Prespecification, Trial Fidelity, and the Causal Interpretation of RCT Results

2026 Annual Meeting of the Society for Clinical Trials
May 19, 2026
Phoenix, Arizona

Roger J. Lewis, MD, PhD
Senior Statistical Editor, JAMA
roger.lewis@jamanetwork.org

New from JAMA

JAMA+ Trials

https://jamanetwork.com/channels/trials

Highlighting the innovations and impact of clinical trial research to advance evidence-based medicine and public health

Editors/Board members: Alison Huang, MD, MAS; Roger J. Lewis, MD, PhD

Scan to sign up for email alerts or visit:
https://jamanetwork.com/channels/jamaplus/pages/trials

Typical Research Manuscript Primary Review Process

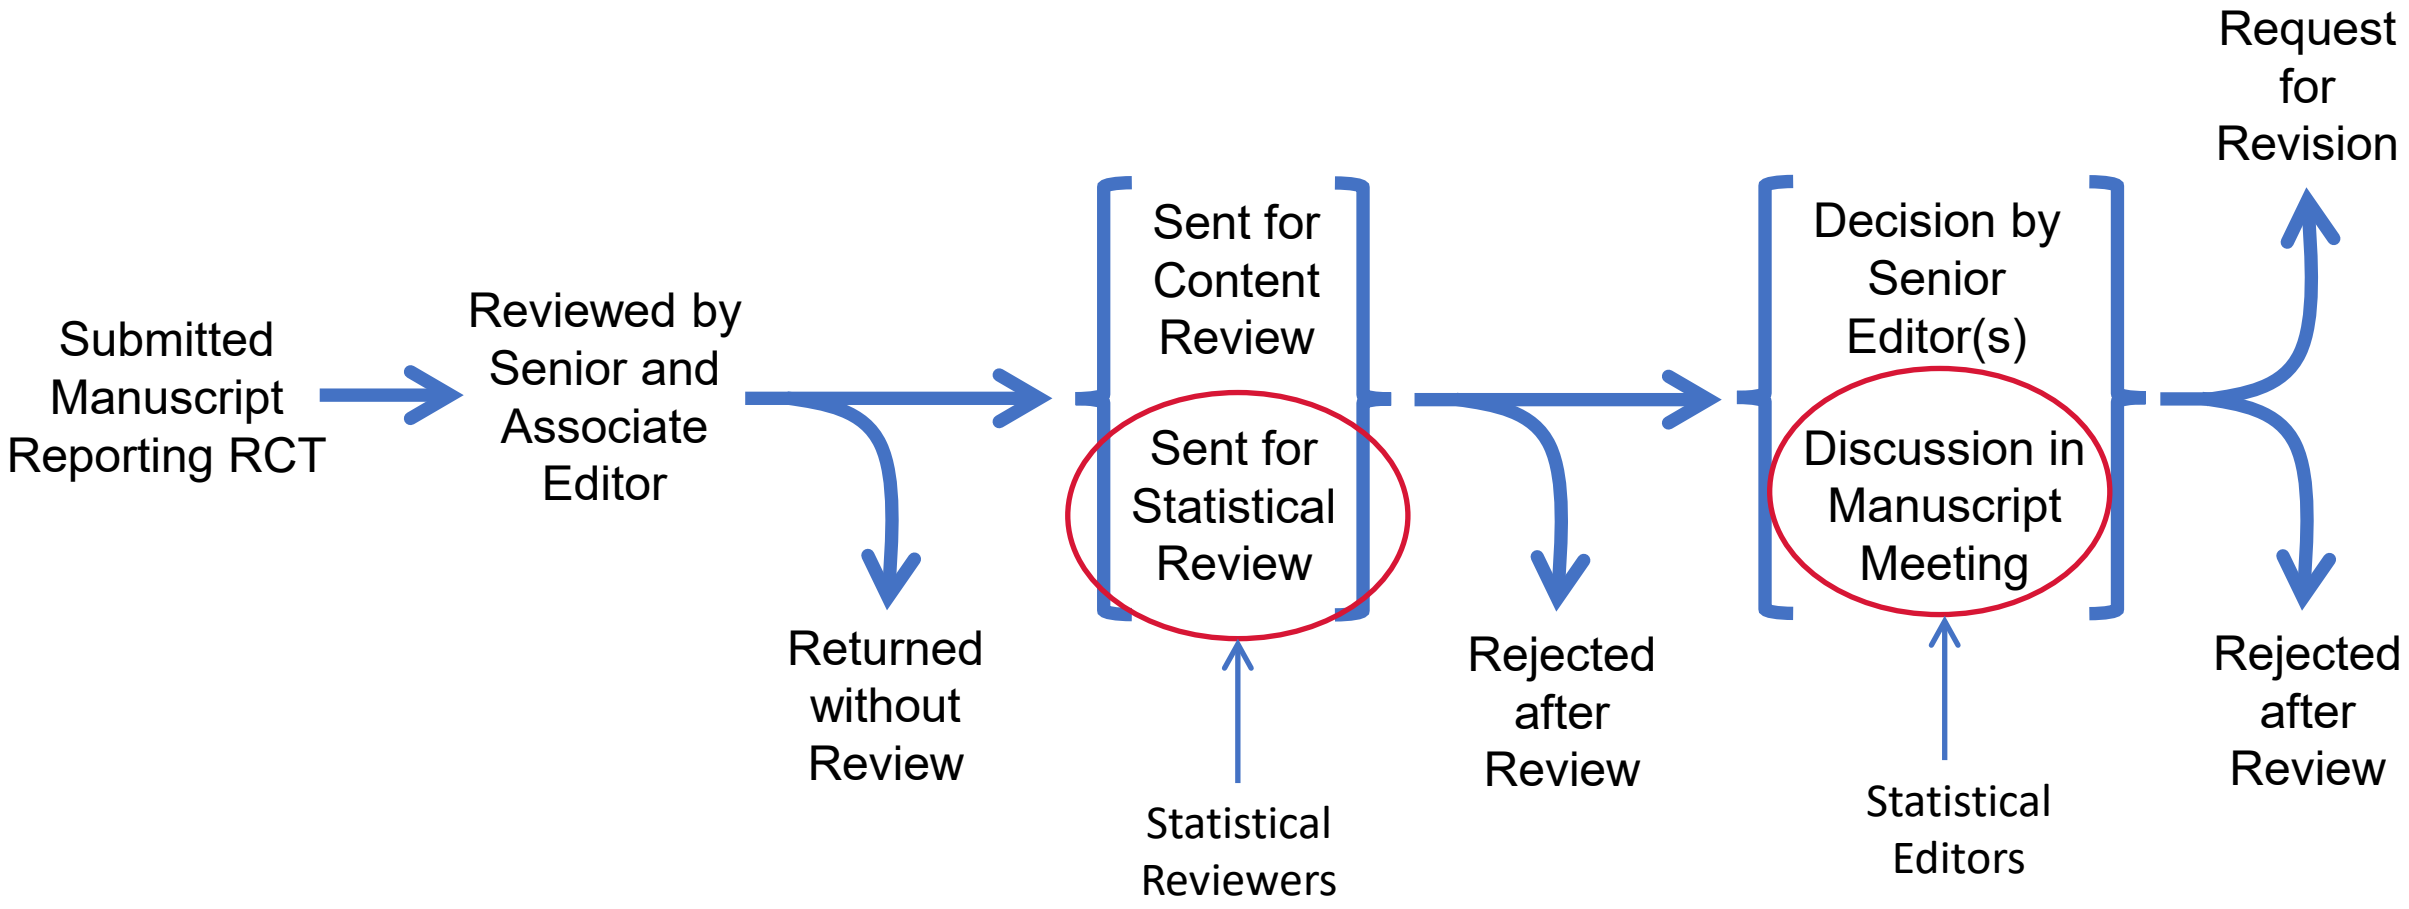

Statistical Editors and Reviewers

- **Statistical Editors:** Roger J. Lewis, MD, PhD, Issa J. Dahabreh, MD, ScD, Laura A. Hatfield, PhD, Patrick J. Heagerty, PhD, Romain Pirracchio, MD, PhD
- **Statistical Reviewers:** Lars W. Andersen, MD, PhD, Susan B. Brogly, PhD, Msc, Sylvie Chevret, MD, PhD, Yu-Han Chiu, MD, ScD, Sofia Dias, PhD, Valerie L. Durkalski-Mauldin, PhD, Thomas J. Hoffmann, PhD, MA, Matthias Kirch, MS, Lan Liu, PhD, David H. Mark, MD, Chaya Moskowitz, PhD, Christopher J. Paciorek, PhD, Michael J. Pencina, PhD, Jing Qian, PhD, Denise M. Scholtens, PhD, Jon Steingrimsson, PhD, Rajeshwari Sundaram, PhD, MStat, Joseph S. Wislar, MS

