* ← …).

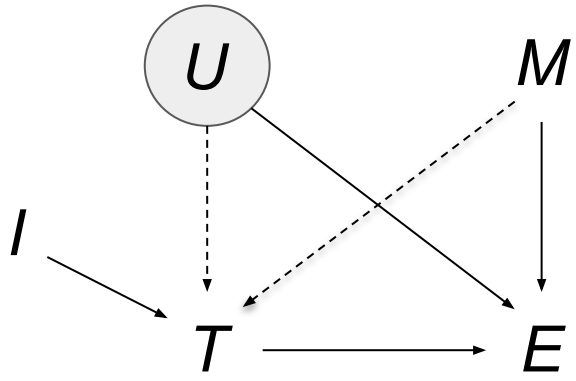

Instrumental Variables

Suppose that

- physicians assign patients to treatment according to their therapeutic judgement
- and only consult a randomizing device (*I*) when they are in equipoise

then *I* is an instrumental variable.

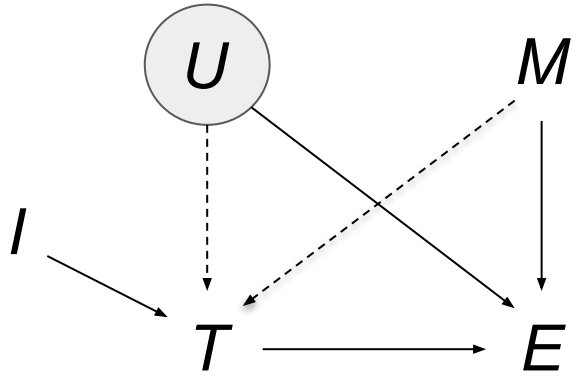

Instrumental Variables

Theorem (Angrist and Imbens 1995): When an instrumental variable satisfies a “monotonicity” condition, then the ATE is **identified** and there is an **unbiased estimator** of the ATE.

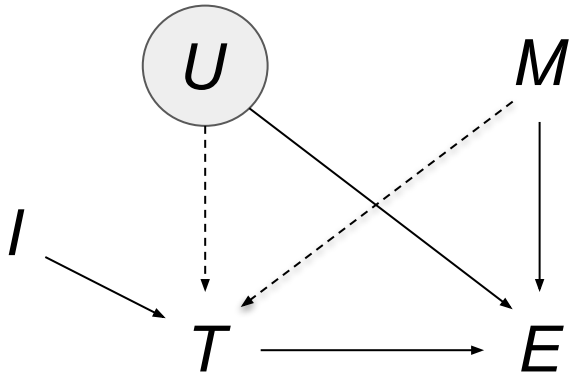

Instrumental Variables

Theorem (Angrist and Imbens 1995): When an instrumental variable satisfies a “monotonicity” condition, then the ATE is **identified** and there is an **unbiased estimator** of the ATE.

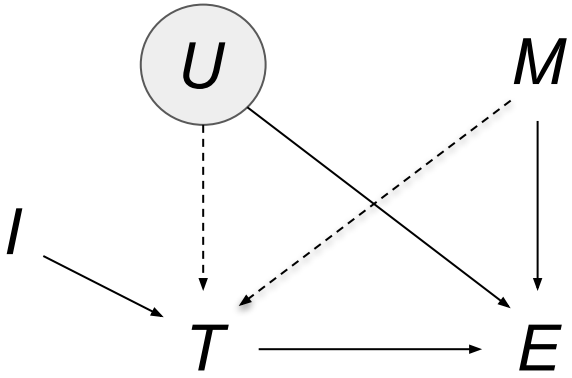

Backdoor Adjustment

M satisfied the backdoor criterion w.r.t (*T, E)* if

- *M* is not a descendant of *T*;
- *M* blocks every path between *T* and *E* that has an arrow into *T.*

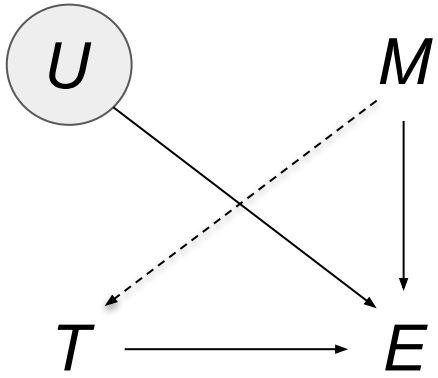

Backdoor Adjustment

Theorem (Pearl, 1993) If there is observed variable *Z* satisfying the backdoor criterion wrt (*T, E)*, then it is possible to construct an unbiased estimate of the causal effect of *T* on *E.*

