Optimized Learning While Doing: The REMAP-CAP Adaptive Platform Trial

Derek C. Angus, MD, MPH

Learning While Doing

- Must do two things simultaneously
  - Do: Treat patients as well as possible
  - Learn: Find out what therapies help

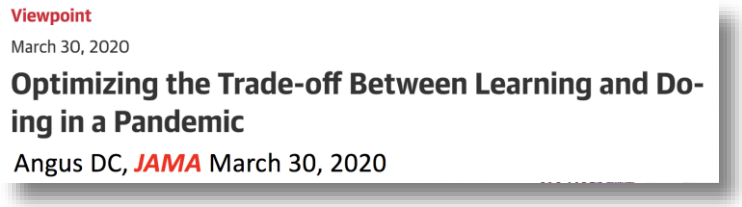

Learning While Doing

- Must do two things simultaneously
  - Do: Treat patients as well as possible
  - Learn: Find out what therapies help

- Framed as a (potentially false) choice

Viewpoint

March 30, 2020

Optimizing the Trade-off Between Learning and Doing in a Pandemic

Angus DC, *JAMA* March 30, 2020

Learning While Doing

- Must do two things simultaneously
  - Do: Treat patients as well as possible
  - Learn: Find out what therapies help

- Framed as a (potentially false) choice

- Classic dilemma in decision-making under uncertainty
  - The 'exploration/exploitation trade-off'
    - James March, Org Sci 1991

- The (elusive) solution is an integrated approach
  - Find the optimal balance to treat patients as well as possible and learn as fast as possible

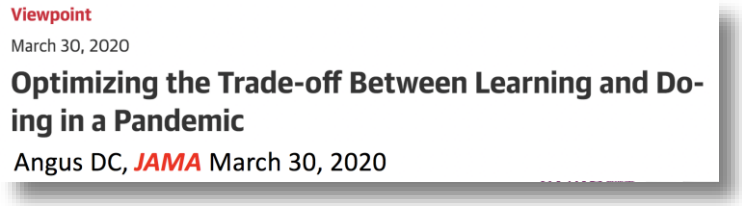
Viewpoint

March 30, 2020

Optimizing the Trade-off Between Learning and Doing in a Pandemic

Angus DC, *JAMA* March 30, 2020

Outside medicine …

- Exploration/exploitation (or ‘Learning While Doing’) is everywhere …
  - Cornerstone of decision-making under uncertainty
- Complex Adaptive Systems research in multiple disciplines
  - Organization science, mathematics, evolutionary biology, economics, social sciences
- Artificial intelligence
  - Reinforcement learning
  - Multi-arm bandits, Markov decision processes, policy evaluations, etc.
- All disciplines exploring the optimal trade-off …

Inside medicine …

- ‘Doing’ (practice) and ‘Learning’ (research) are separate
  - Many reasons, including Belmont Report
  - Separate organizations, cultures, people, funding, procedures, and goals

- Consequence: no one really empowered to find the optimal trade-off
  - Always true, but particularly obvious during a pandemic

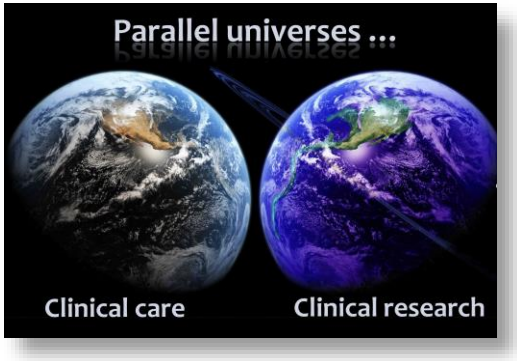

Best learning tool is the RCT, but 3 major challenges in a pandemic …

- Randomization is very uncomfortable
  - Physician feels responsible for patient outcomes, consequences are immediate
  - Physician feels less responsible for research, consequences are remote
- RCTs are very cumbersome
  - Slow to start
  - Intrusive to execute
- Little coordination in the clinical research enterprise
  - >100 RCTs registered for HCQ; few likely to be completed
  - AMCs bombarded with 100s of requests to participate in trials; no national or global prioritization

3 solutions from the clinical research enterprise, designed to ‘lean in’ to the realities of clinical care …

- Make randomization more comfortable
  - Multiple arms, only one is control
  - Adaptive randomization, preferentially assign to best therapy over time

- Make entry into clinical trials ‘1-stop shopping’
  - Simplify interface between clinical practice and clinical research
  - Use master protocols with standard entry criteria, outcomes, etc.
  - Essentially, combine trials/study questions

- Sacrifice ‘sacred cows’ of research
  - Don’t let perfection be the enemy of the good
  - Ex. placebo probably overrated in a pandemic; added rigor not worth the burden

REMAP-CAP Executive Summary

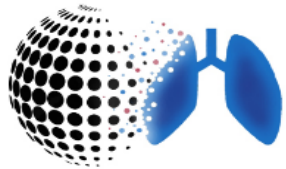

- A global adaptive platform trial

- Designed to determine best treatment for severe pneumonia
  - Randomizes multiple interventions simultaneously, nested within domains
  - Uses a multifactorial Bayesian inference model
  - Uses response-adaptive randomization

- Assesses both interpandemic AND pandemic forms of pneumonia
  - Pre-set rules to switch into pandemic mode

- Entered pandemic mode (termed ‘REMAP-COVID’) in February 2020

Adaptive Platform Trials

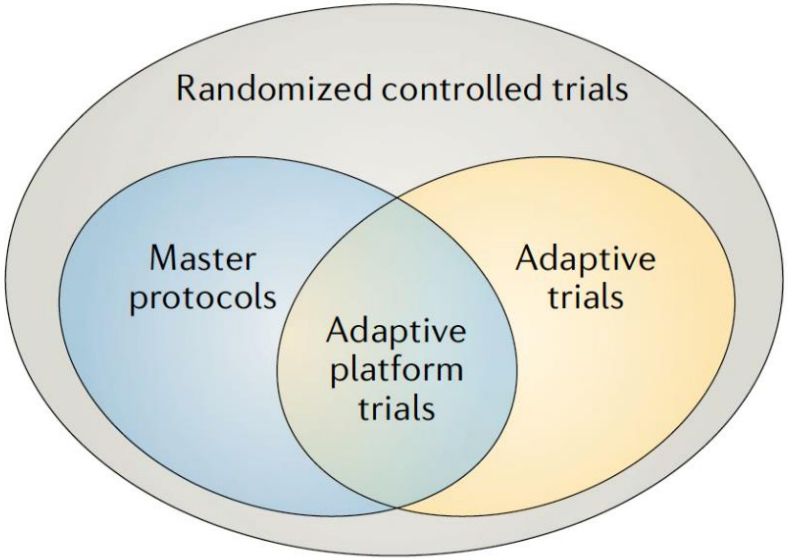

- **Typically, have focused on pre-approval space**
  - Emphasis on efficiency with (very) small sample sizes
  - Different therapies 'graduate' to next phase while trial continues