Prespecification, Rigor, and Reproducibility in Clinical Trials

- What are some of the reasons a clinical research result might not be rigorous?
  - Lack of study elements to reduce bias, e.g., randomization, double-blinded design, blinded outcome assessment, covariate adjustment
- What are some of the reasons a clinical research result might not be reproducible?
  - Lack of clear and detailed definitions of study protocol, procedures, and planned analysis
  - **Failure to <u>adhere</u> to a specific study protocol, procedures, and planned analysis, *especially* if the deviations are *non-random*, i.e., driven by the specific study data that have been observed**

Study Data and Random Variability

- Consider a randomized, double-blinded study evaluating a treatment X in patients with characteristics Y in modifying the primary outcome Z
- The observed data include the following components:
  - **The true effect of X on Z**
  - The effects of patient characteristics Y on Z (confounding)
  - Random variation of Z (noise)
- Signal and noise look the same in a dataset
- To avoid being misled by random variation, the analysis must not be driven by the observed data with its random components

Prespecified Analysis

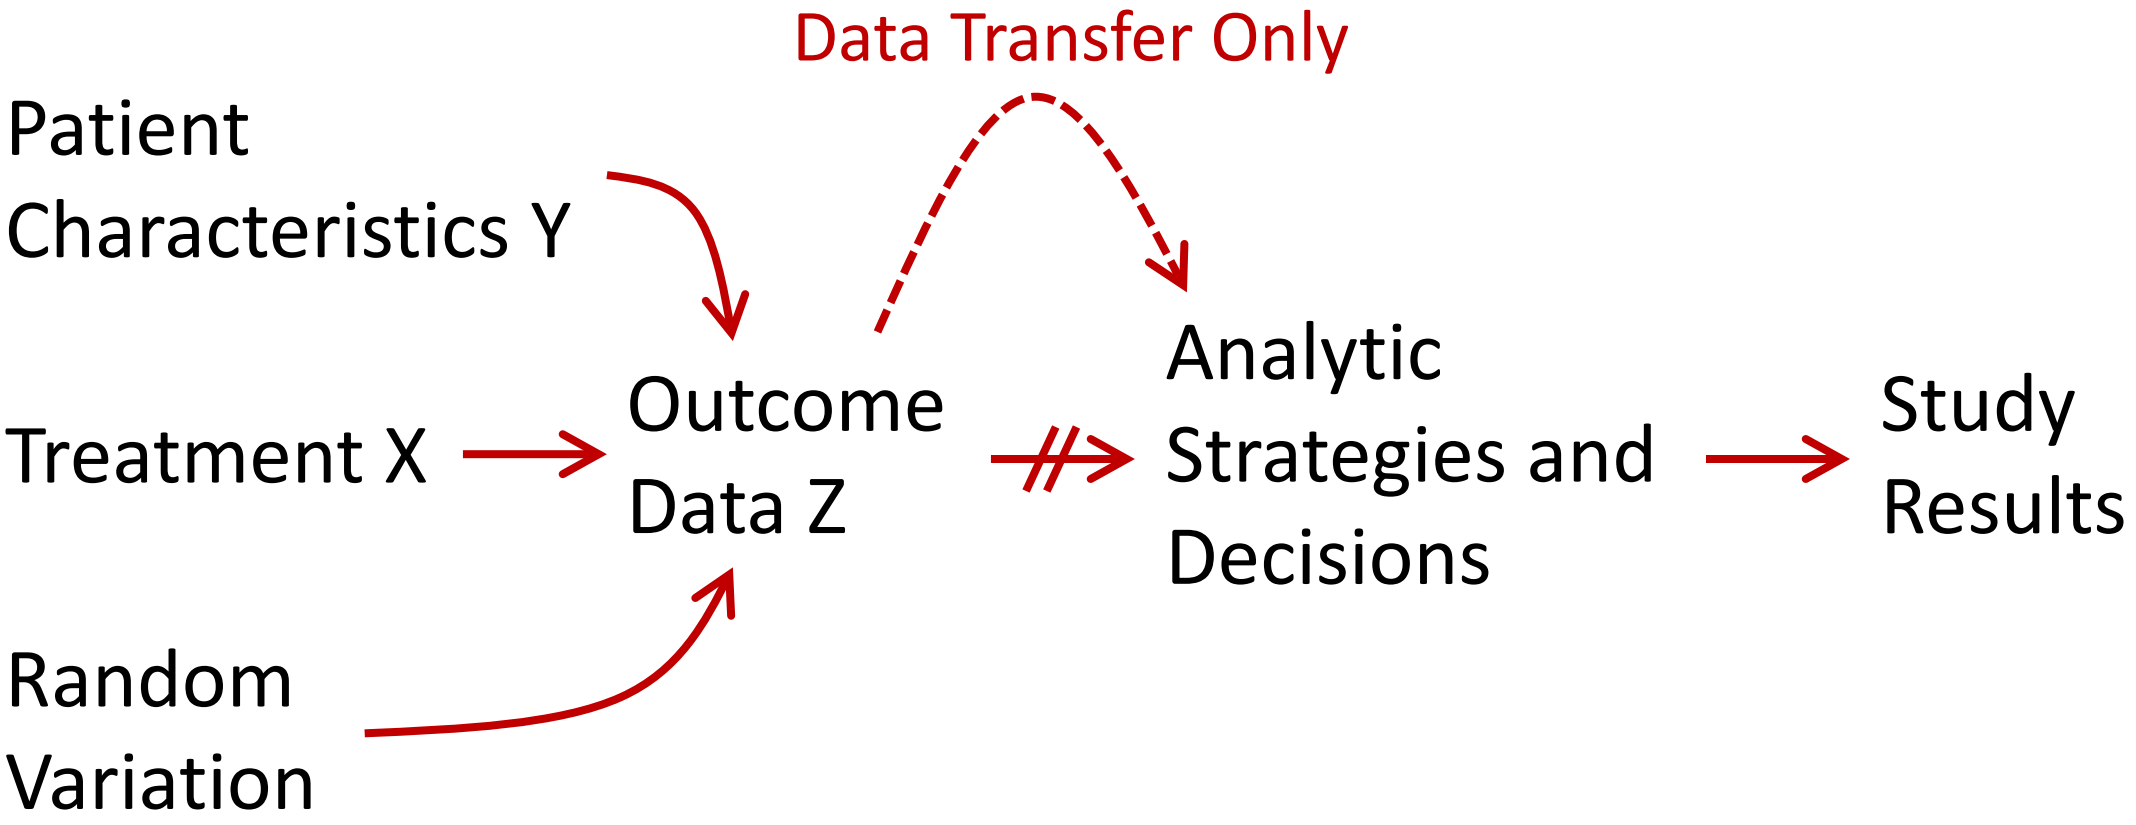

Non- or Inadequately Prespecified Analysis

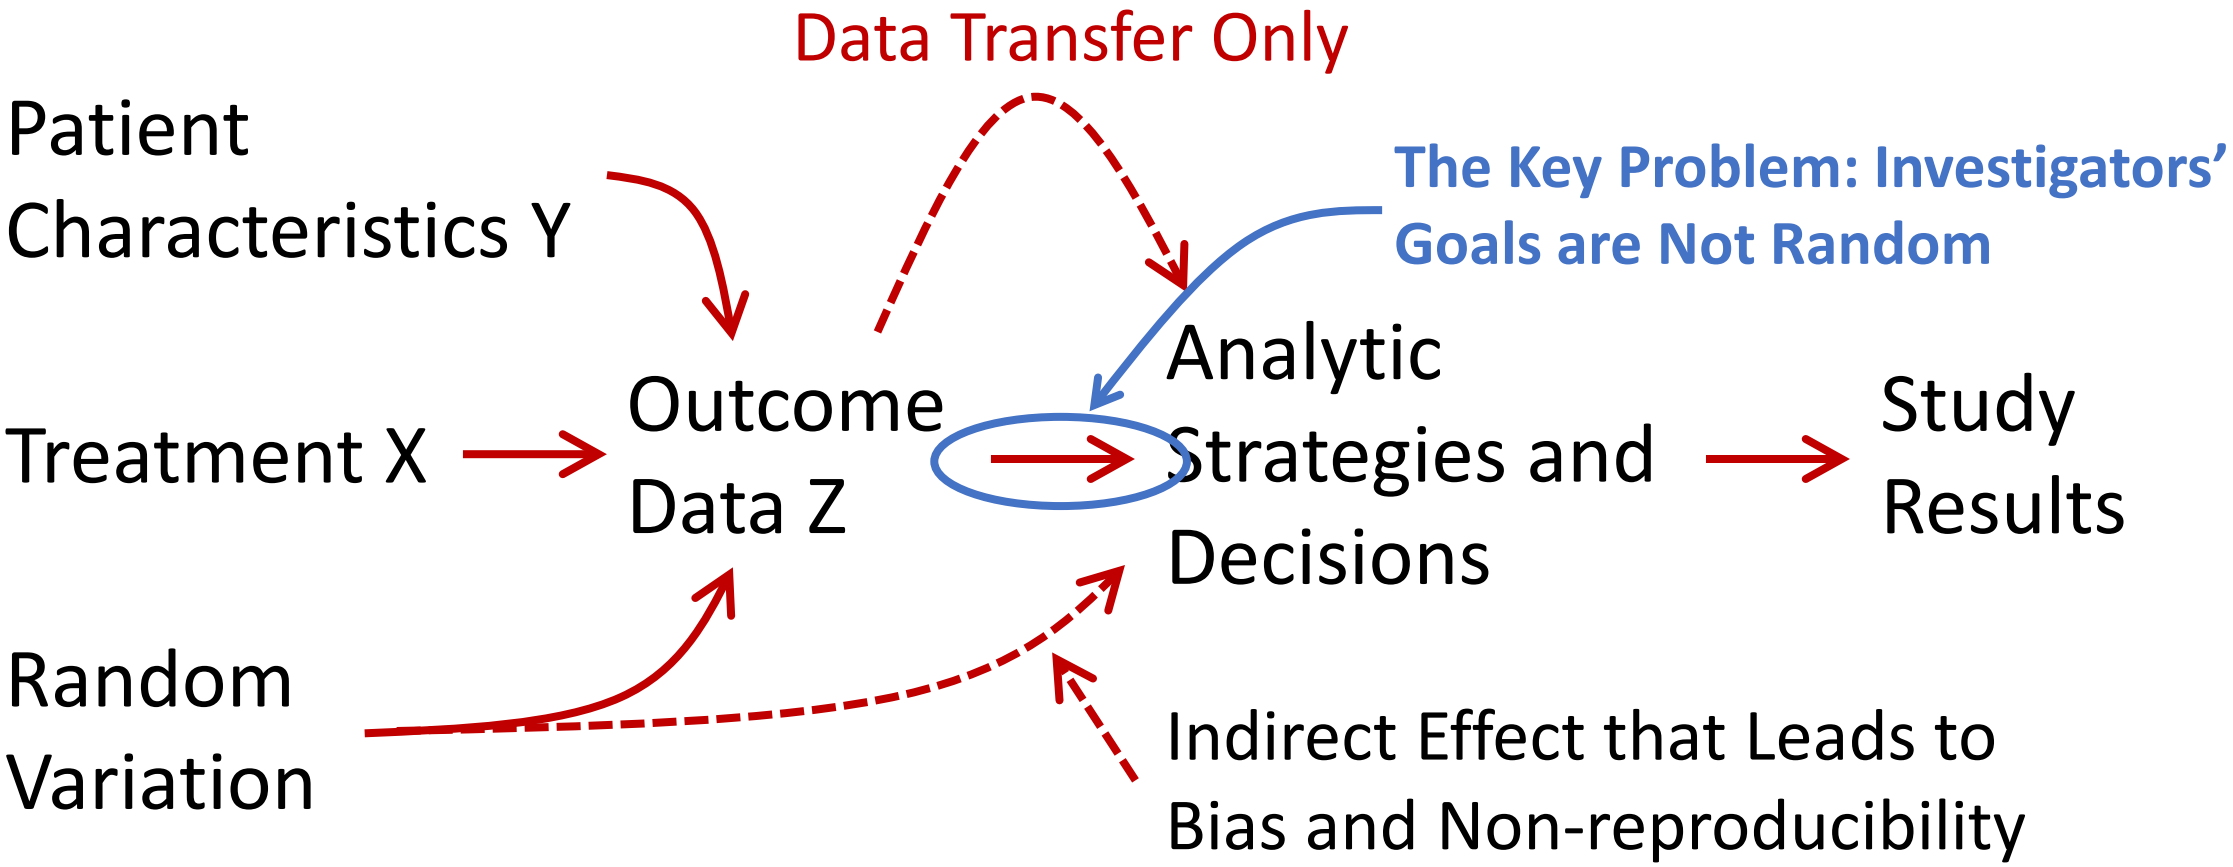

The Life and Evolution of a Typical Clinical Trial

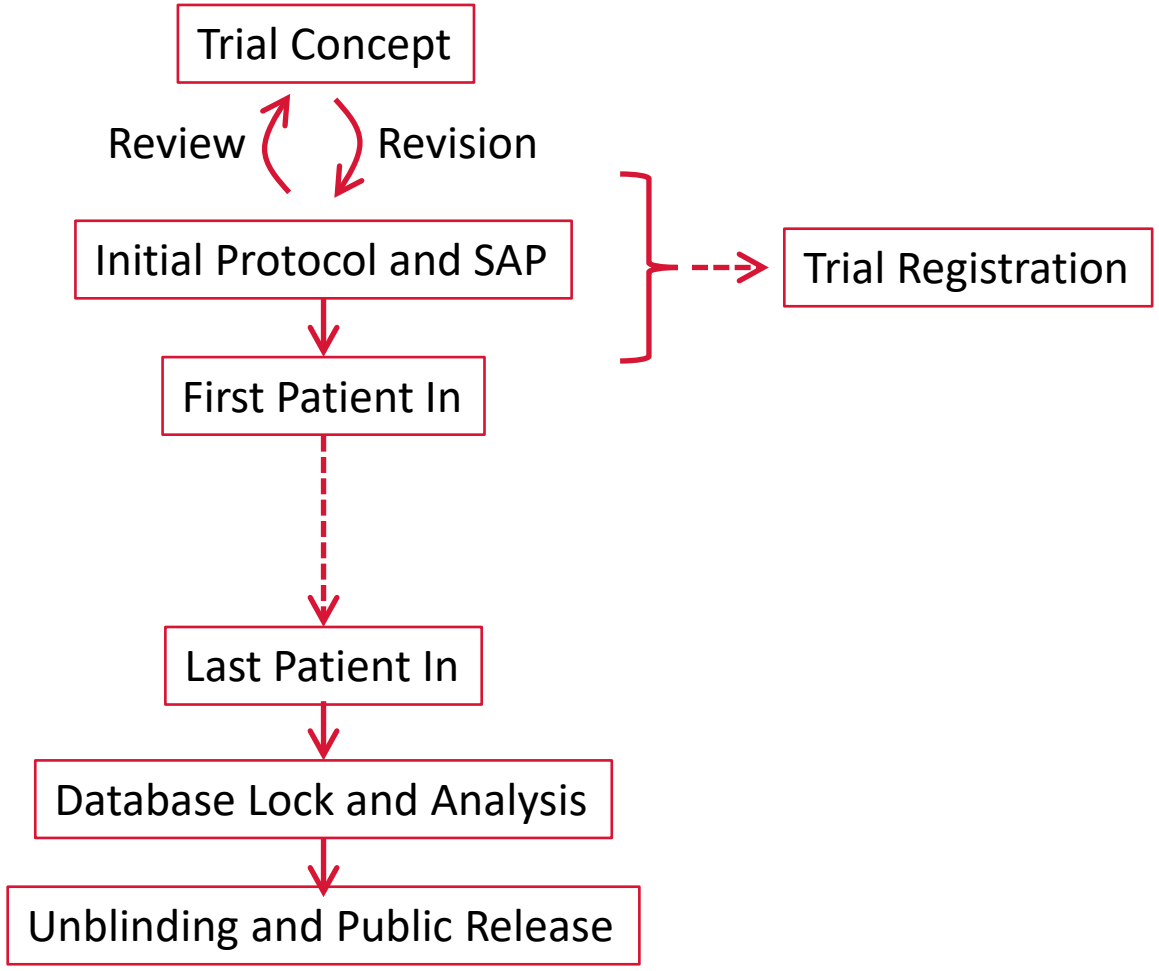

Purpose of Trial Registration

"The purpose of clinical trial registration is to prevent selective publication and selective reporting of research outcomes, to prevent unnecessary duplication of research effort, to help patients and the public know what trials are planned or ongoing into which they might want to enroll, and to help give ethics review boards considering approval of new studies a view of similar work and data relevant to the research they are considering. Retrospective registration, for example at the time of manuscript submission, meets none of these purposes."