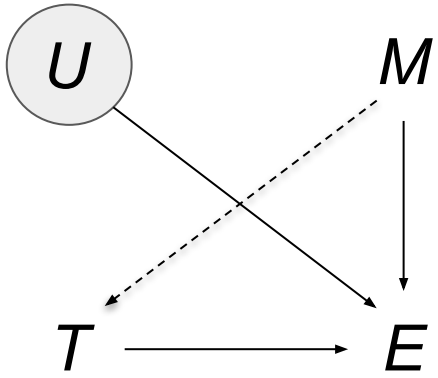

Backdoor Adjustment

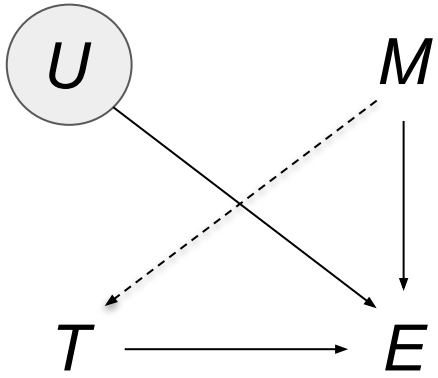

Suppose that

- physicians make assignment to treatment **only** on the basis of observed covariates *M*,

then *M* satisfies the backdoor criterion wrt (*T, E*).

Backdoor Adjustment

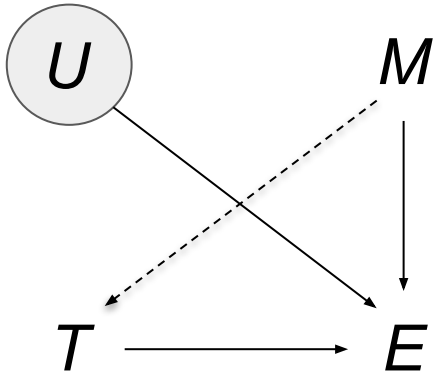

Suppose that

- physicians make assignment to treatment **only** on the basis of observed covariates *M*,

then *M* satisfies the backdoor criterion wrt (*T, E*).

Cut-off Designs

Clinicians rate patients on a continuous scale according to disease severity then

- assign low/high severity patients to less/more aggressive treatment, respectively;

- Randomize patients with moderate severity.

Estimate the ATE by **ridge regression.**

(See Cappelleri, 1995).