Woodcock and Lavange. *NEJM* 2017

Adaptive Platform Trials Coalition. *Nature Drug Discovery* 2019

New Breast Cancer Results Illustrate Promise and Potential of I-SPY 2 Trial

Trial Identifies Breast Cancer Patients Likely to Benefit from Experimental Drug

By Elizabeth Fernandez on April 07, 2014

Physicians

I-SPY 2 is a collaborative research effort that uses genetic and biological markers from individual patient's tumors to screen several promising new treatments simultaneously and allows doctors to quickly measure the effectiveness of the treatment prior to removing the tumor.

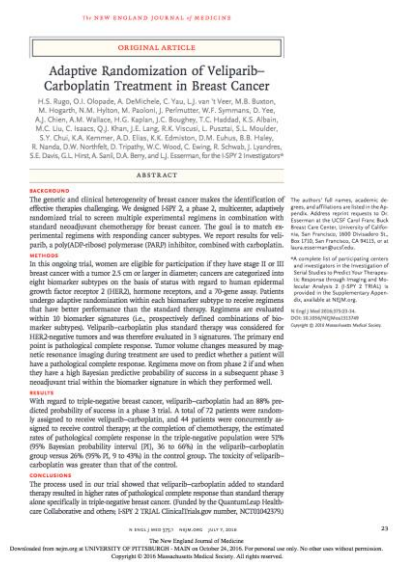

The NEW ENGLAND JOURNAL of MEDICINE

ORIGINAL ARTICLE

Adaptive Randomization of Veliparib–Carboplatin Treatment in Breast Cancer

H.S. Rugo, O.I. Olopade, A. DeMichele, C. Yau, L.J. van 't Veer, M.B. Buxton, M. Hogarth, N.M. Hylton, M. Paoloni, J. Perlmutter, W.F. Symmans, D. Yee, A.J. Chien, A.M. Wallace, H.G. Kaplan, J.C. Boughey, T.C. Haddad, K.S. Albain, M.C. Liu, C. Isaacs, Q.J. Khan, J.E. Lang, R.K. Viscusi, L. Pusztai, S.L. Moulder, S.Y. Chui, K.A. Kemmer, A.D. Elias, K.K. Edmiston, D.M. Euhus, B.B. Haley, R. Nanda, D.W. Northfelt, D. Tripathy, W.C. Wood, C. Ewing, R. Schwab, J. Lyandres, S.E. Davis, G.L. Hirst, A. Sanil, D.A. Berry, and L.J. Esserman, for the I-SPY 2 Investigators*

ABSTRACT

BACKGROUND

The genetic and clinical heterogeneity of breast cancer makes the identification of effective therapies challenging. We designed I-SPY 2, a phase 2, multicenter, adaptively randomized trial to screen multiple experimental regimens in combination with standard neoadjuvant chemotherapy for breast cancer. The goal is to match experimental regimens with responding cancer subtypes. We report results for veliparib, a poly(ADP-ribose) polymerase (PARP) inhibitor, combined with carboplatin.

METHODS

In this ongoing trial, women are eligible for participation if they have stage II or III breast cancer with a tumor 2.5 cm or larger in diameter; cancers are categorized into eight biomarker subtypes on the basis of status with regard to human epidermal growth factor receptor 2 (HER2), hormone receptors, and a 70-gene assay. Patients undergo adaptive randomization within each biomarker subtype to receive regimens that have better performance than the standard therapy. Regimens are evaluated within 10 biomarker signatures (i.e., prospectively defined combinations of biomarker subtypes). Veliparib–carboplatin plus standard therapy was considered for HER2-negative tumors and was therefore evaluated in 3 signatures. The primary end point is pathological complete response. Tumor volume changes measured by magnetic resonance imaging during treatment are used to predict whether a patient will have a pathological complete response. Regimens move on from phase 2 if and when they have a high Bayesian predictive probability of success in a subsequent phase 3 neoadjuvant trial within the biomarker signature in which they performed well.

RESULTS

With regard to triple-negative breast cancer, veliparib–carboplatin had an 88% predicted probability of success in a phase 3 trial. A total of 72 patients were randomly assigned to receive veliparib–carboplatin, and 44 patients were concurrently assigned to receive control therapy; at the completion of chemotherapy, the estimated rates of pathological complete response in the triple-negative population were 53% (95% Bayesian probability interval [PI], 36 to 66%) in the veliparib–carboplatin group versus 26% (95% PI, 9 to 43%) in the control group. The toxicity of veliparib–carboplatin was greater than that of the control.

CONCLUSIONS

The process used in our trial showed that veliparib–carboplatin added to standard therapy resulted in higher rates of pathological complete response than standard therapy alone specifically in triple-negative breast cancer. (Funded by the QuantumLeap Healthcare Collaborative and others; I-SPY 2 TRIAL ClinicalTrials.gov number, NCT01042379.)

The authors' full names, academic degrees, and affiliations are listed in the Appendix. Address reprint requests to Dr. Esserman at the UCSF Carol Franc Buck Breast Care Center, University of California, San Francisco, 2880 Divisadero St., Box 1710, San Francisco, CA 94115, or at laura.esserman@ucsfmbi.

*A complete list of participating centers and investigators in the Investigation of Serial Studies to Predict Your Therapeutic Response through Imaging and Molecular Analysis 2 (I-SPY 2 TRIAL) is provided in the Supplementary Appendix, available at NEJM.org.