https://www.icmje.org/recommendations/browse/publishing-and-editorial-issues/clinical-trial-registration.html

The Life and Evolution of a Typical Clinical Trial

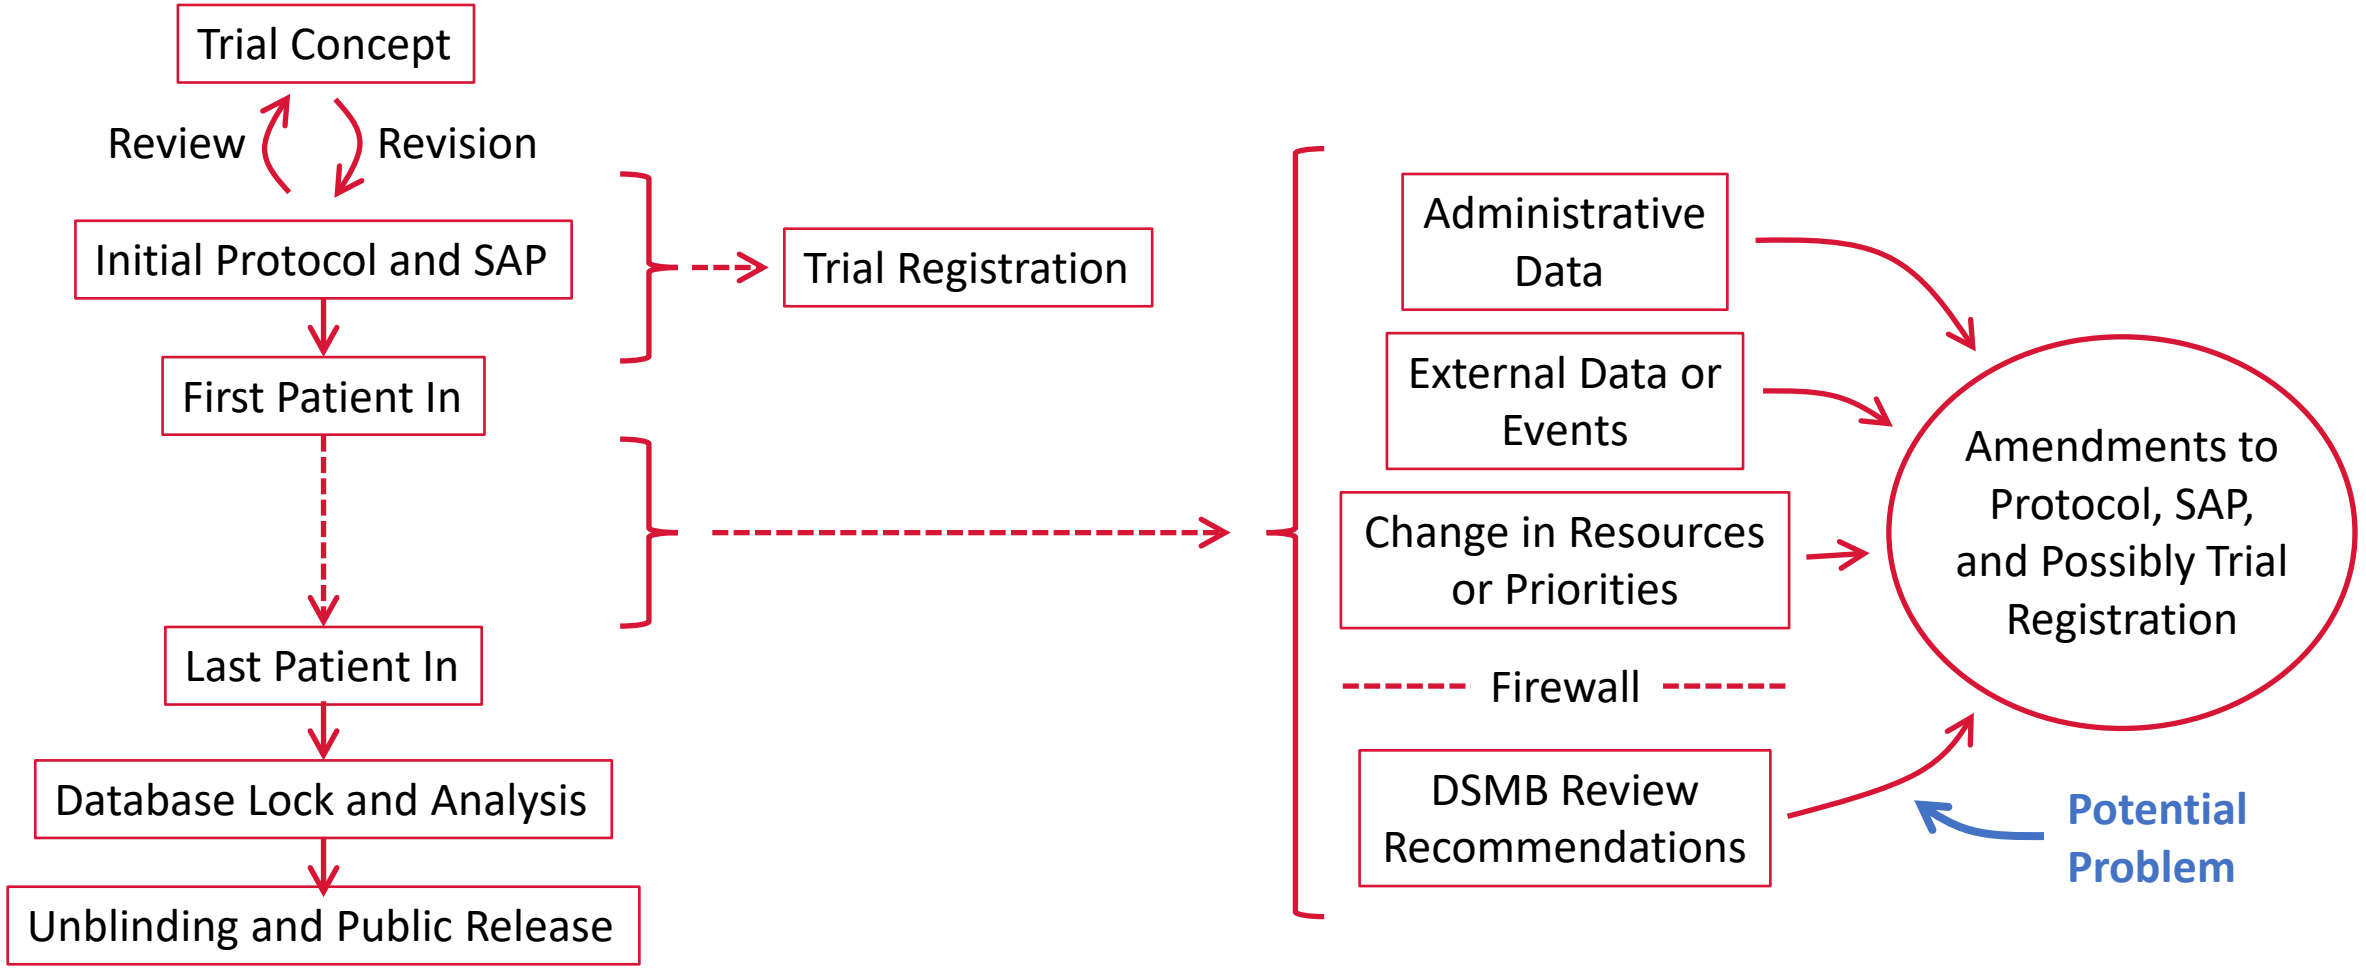

Conceptual Threats to Validity in RCT Reporting

- Bias from design decisions or changes informed by knowledge of unblinded data (e.g., by DSMB)