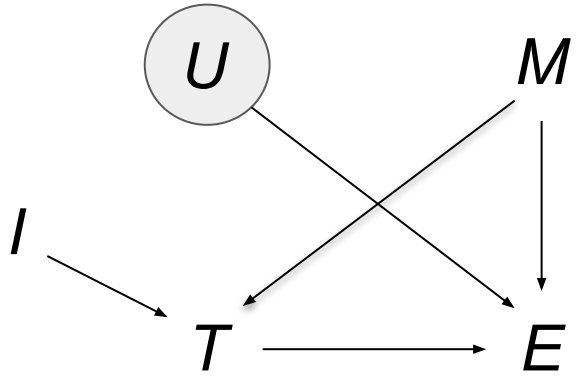

Cut-off Designs

Senn (2008) criticizes this design on grounds of **efficiency.**

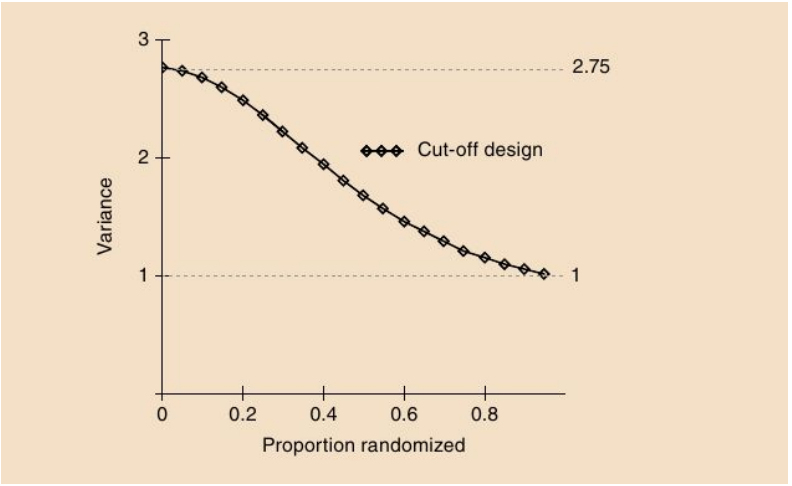

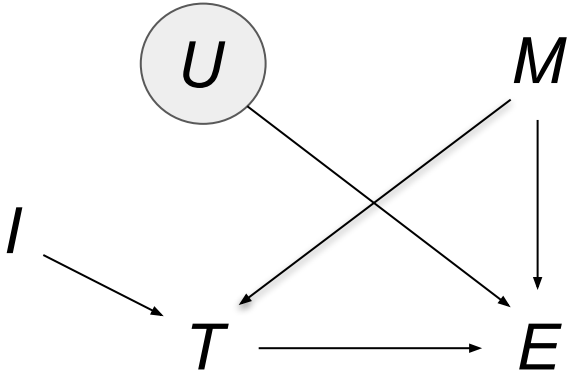

Randomization On its Own Terms

Neither guaranteeing that

1. the ATE is identified, nor that
2. there exists an unbiased estimator of the ATE,

is sufficient to justify randomization.

Other designs get the same goods and are less hostile to individualized treatment.

Randomization On its Own Terms

If there is a frequentists argument justifying randomization over other methods, it **cannot** be framed in terms of identification or unbiased estimation.

It must be about **efficiency**.

I.e. the **variance** of the estimator.

Randomization On its Own Terms

Are there such arguments?

Minimax Justifications

A series of somewhat neglected papers (Wu 1981; Li 1983; Waite and Woods 2020) develops a **minimax risk** argument for randomization.

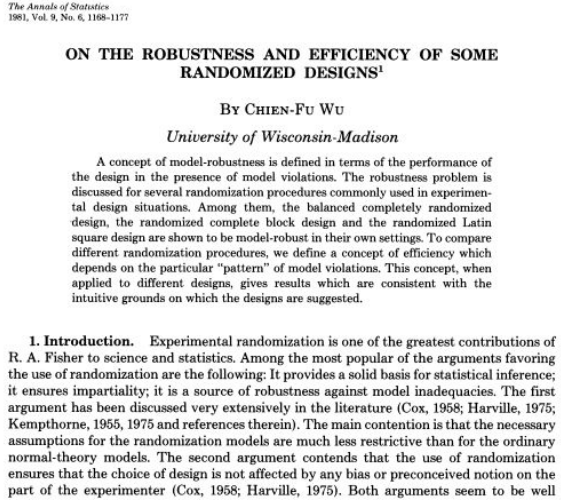

The Annals of Statistics
1981, Vol. 9, No. 6, 1168–1177

ON THE ROBUSTNESS AND EFFICIENCY OF SOME RANDOMIZED DESIGNS[1]

BY CHIEN-FU WU

University of Wisconsin-Madison

A concept of model-robustness is defined in terms of the performance of the design in the presence of model violations. The robustness problem is discussed for several randomization procedures commonly used in experimental design situations. Among them, the balanced completely randomized design, the randomized complete block design and the randomized Latin square design are shown to be model-robust in their own settings. To compare different randomization procedures, we define a concept of efficiency which depends on the particular "pattern" of model violations. This concept, when applied to different designs, gives results which are consistent with the intuitive grounds on which the designs are suggested.