N Engl J Med 2016;375:23-34.
DOI: 10.1056/NEJMoa1510749
Copyright © 2016 Massachusetts Medical Society.

I-SPY 2 is an **Innovative** public-private collaboration that combines **Personalized Medicine** & **Novel Trial Design** to develop new cancer treatments much faster and for much less cost

Response-adaptive randomization

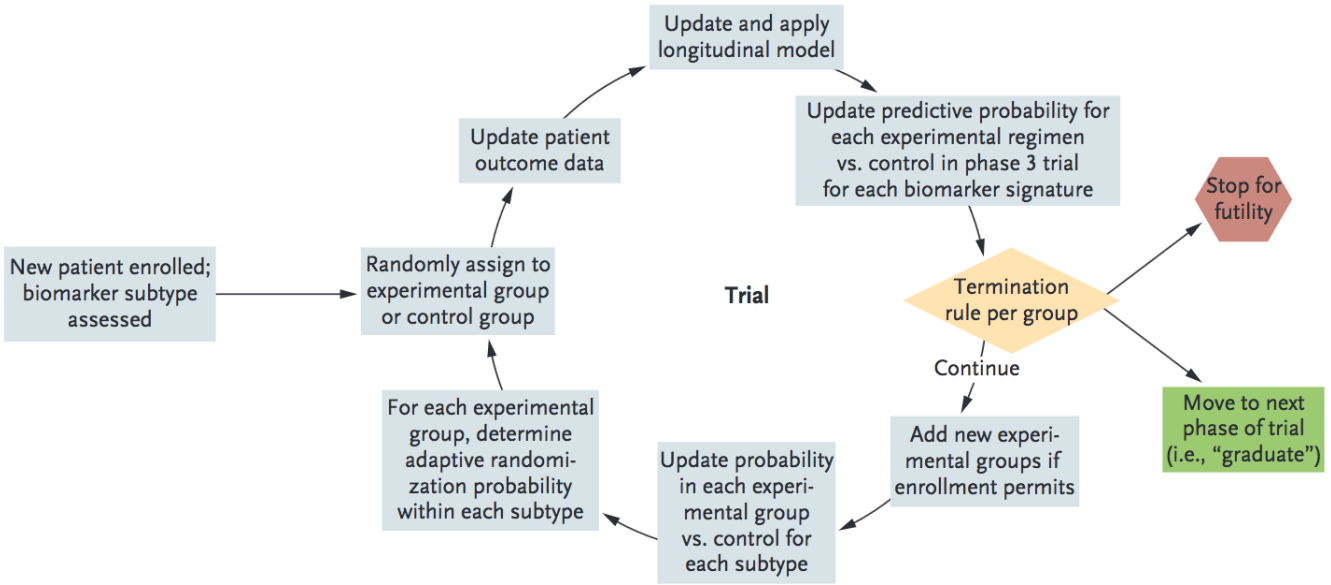

Rugo et al. NEJM 2016

The traditional RCT ...

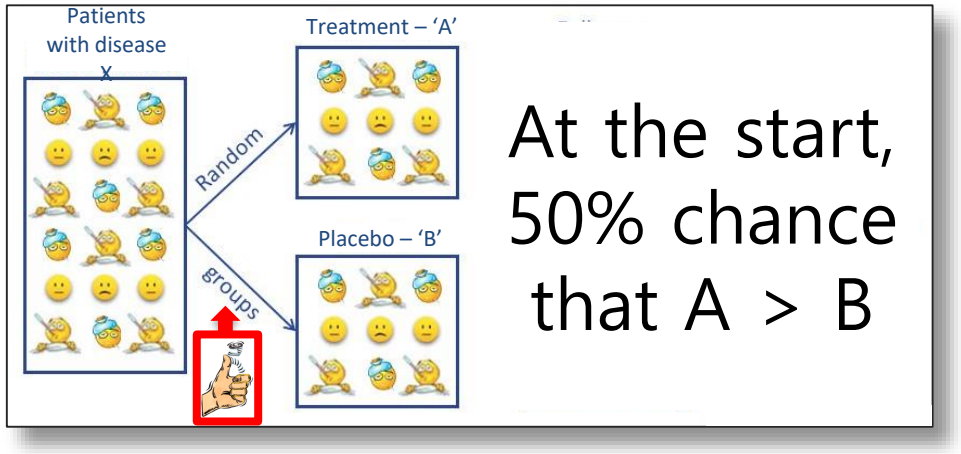

The traditional RCT …

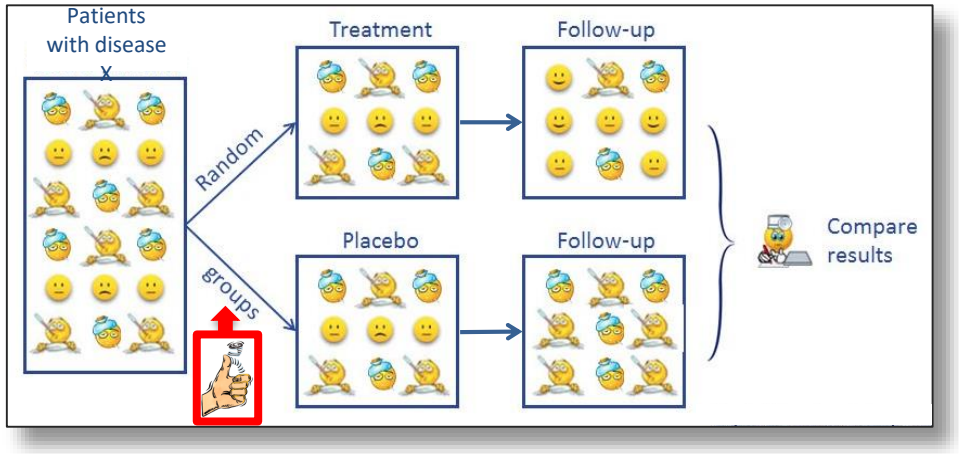

At the end, >99% sure that A > B

What about in the middle?