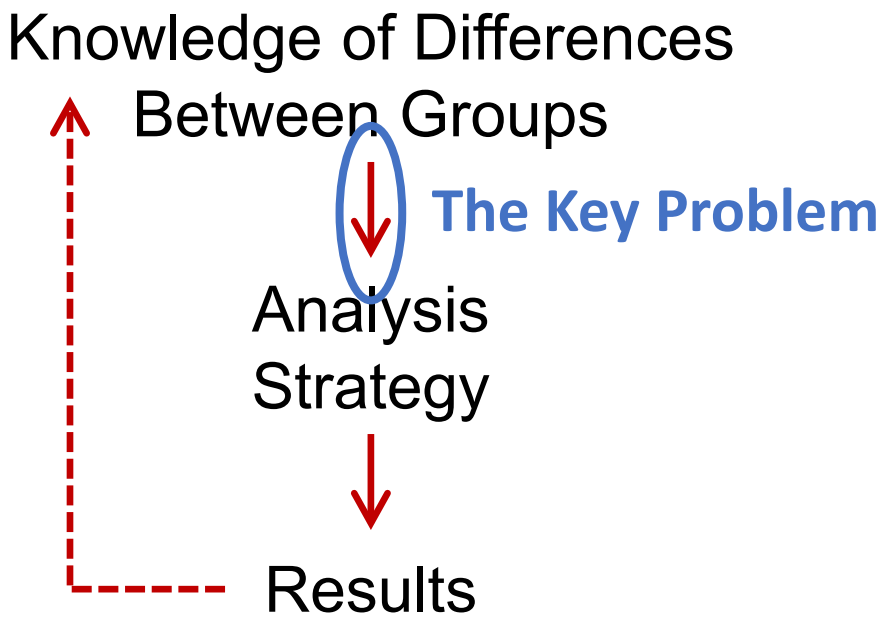

Conceptual Threats to Validity in RCT Reporting

- Bias from design decisions or changes informed by knowledge of unblinded data (e.g., by DSMB)
- Multiplicity, especially "hidden multiplicities"

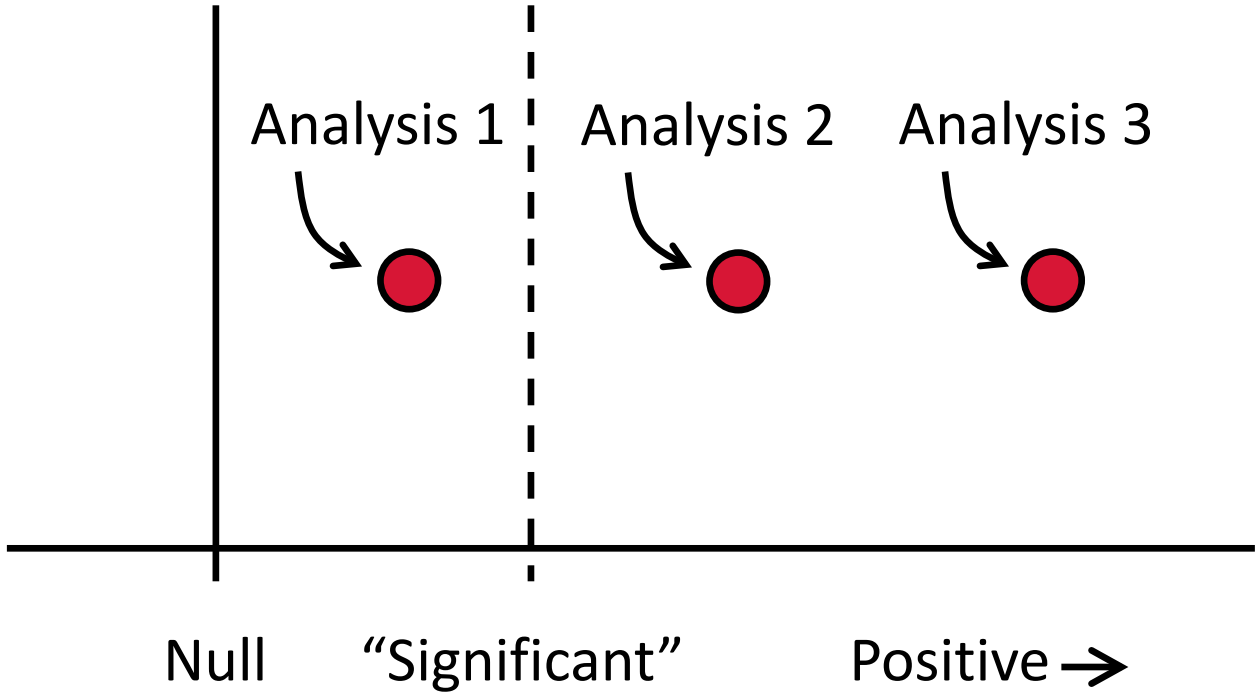

Conceptual Threats to Validity in RCT Reporting

- Bias from design decisions or changes informed by knowledge of unblinded data (e.g., by DSMB)
- Multiplicity, especially "**hidden multiplicities**"

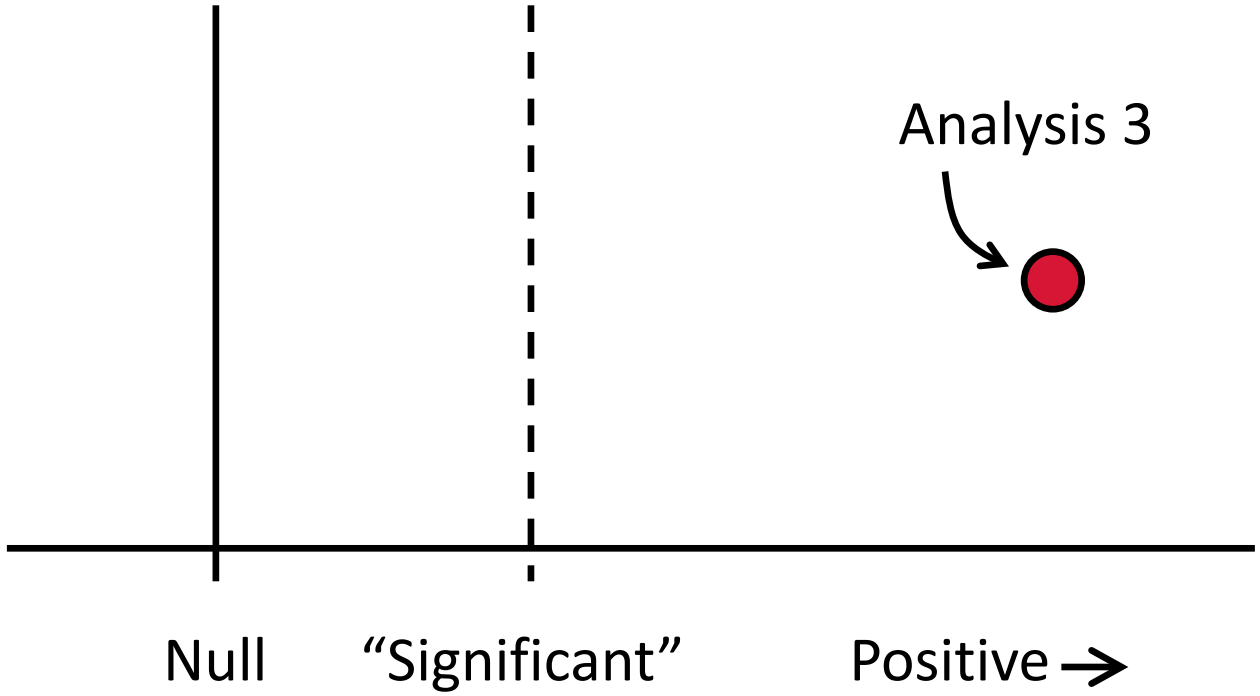

Consistency of Trial Documents

Design Publication	Manuscript: Methods	Study Protocol	SAP
Eligibility	Eligibility	**Eligibility**	Eligibility
Randomization	Randomization	**Randomization**	Randomization
Treatments	Treatments	**Treatments**	Treatments
Primary Endpoints	Primary Endpoints	**Primary Endpoints**	**Primary Endpoints**
Secondary Endpoints	Secondary Endpoints	**Secondary Endpoints**	**Secondary Endpoints**
Exploratory Endpoints	Exploratory Endpoints	Exploratory Endpoints	**Exploratory Endpoints**
Analysis Models	Analysis Models	Analysis Models	**Analysis Models**
Interim Analysis Plans	Interim Analysis Plans	Interim Analysis Plans	**Interim Analysis Plans**

Trial Registration

CONSORT

DSMB Charter
Primary Endpoints
Secondary Endpoints
Analysis Models
Interim Analysis Plans

Extending Inferences from Randomized Trials to Target Populations: Why This Is a Causal Problem

Yu-Han Chiu, MD, ScD

Department of Epidemiology

Brown University School of Public Health

May 19, 2026

Acknowledgements

- Drs. Issa Dahabreh, Lawson Ung
- Selected reference:
  - Dahabreh et al. (*Stat Med*, 2020)
  - Chiu and Dahabreh (arXiv:2209.08758)
  - Ung et al. (*Epidemiology*, 2025)
- Funding:
  - NIGMS R35GM154888 (Y-HC)
  - NHLBI R01HL136708, NLM R01LM013616, PCORI ME-2021C2-22365 (IJD)

Why RCT results may not be enough

- Trial participants are often
  - healthier
  - more adherent
  - more closely monitored
- Trials may be conducted
  - outside the US
  - in different healthcare systems
- Therefore, trial results **may not directly apply to** patients seen in routine clinical practice
- Key question: What is the treatment effect in **the population we actually care about**?