1. Introduction. Experimental randomization is one of the greatest contributions of R. A. Fisher to science and statistics. Among the most popular of the arguments favoring the use of randomization are the following: It provides a solid basis for statistical inference; it ensures impartiality; it is a source of robustness against model inadequacies. The first argument has been discussed very extensively in the literature (Cox, 1958; Harville, 1975; Kempthorne, 1955, 1975 and references therein). The main contention is that the necessary assumptions for the randomization models are much less restrictive than for the ordinary normal-theory models. The second argument contends that the use of randomization ensures that the choice of design is not affected by any bias or preconceived notion on the part of the experimenter (Cox, 1958; Harville, 1975). Both arguments seem to be well

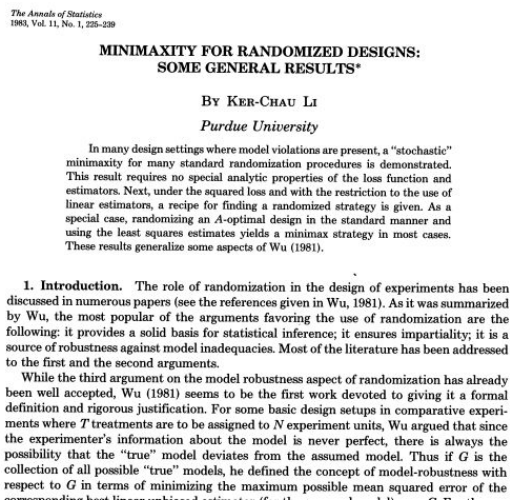

The Annals of Statistics
1983, Vol. 11, No. 1, 225–239

MINIMAXITY FOR RANDOMIZED DESIGNS: SOME GENERAL RESULTS*

BY KER-CHAU LI

Purdue University

In many design settings where model violations are present, a "stochastic" minimaxity for many standard randomization procedures is demonstrated. This result requires no special analytic properties of the loss function and estimators. Next, under the squared loss and with the restriction to the use of linear estimators, a recipe for finding a randomized strategy is given. As a special case, randomizing an A-optimal design in the standard manner and using the least squares estimates yields a minimax strategy in most cases. These results generalize some aspects of Wu (1981).

1. Introduction. The role of randomization in the design of experiments has been discussed in numerous papers (see the references given in Wu, 1981). As it was summarized by Wu, the most popular of the arguments favoring the use of randomization are the following: it provides a solid basis for statistical inference; it ensures impartiality; it is a source of robustness against model inadequacies. Most of the literature has been addressed to the first and the second arguments.

While the third argument on the model robustness aspect of randomization has already been well accepted, Wu (1981) seems to be the first work devoted to giving it a formal definition and rigorous justification. For some basic design setups in comparative experiments where T treatments are to be assigned to N experiment units, Wu argued that since the experimenter's information about the model is never perfect, there is always the possibility that the "true" model deviates from the assumed model. Thus if G is the collection of all possible "true" models, he defined the concept of model-robustness with respect to G in terms of minimizing the maximum possible mean squared error of the

JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION
2021, AHEAD-OF-PRINT, 1-14
https://doi.org/10.1080/01621459.2020.1863221

Taylor & Francis
Taylor & Francis Group

Open access

Minimax Efficient Random Experimental Design Strategies With Application to Model-Robust Design for Prediction

Timothy W. Waite[a] and David C. Woods[b]

[a] Department of Mathematics, University of Manchester, Manchester, UK; [b] Statistical Sciences Research Institute, University of Southampton, Southampton, UK

ABSTRACT

In game theory and statistical decision theory, a random (i.e., mixed) decision strategy often outperforms a deterministic strategy in minimax expected loss. As experimental design can be viewed as a game pitting the Statistician against Nature, the use of a random strategy to choose a design will often be beneficial. However, the topic of minimax-efficient random strategies for design selection is mostly unexplored, with consideration limited to Fisherian randomization of the allocation of a predetermined set of treatments to experimental units. Here, for the first time, novel and more flexible random design strategies are shown to have better properties than their deterministic counterparts in linear model estimation and prediction, including stronger bounds on both the expectation and survivor function of the loss distribution. Design strategies are considered for three important statistical problems: (i) parameter estimation in linear potential outcomes models, (ii) point prediction from a correct linear model, and (iii) global prediction from a linear model taking into account an L_2-class of possible model discrepancy functions. The new random design strategies proposed for (iii) give a finite bound on the expected loss, a dramatic improvement compared to existing deterministic exact designs for which the expected loss is unbounded. Supplementary materials for this article are available online.

Minimax Justifications

Suppose that for each patient, the effect of treatment is given by:

$$E_{it} = \alpha_t + g_i + \epsilon_{ut}$$

effect of treatment *t* on patient *i*

fixed effect of treatment *t*

patient effect

Independent mean-zero noise

Minimax Justifications

Let **g** = (g_1, g_2, …, g_N) be an assignment of patient effects to individuals.