A planned trial of A vs. B in 400 patients

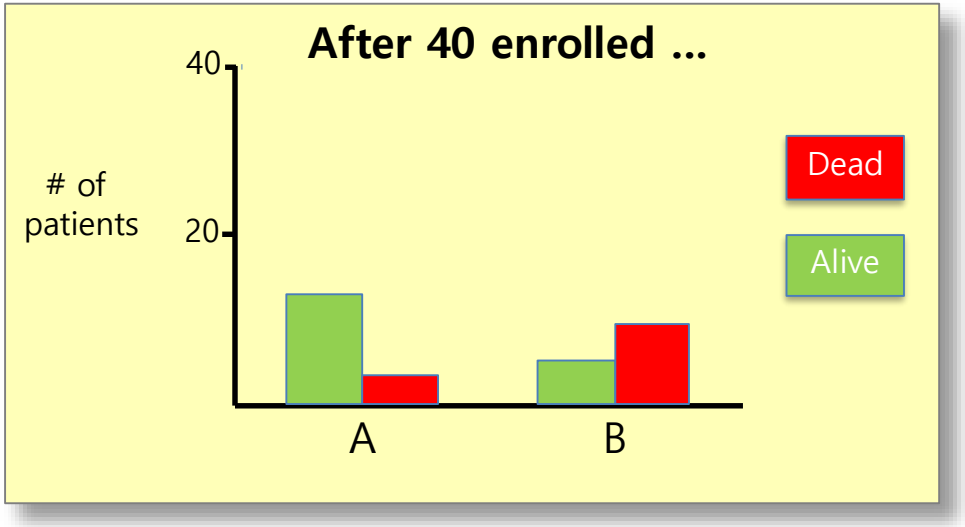

The probability that A > B = 78%

Start randomizing MORE patients to A than B …

After 80 patients …

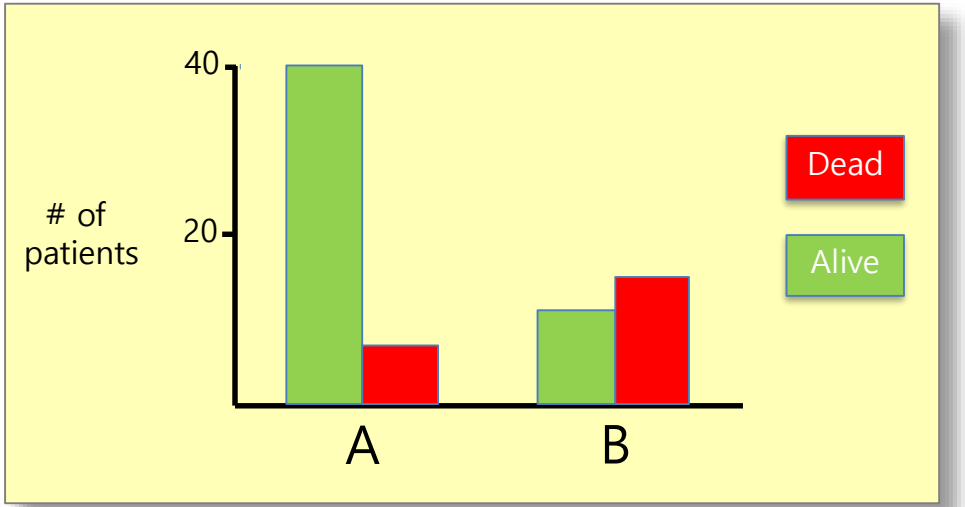

Now, the probability that A > B = 99.9%

Stop the trial!

Caveats

- If the ‘second’ 40 was flat or opposite direction …
  - Trial continues and the next ‘bet’ swings back closer to 50:50
- When 2 groups, power driven by the smaller group
- So, NOT very helpful if …
  - Single homogenous cohort
  - Two arms
- But, becomes VERY interesting when …
  - Multiple arms
  - Multiple subgroups