Possible Solutions

1. **Conduct a large pragmatic trial**
   - costly, slow, sometimes infeasible
2. **Emulate a target trial using observational data**
   - possible unmeasured confounding
   - treatment or outcome information may be unavailable or unreliable
3. **Extend inference from an RCT to a target population**

Extending Inference to a Target Population

Key idea: combine two sources of information

- A randomized trial with:
  - treatment
  - outcomes
  - baseline covariates
- A sample from the target population with:
  - baseline covariates

Often called **generalizability** or **transportability** analysis.

A running example

ARISTOTLE trial

- **Population**: Patients with atrial fibrillation and at least one risk factor for stroke
- **Randomized**: Apixaban vs warfarin
- **Outcome**: Stroke or systemic embolism
- **Main findings**: Apixaban was superior to warfarin

Target population of interest

Target population

Medicare beneficiaries seen in routine clinical practice who would have meet trial eligible criteria

Compared with eligible Medicare beneficiaries, trial participants were:

- younger
- less cardiovascular comorbidity

Goal: Extend inferences from ARISTOTLE trial to trial-eligible Medicare beneficiaries

Notation

- S: trial participation indicator
  - $S = 1$: randomized
  - $S = 0$: nonrandomized
- A: treatment assignment (apixaban vs warfarin)
- Y: stroke or systemic embolism (yes, no)
- X: baseline covariates

Assume perfect adherence, no missing data.

Composite data structure

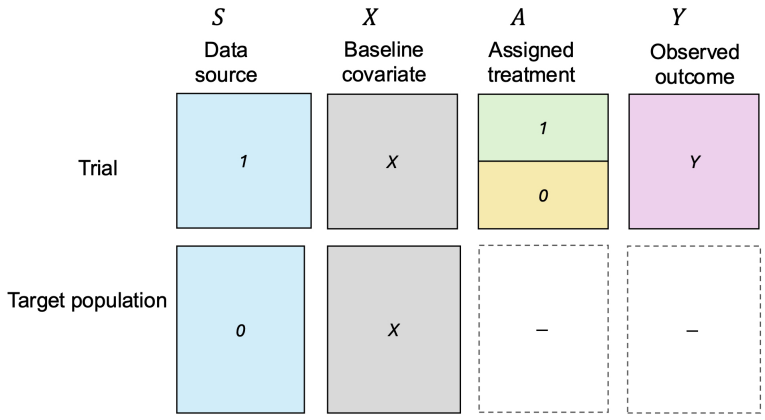

Treatment and outcome data come from the trial.
The target population contributes baseline covariates only.

Causal estimand

In our example, we are interested in:

- The average effect of apixaban versus warfarin among trial-eligible Medicare beneficiaries

This compares:

- outcomes if all trial-eligible Medicare beneficiaries received **apixaban**
- outcomes if all trial-eligible Medicare beneficiaries received **warfarin**

Key identification conditions

Within-trial conditions

- Trial participation does not affect the outcome except through treatment (no Hawthorne effect)
- Treatment is randomized within the trial, conditional on covariates.

Across-population conditions

- Conditional on baseline covariates, trial participants and the target population have similar expected outcomes under each treatment
- All covariate patterns in the target population are represented in the trial

Identification results

Under the conditions, the causal effect is identified by

$$E\left[E(Y \mid S=1, A=1, X) - E(Y \mid S=1, A=0, X) \,\middle|\, S=0\right]$$

- Learn treatment effect from trial ($S = 1$)
- Standardize to the target population ($S = 0$)

Why this is not merely standardization

At first glance, extending inference is just a standardization exercise.

Two deeper issues

1. Sampling the target population
2. Trial participation may itself affect outcomes

This is not merely a standardization problem; it depends on the underlying causal structure.

How to sample the target population?

In the ARISTOTLE example

- Target population:
  - trial-eligible Medicare beneficiaries
- Practical challenge:
  - an individual may meet eligibility criteria multiple times

A common but problematic shortcut

Sample individuals based on treatment initiation, for example:

- warfarin initiators
- direct oral anticoagulant initiators

Standardize trial results to individuals who **met eligiblity criteria and initiated treatment**.

This changes the target population sample to the **treated trial-eligible individuals**.

Can we standardize trial results to any convenient sample of the target population?

Can we standardize trial results to any convenient sample of the target population?

The short answer: No

Conditioning on treatment can induce bias

Inside the trial: Treatment assignment is randomized.

Outside the trial: Treatment reflects clinical indications, patient characteristics, unmeasured prognostic factors, ...

Once we **condition on treatment**, we created a selected subgroup whose treatment arose from different mechanisms inside and outside the trial.

This can induce (selection) bias.

Conditioning on treatment can induce bias

Inside the trial: Treatment assignment is randomized.

Outside the trial: Treatment reflects clinical indications, patient characteristics, unmeasured prognostic factors, ...

Once we **condition on treatment**, we created a selected subgroup whose treatment arose from different mechanisms inside and outside the trial.

This can induce (selection) bias.

Recommendation

Sample all trial-eligible individuals, regardless of actual treatment received.

The role of trial participation

- Classical transportability analyses typically treat trial participation S as a **selection indicator**
  - Who enters the trial ($S = 1$)
  - Who belongs to the target population ($S = 0$)
- Trial participation S may itself be part of the **intervention**.

Trial engagement effects

So far, we have assumed that trial participation affects outcomes only through treatment assignment.

- But trial participation may affect outcomes through paths not mediated by treatment assignment:
  - more intensive follow-up
  - protocol-driven monitoring
  - adherence support
  - behavioral changes related to study participation
- These are **trial engagement effects**
- **Key ideas**: **Trial participation** itself may become part of the **intervention**.

New causal estimands: with trial engagement effects

Treatment effect under trial protocol-based care.

This compares outcomes if all trial-eligible Medicare beneficiaries receive

- apixaban and under trial protcol-based care
- warfarin and under trial protocol-based care

Treatment effect under usual care: This compares outcomes if all trial-eligible Medicare beneficiaries receive

- apixaban and under usual care
- warfarin and under usual care

Identification under trial engagement effects

Treatment effect under usual care

Additional assumption

No interaction between:

- trial participation
- treatment

Then the estimand is identified by the same identification formula

$$E\Big[E(Y \mid S=1, A=1, X) - E(Y \mid S=1, A=0, X) \mid S=0\Big]$$

Key lesson

The same identification formula can represent different causal questions under different assumptions

What have we learned so far?

- Trial participation may act as:
  - a selection mechanism
  - an intervention
- The causal interpretation depends on:
  - how we define the target population
  - whether trial participation itself affects outcomes

Broader Implications

This framework underlies many important problems requiring combing randomized and nonrandomized informaiton:

- Comparing results across trials
- Benchmarking observational analyses against randomized trials
- Indirect treatment comparisons
- External control arms
- Causally interpretable meta-analysis

Summary

Extending inference from randomized trials requires:

- Carefully considering potential **trial engagement effects**
- Defining the **target estimand** and causal structure
- **Avoiding conditioning on treatment** when sampling the target population

In addition to trial-specific estimates, randomized trials may consider reporting treatment effects in clinically relevant target populations.

Take-home message

Extending inference is fundamentally a causal problem,

not merely a standardization exercise.