Let **G** be the set of all assignments consistent with background knowledge.

Symmetry assumption: if **g** is in **G**, then so is every permutation of **g**.

Minimax Justifications

Theorem (Wu, 1981) The fully randomized design minimizes the maximum MSE of the estimate of the α_t over all possible values of **G**.

So the fully randomized design has the best worst-case efficiency.

Research Questions

Can these minimax arguments be generalized away from the linearity assumptions?

Generalized Minimax Justifications

Can these minimax arguments be generalized away from the linearity assumptions?

Yes, to some extent.

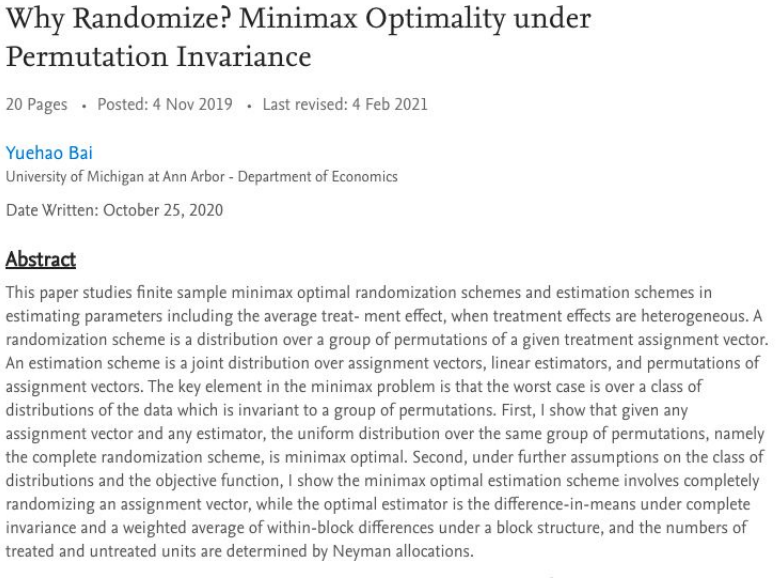

Why Randomize? Minimax Optimality under Permutation Invariance

20 Pages • Posted: 4 Nov 2019 • Last revised: 4 Feb 2021

Yuehao Bai
University of Michigan at Ann Arbor - Department of Economics

Date Written: October 25, 2020

Abstract

This paper studies finite sample minimax optimal randomization schemes and estimation schemes in estimating parameters including the average treat- ment effect, when treatment effects are heterogeneous. A randomization scheme is a distribution over a group of permutations of a given treatment assignment vector. An estimation scheme is a joint distribution over assignment vectors, linear estimators, and permutations of assignment vectors. The key element in the minimax problem is that the worst case is over a class of distributions of the data which is invariant to a group of permutations. First, I show that given any assignment vector and any estimator, the uniform distribution over the same group of permutations, namely the complete randomization scheme, is minimax optimal. Second, under further assumptions on the class of distributions and the objective function, I show the minimax optimal estimation scheme involves completely randomizing an assignment vector, while the optimal estimator is the difference-in-means under complete invariance and a weighted average of within-block differences under a block structure, and the numbers of treated and untreated units are determined by Neyman allocations.

Generalized Minimax Justifications

Can these minimax arguments be generalized away from the linearity assumptions?

Yes, to some extent.

But the nature of this justification is very different from that suggested by the tragic view!