Response-adaptive randomization

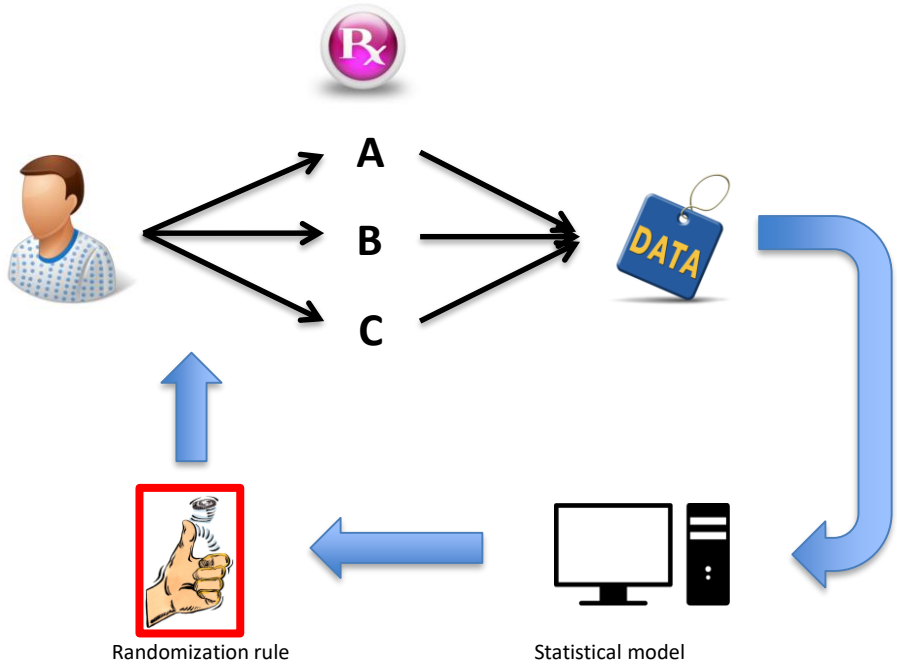

Response-adaptive randomization

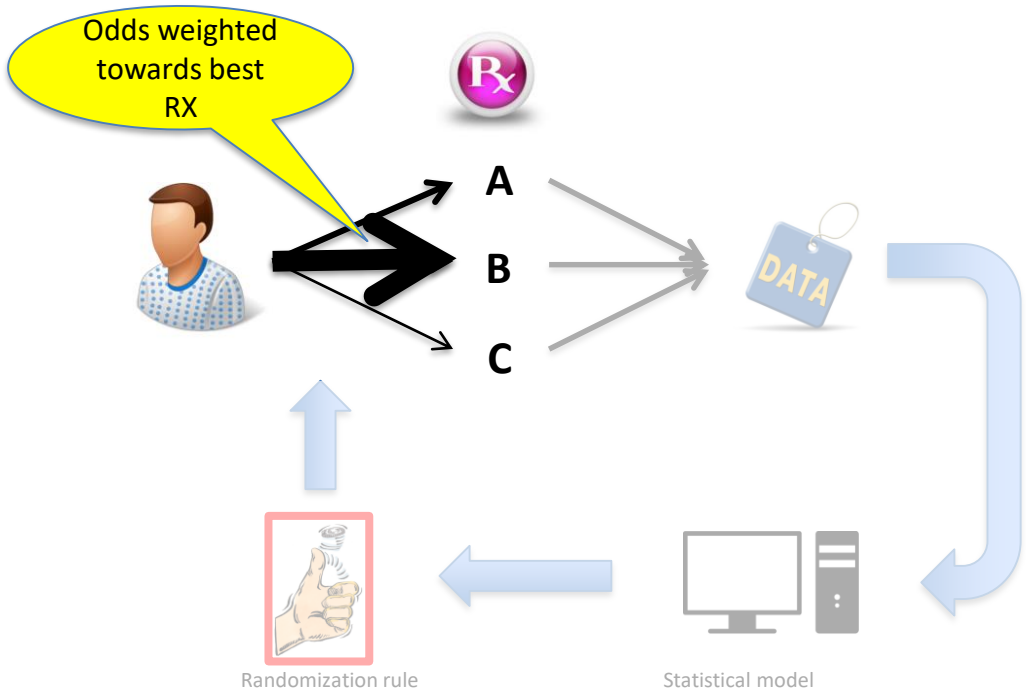

Response-adaptive randomization

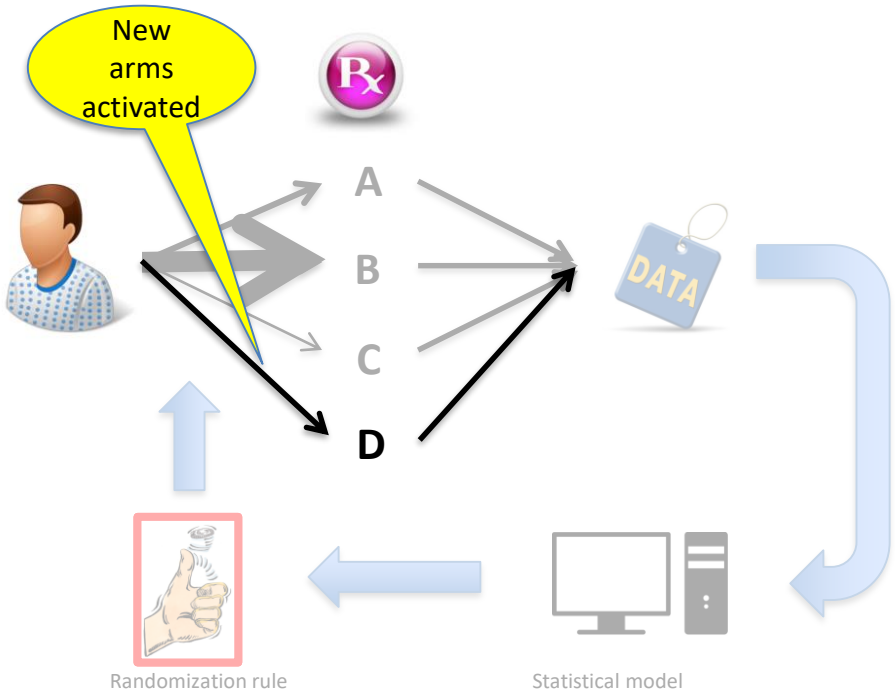

Response-adaptive randomization

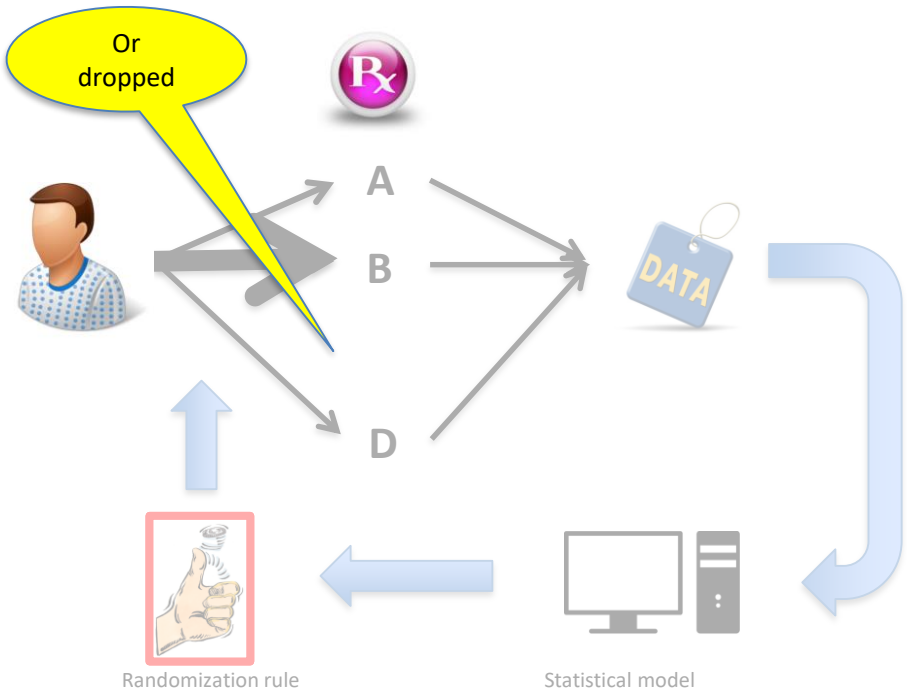

Response-adaptive randomization