Backup slides

DAG under classical transportability

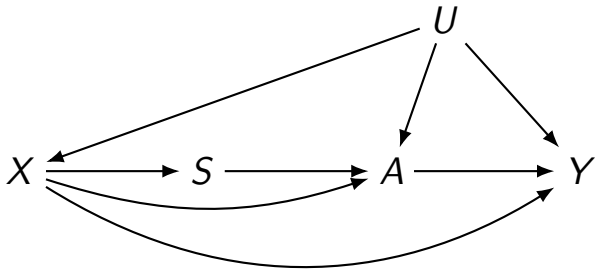

- No direct arrow from S to Y
- Trial participation affects treatment assignment, but not the outcome directly

DAG (trial engagement effect)

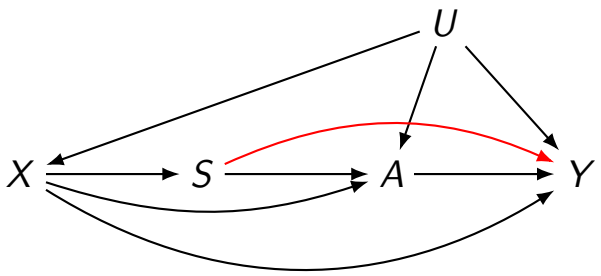

- Direct arrow from S to Y

Conditioning on treatment can induce bias

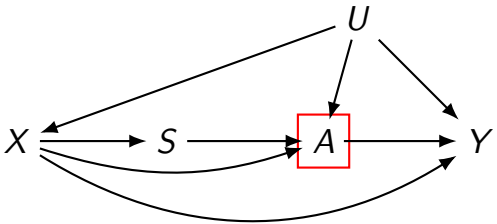

- In the trial, treatment is randomized
- Outside the trial, treatment reflects clinical indications and unmeasured prognostic factors.
- Conditioning on treatment opens a collider path $S \rightarrow A \leftarrow U \rightarrow Y$

Conditioning on treatment can induce bias

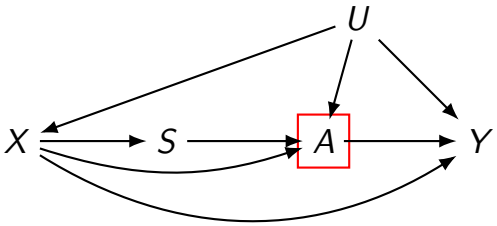

- In the trial, treatment is randomized
- Outside the trial, treatment reflects clinical indications and unmeasured prognostic factors.
- Conditioning on treatment opens a collider path $S \rightarrow A \leftarrow U \rightarrow Y$

Recommendation

Sample all trial-eligible individuals, regardless of actual treatment received.

Augmenting trials using external data: randomization-aware estimation

Issa Dahabreh
CAUSALab, Department of Epidemiology
Harvard T.H. Chan School of Public Health

3 ideas

- Using models and external data in clinical trials
- What to do when you are not feeling lucky (about external data)
- Trial augmentation and simulation-informed design

Optimal trial analyses for average treatment effects

Without assumptions beyond those justified by randomization,…

… **covariate-adjusted, model-robust** analyses are the most efficient approach for **marginal treatment effects** in **large samples**

What does this even mean ?

Optimal trial analyses

Covariate-adjusted = incorporate information on covariate–outcome association, e.g., through an outcome model

Model-robust = does not requiring the model to be "correct" to recover the treatment effect (consistent under misspecification)

Marginal treatment effect = population-averaged effect, as usual

Large samples = as the sample size of the trial grows to infinity

[we can also get improvements in finite samples]

The problem with these beautiful results

Robustness **does NOT** depend on model correctness

Efficiency **does** depend on model correctness

Trial budget runs out before optimality can be achieved

- It is hard to “learn” good models using only trial data
- We need a good model, we cannot get it from the trial alone

Possible solution: use external data

Use rich and voluminous external data (prior trials, registries, EHR)

External controls: replace an entire group with external data + assumptions

Augmented/hybrid/fusion designs: supplement one or more assignment groups with external data (± assumptions)

Do you feel lucky?

- Beyond rigorous data requirements

- Assumptions about alignment between the external and experimental data
  - No engagement effects
  - No confounding + transportability between sources
  - Overlap in distributions
  - Correct model specification

What to do when not feeling lucky?

- Take a chance!
- Use design (e.g., prospective non-experimental study; design experimental to "match" non-experimental data!)
- Use randomization, in smart way to reduce sample size requirements

Solution: trial augmentation

Exploit randomization **for robustness**

No assumptions beyond those justified by the design

Exploit external data for **efficiency improvement**

Massive data allow for flexible modeling, possibly adapted to the trial

Trial augmentation – sketch of method

"Learn" regression models in the external data; "fixed" before trial

- No multiple looks, data snooping, etc.
- Pre-specify the trial analysis

Make 2 component estimators…

- Traditional (efficient, model-robust) estimator
- New estimator using the externally developed models

… and develop a combined estimator

- Guaranteed to do no-worse than each component
- Adaptively decide how much to rely on the external data

Notation

X covariates

S indicator for information source (1 for trial; 0 for external data)

A treatment assignment

Y outcome

Y^a potential outcome under assignment a

Data structure

<table>
<tr><td rowspan="3">X</td><td rowspan="2">$S = 1$</td><td>$A = 1$</td><td rowspan="3">Y</td></tr>
<tr><td>$A = 0$</td></tr>
<tr><td>$S = 0$</td><td>$A = 0$</td></tr>
</table>

$(X_i, S_i = s, A_i, Y_i)$; $i = 1, \ldots, n_s$; n_s is sample size for $s = 0,1$

$S = 0 \Longrightarrow A = 0$

$n = n_0 + n_1$

Identifiability conditions

1. Consistency of potential outcomes: if $A = a$, the $Y^a = Y$, for $a = 0,1$.

2. No confounding in the trial: $Y^a \perp A|(X, S = 1)$, for $a = 0,1$.

3. Positive probability of assignment in the trial: For $a = 0,1$, $\Pr[A = a|X = x, S = 1] > 0$, for each covariate level x with positive density in the trial.

In the trial, association is causation

Under conditions 1 – 3, the potential outcome mean under assignment to a is identifiable:

$$\mathrm{E}[Y^a \mid S = 1] = \mathrm{E}[Y^a \mid S = 1, A = a] = \mathrm{E}[Y \mid S = 1, A = a]$$

Estimation in the trial

$$\widehat{\phi}_a = \left(\sum_{i=1}^{n} S_i\right)^{-1} \sum_{i=1}^{n} S_i \left[\frac{\mathbf{1}(A_i = a)}{\widehat{e}_a(X_i)} \left(Y_i - \widehat{g}_a(X_i)\right) + \widehat{g}_a(X_i)\right]$$

$e_a(X) = \Pr[A = a|X, S = 1]$

$\hat{g}_a(X)$ is an estimator for $\mathrm{E}[Y|X, S = 1, A = a]$

We can estimate the treatment effect as $\hat{\tau}(\hat{g}) = \hat{\varphi}_1 - \hat{\varphi}_0$

https://arxiv.org/abs/2406.17971 (*in-press*, Biometrics 2026)

A family of estimators

Consider the class of estimators: $\hat{\tau}(h) = \hat{\psi}_1(h) - \hat{\psi}_0(h)$,

$$\widehat{\psi}_a(\pi, h) = \left(\sum_{i=1}^{n} S_i \right)^{-1} \sum_{i=1}^{n} S_i \left[\frac{\mathbf{1}(A_i = a)}{\pi(X_i)} \{Y_i - h(X_i)\} + h(X_i) \right]$$

$\pi(X)$, $h(X)$: functions of covariates

In randomized trials the assignment probability, $\pi(X)$, is under investigator control (known).

https://arxiv.org/abs/2406.17971 (*in-press*, Biometrics 2026)

2 options

Ignore the external data when learning $h(X)$; use a trial-only estimator, $\hat{\tau}(\hat{g})$.

Use the external data and data-adaptive methods (e.g., machine learning/AI) to learn a "good" $h(X)$, say, h^*, and construct a new estimator, $\hat{\tau}(h^*)$.

https://arxiv.org/abs/2406.17971 (*in-press*, Biometrics 2026)

Linear combination of estimators

Consider the class of estimators indexed by λ,

$$\hat{\tau}(\lambda) = \lambda\hat{\tau}(\hat{g}) + (1-\lambda)\hat{\tau}(h^*) = \hat{\tau}(\lambda\hat{g} + (1-\lambda)h^*)$$

The asymptotic variance of $\hat{\tau}(\lambda)$ is minimized at $\lambda^* = \dfrac{v_g^2 - v_{g,h^*}}{v_g^2 + v_{h^*}^2 - 2v_{g,h^*}}$

Here, v_g^2 is the asymptotic variance of $\hat{\tau}(\hat{g})$, $v_{h^*}^2$ is the asymptotic variance of $\hat{\tau}(h)$, and v_{g,h^*} is their asymptotic covariance.

https://arxiv.org/abs/2406.17971 (*in-press*, Biometrics 2026)

Combined estimator

We obtain a combined estimator $\hat{\tau}(\hat{\lambda}^*)$ by replacing the asymptotic variance—covariance terms in the expression for λ^* with corresponding estimators.