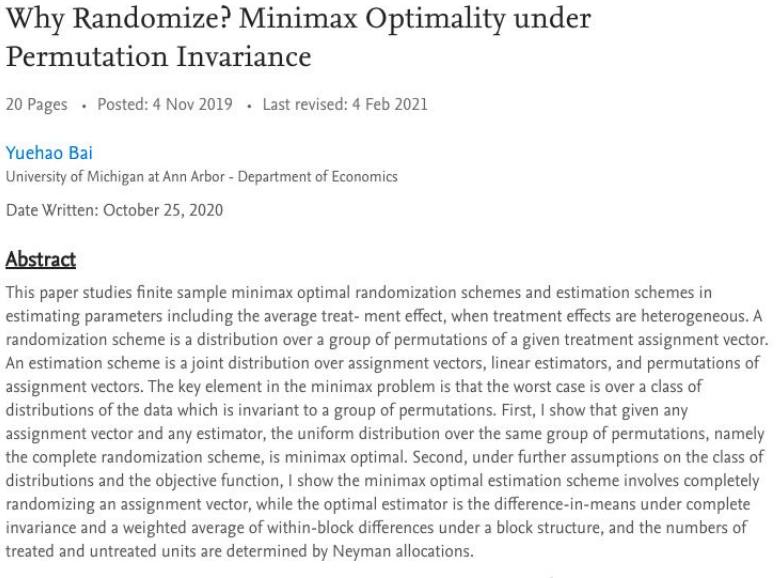

Why Randomize? Minimax Optimality under Permutation Invariance

20 Pages • Posted: 4 Nov 2019 • Last revised: 4 Feb 2021

Yuehao Bai
University of Michigan at Ann Arbor - Department of Economics

Date Written: October 25, 2020

Abstract

This paper studies finite sample minimax optimal randomization schemes and estimation schemes in estimating parameters including the average treat- ment effect, when treatment effects are heterogeneous. A randomization scheme is a distribution over a group of permutations of a given treatment assignment vector. An estimation scheme is a joint distribution over assignment vectors, linear estimators, and permutations of assignment vectors. The key element in the minimax problem is that the worst case is over a class of distributions of the data which is invariant to a group of permutations. First, I show that given any assignment vector and any estimator, the uniform distribution over the same group of permutations, namely the complete randomization scheme, is minimax optimal. Second, under further assumptions on the class of distributions and the objective function, I show the minimax optimal estimation scheme involves completely randomizing an assignment vector, while the optimal estimator is the difference-in-means under complete invariance and a weighted average of within-block differences under a block structure, and the numbers of treated and untreated units are determined by Neyman allocations.

Generalized Minimax Justifications

1. Suppose that you have *M* groups of patients. Within the groups the patients are clinically indistinguishable.

2. Suppose that for each group you are required to test the new treatment on exactly n_i patients in group *i.*

3. Then, uniformly randomizing the assignment is minimax optimal.

Generalized Minimax Justifications

1. Suppose that you have *M* groups of patients. Within the groups the patients are clinically indistinguishable.

2. Suppose that for each group you are required to test the new treatment on exactly n_i patients in group *i.*

3. Then, uniformly randomizing the assignment is minimax optimal.

But this is not in conflict with therapeutic obligation!

Generalized Minimax Justifications

1. Suppose that you have *M* groups of patients. Within the groups the patients are clinically indistinguishable.

2. Suppose that for each group you are required to test the new treatment on exactly n_i patients in group *i.*

3. Then, uniformly randomizing the assignment is minimax optimal.

4. The optimal n_i depend on worst-case outcome variances in the group.

Generalized Minimax Justifications

1. Suppose that you have *M* groups of patients. Within the groups the patients are clinically indistinguishable.

2. Suppose that for each group you are required to test the new treatment on exactly n_i patients in group *i.*

3. Then, uniformly randomizing the assignment is minimax optimal.

4. The optimal n_i depend on worst-case outcome variances in the group.

But this is also not in conflict with therapeutic obligation!

Research Questions

What is the **precise tradeoff** between individualized treatment and worst-case efficiency of the estimator?

e.g. If we want X% of patients to get individualized treatment, how many more trial participants (*N'*) would we need to achieve the same efficiency as a fully randomized RCT with *N*?

What is more important: giving (most) participants individualized treatment, or getting informative results with fewer participants?

Research Questions

What is the **precise tradeoff** between individualized treatment and worst-case efficiency of the estimator?

If we want X% of patients to get individualized treatment, how many more trial participants (*N'*) would we need to recruit in order to achieve the same efficiency as a fully randomized RCT with *N* patients?

Research Questions

What is the **precise tradeoff** between individualized treatment and worst-case efficiency of the estimator?

What is more important: giving (most) participants individualized treatment, or getting informative results with fewer participants?

Takeaway

If individualized treatment and estimation efficiency trade off, we should be able to say something quantitative about the nature of the trade-off.

The existence of **some** trade-off does not justify abandoning all therapeutic obligations.