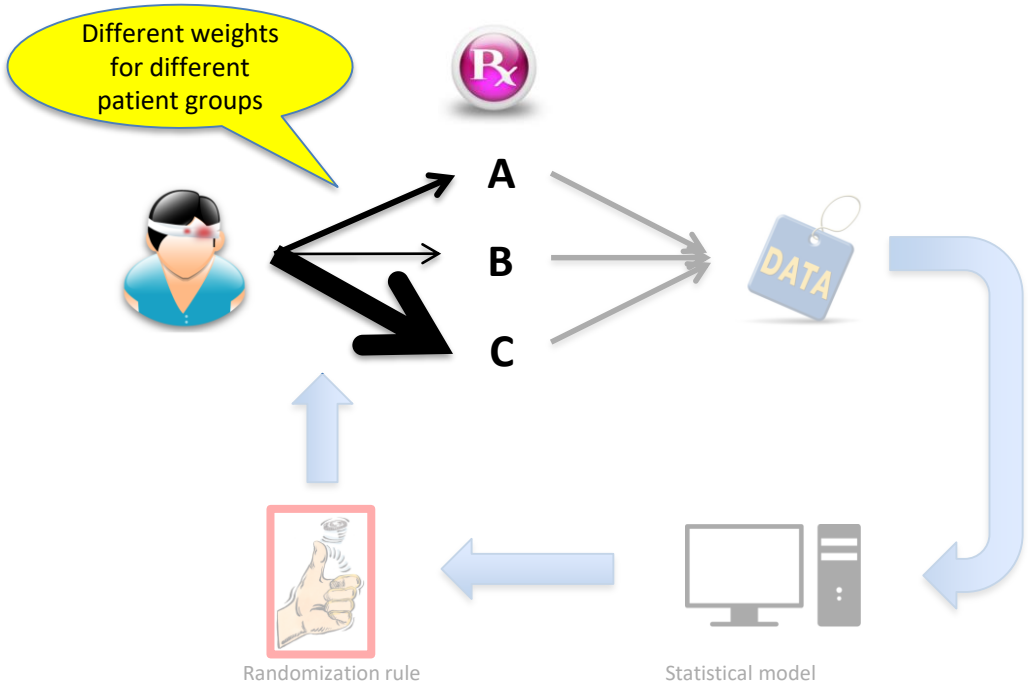


A
B
C
DATA
'Weighted' dice
Statistical model

RANDOMIZED	**Allow CAUSAL inference**
EMBEDDED	**Align with care; leverage the EHR**
MULTIFACTORIAL	**Multiple treatments and subgroups**
ADAPTIVE	**Match odds of success to odds of assignment**
PLATFORM	**Perpetual enrollment; continuous learning**