The combined estimator $\hat{\tau}(\hat{\lambda}^*)$ is **consistent** for the target parameter, **asymptotically normal**, and has **asymptotic variance that is no greater than the lowest** of $\hat{\tau}(\hat{g})$ or $\hat{\tau}(h^*)$.

https://arxiv.org/abs/2406.17971 (*in-press*, Biometrics 2026)

Data adaptation: how much weight for the external data

With sample-splitting and cross-fitting:

- Finite sample unbiasedness (for linear contrasts)
- Consistent, asym. normal
- Nominal alpha
- Nominal coverage
- Exact tests possible (for the purists!)

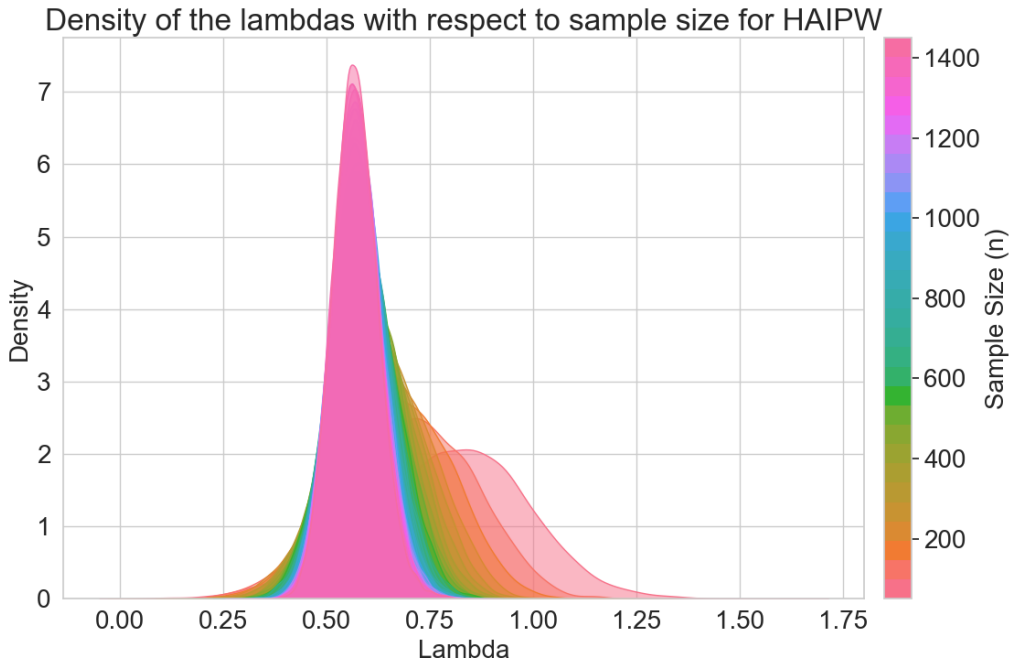

Data adaptation: how much weight for the external data

With sample-splitting and cross-fitting:

- Finite sample unbiasedness (for linear contrasts)
- Consistent, asym. normal
- Nominal alpha
- Nominal coverage
- Exact tests possible (for the purists!)

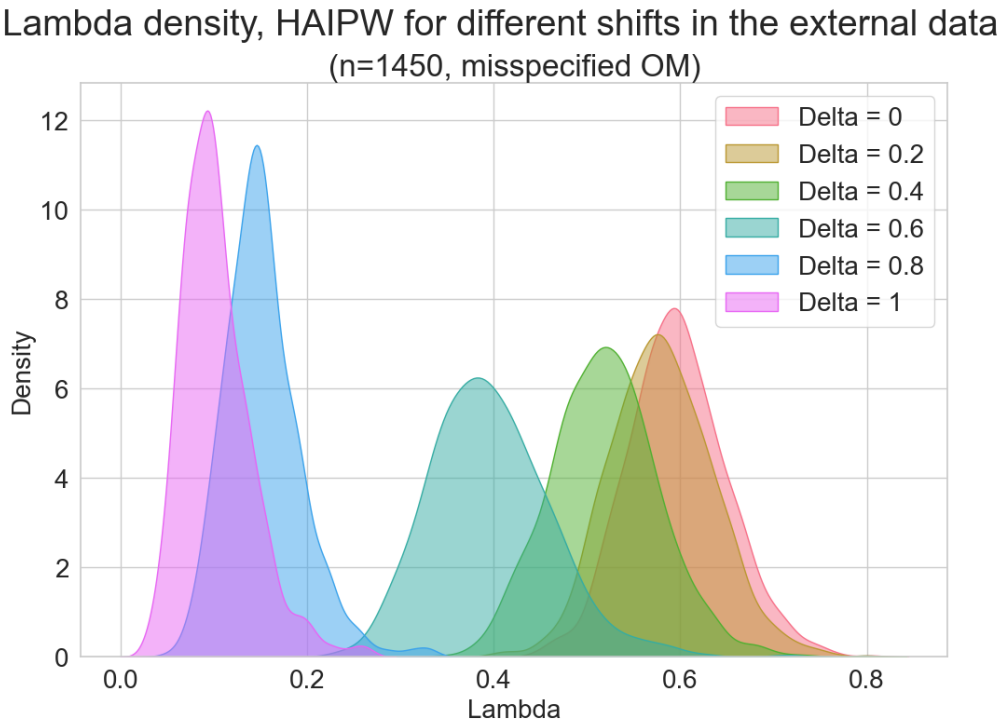

Comparison of efficiency

Data from the CASS trial

Using the external cohort for augmentation

Subsampling the trial arms to examine efficiency gains

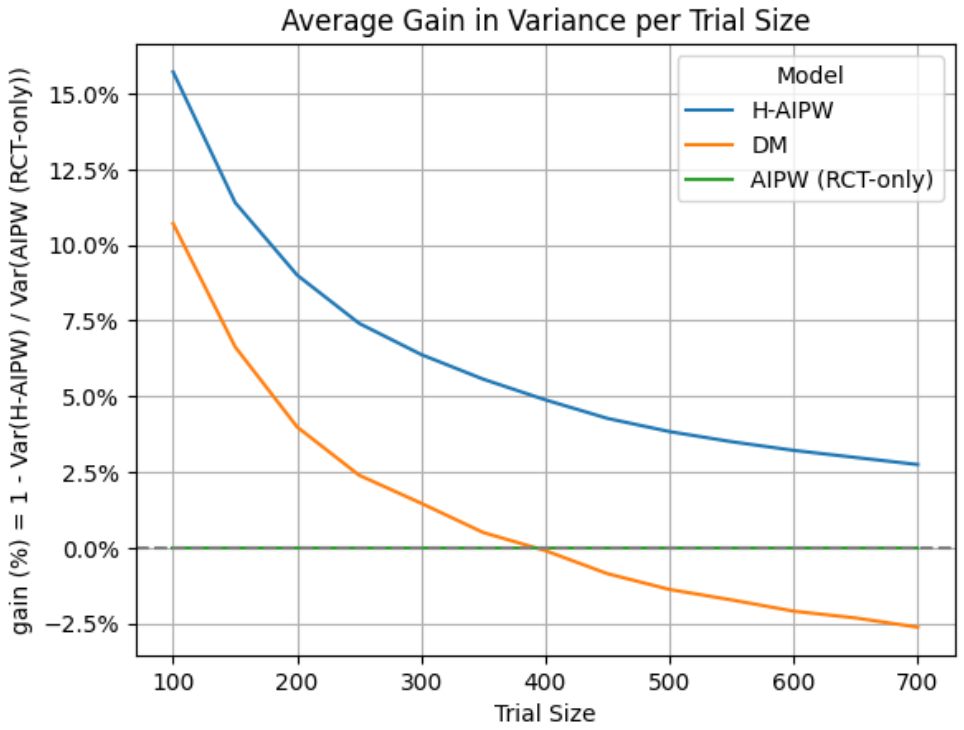

How do we plan a study like this?

Simulation-informed design

- Generative models can produce realistic data
- Implement the methods to select modeling approaches, estimators, optimization strategies, regularization for tuning, handling of complications (e.g., missing data)

Examine performance of as many designs as needed

- Classical approach: power, coverage, false positive error rate

Pre-specify all analyses

Selected references

- Dahabreh, 2025. Trial emulation, simulation, and augmentation using electronic health records and generative ai. *NEJM AI*, *2*(10), p.AIe2500894.
- De Bartolomeis, 2025. Efficient randomized experiments using foundation models. *NeurIPS*.
- Karlsson, 2024. Robust integration of external control data in randomized trials (in-press, *Biometrics 2026*; *arXiv:2406.17971*).
- Karlsson, 2025. Robust estimation of heterogeneous treatment effects in randomized trials leveraging external data. AISTATS; *arXiv preprint arXiv:2507.03681*.