= REMAP

Angus DC. JAMA 2015

Scenario: 2 of 8 regimens are best

‘True’ mortality

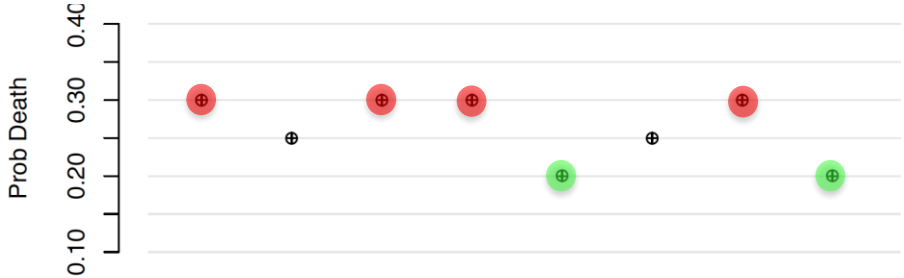

Scenario: 2 of 8 regimens are best

'True' mortality

For a 2,000 patient trial …

Average results from 1,000s of simulations

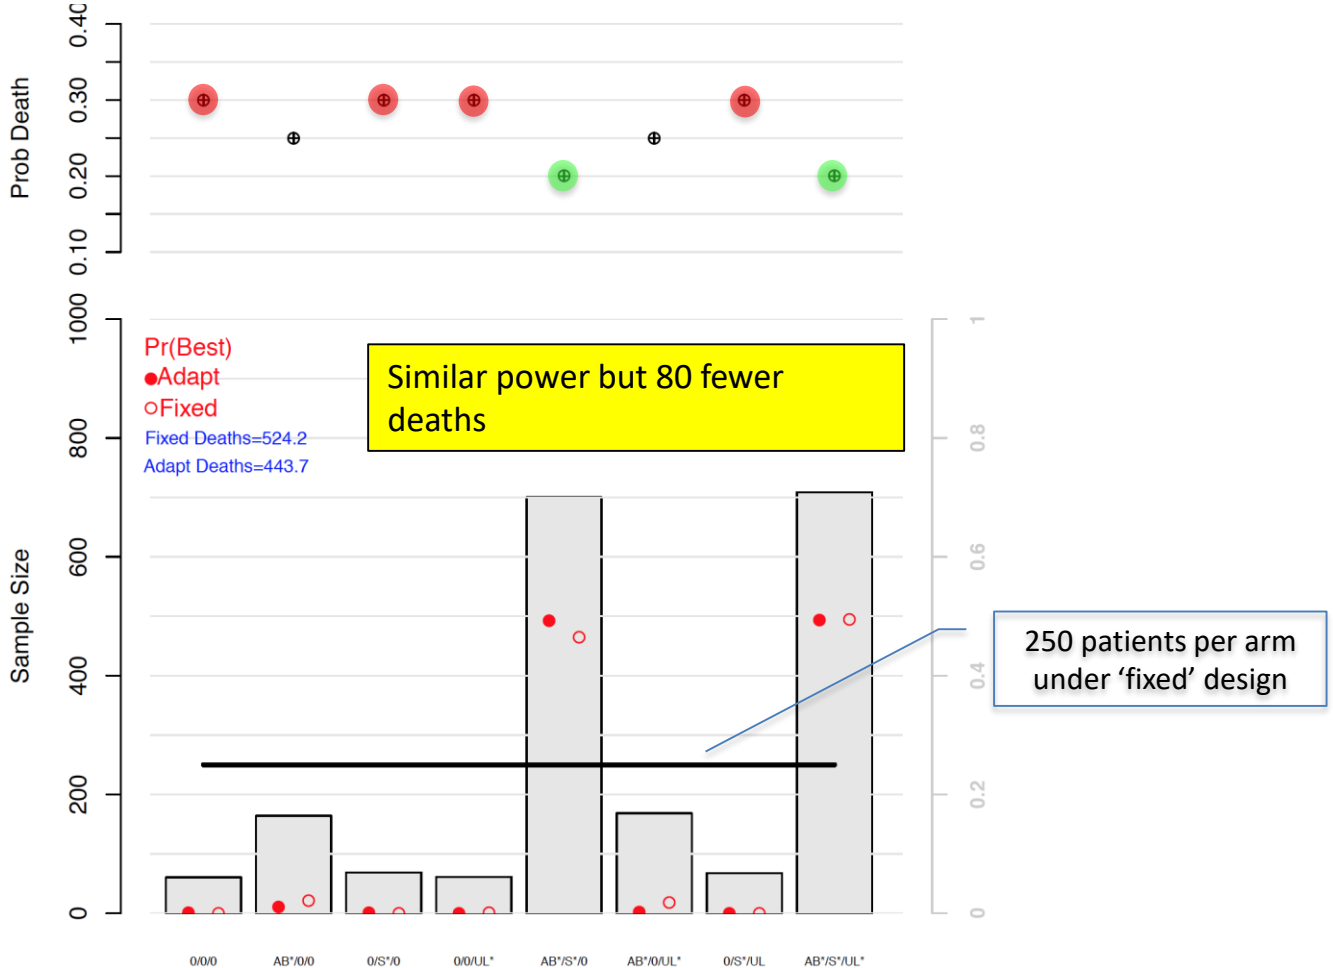

REMAP designs …

- Smart
  - Consider many different treatment options
  - Vary the options depending on the patient

- Safe
  - Probably ‘play’ what is probably the ‘winner’
  - On average, safer ‘in’ the trial than out of it …

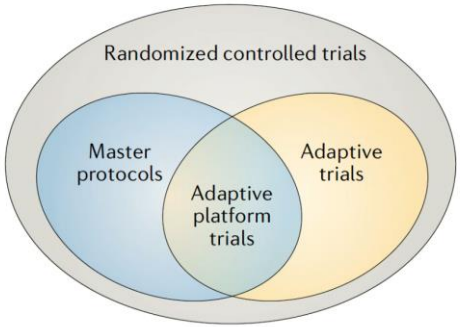

REMAP-COVID,
a ’sub-platform’ of REMAP-CAP

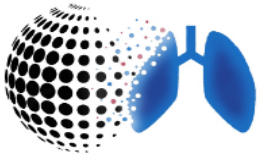

- Expanded to all hospitalized patients with COVID-19, in 2 strata
  - Moderate (hospitalized but not severe)
  - Severe (requiring ICU care for respiratory failure or shock)

REMAP-COVID,
a 'sub-platform' of REMAP-CAP

- Expanded to all hospitalized patients with COVID-19, in 2 strata
  - Moderate (hospitalized but not severe)
  - Severe (requiring ICU care for respiratory failure or shock)

- 1° endpoint: organ failure-free days
  - Death worst outcome, followed by number of days free of ICU-based cardiovascular or respiratory support through 21 days
  - Modeled with cumulative logistic proportional odds model

$$\log\left(\frac{\pi_y}{1-\pi_y}\right) = [Site] + [Time] + [Age] + \sum_{i=1}^{k}[Intervention] + \sum[IxI\ Interactions]$$

- 2° endpoints: mortality, WHO ordinal scale, safety

REMAP elements

- Domain – an area where a question is asked …
  - Domain #1 – choice of antibiotic
  - Domain #2 – whether to give steroids or not
  - Domain #4 – choice of ventilator strategy
  - Etc. ….
- Intervention
  - Any option within a domain …
- Regimen
  - Unique combination of interventions within a domain …
- Stratum
  - Baseline subgroup
  - Ex. Moderate vs. Severe COVID19 at presentation

Multifactorial intervention assignments

Regimen = set of domain-specific interventions
Effect of an intervention is conditional upon

- Stratum
- Interventions within other domains

Regimen	Domain A	Domain B	Domain C
#1	A1	B1	C1
#2	A1	B1	C2
#3	A1	B2	C1
#4	A1	B2	C2
#5	A2	B1	C1
…..			
#*n*	A*n*	B*n*	C*n*

REMAP-COVID domains/interventions

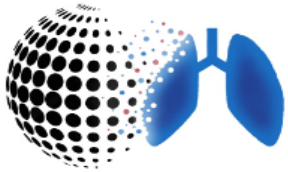

- Current COVID19-specific domains
  - Antiviral agents (NONE, HCQ, kaletra, HCQ/kaletra combo)
  - Corticosteroids (NONE, 3 doses)
  - Targeted innate immune modulation (NONE, IL1ra, 2 X IL6ra, IFNbeta, others)
  - Immunoglobulin therapy (NONE, CP, with synthetic IGs to be added later)

- Additional funded domains about to launch
  - Coagulation modulation (prophylaxis only, heparin, possibly dipyridamole)
  - High dose vitamin C (NONE, vitamin C)
  - Statin (NONE, simvastatin)

- Once these 7 domains all running, there are 1,280 separate regimens (recipes) …
  - Plus, more under development
    - ACE2 modulation (3 subdomains for binding and downstream activation)
    - Ventilation

What does background care look like?

- Surviving Sepsis Campaign Guidelines for COVID19
  - 54 separate care statements
    - Uncertainty regarding every statement

- Even if there are only 2 choices for each of these 54 statements …
  - 2^{54} care ‘regimens’

- In other words, all RCTs are taking place on potentially mammoth scale of background variation in care

REMAP-COVID design

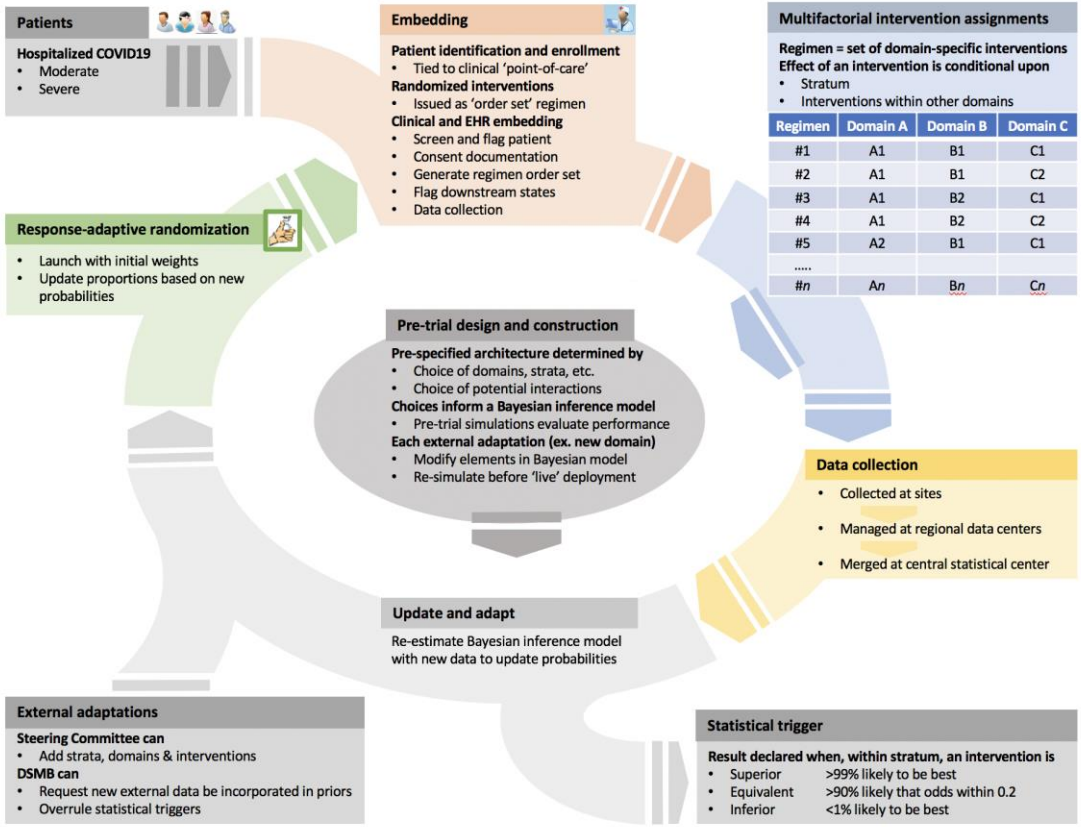

REMAP-COVID design

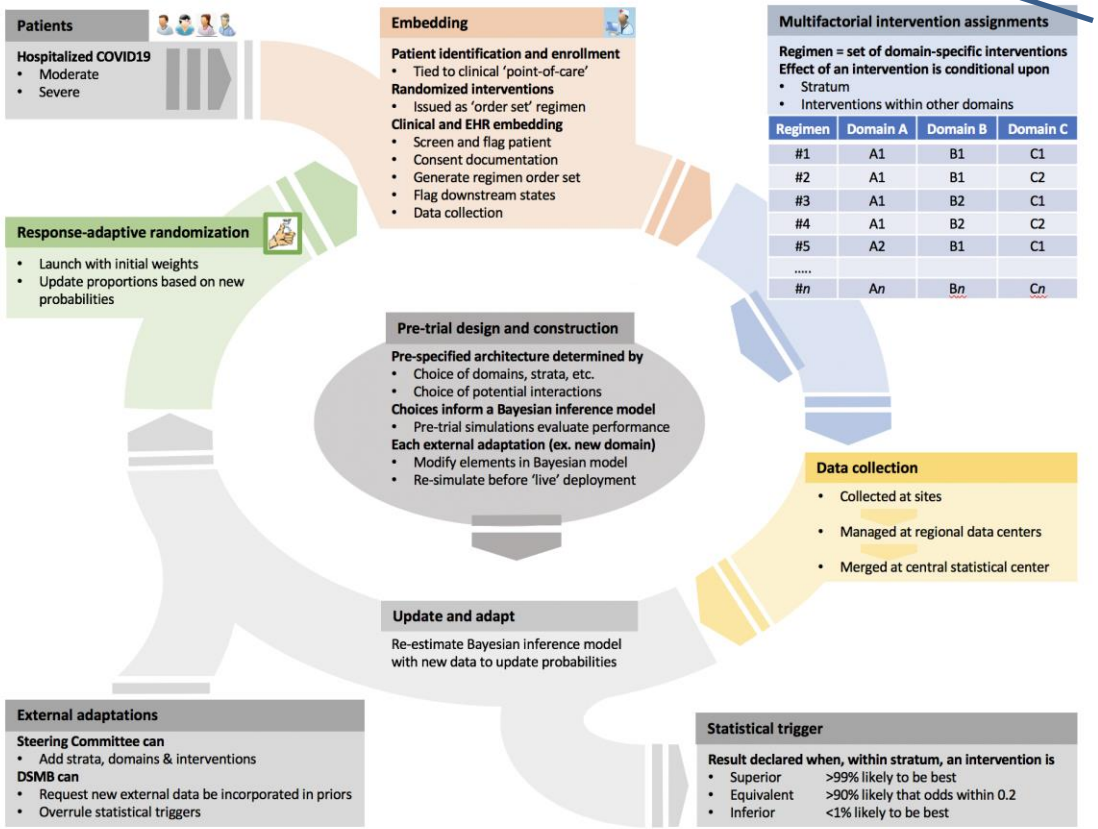

Regimens, domains and interventions

- Many domains can be added
- ~4 interventions can be tested within any 1 domain @ 1 time
- Interventions can be tested as a 'nest'
  - Ex. all IL-6 blocking agents vs. none
- A priori consideration re: interactions
- Each domain has a control arm
- If usual care inferior, can be dropped
  - Ex. if all IL6 blockers superior to 'none'

REMAP-CAP/COVID is global

- A federation of several highly successful clinical trial networks and coordinating centers
  - >100 sites and 13 countries ‘live’
  - New COVID-specific grants from EU, the Netherlands
    France, Germany, UK, Ireland, Canada, Australia and NZ

- Scaling up rapidly across the world
  - Funded to expand to >200 sites this month
  - Adding sites in Middle East and South America
  - Discussions for further expansion in Asia (e.g., Japan) and Africa

- Advantage – global positioning allows capture of patients across the globe

Simulations and power

- For ‘head-to-head’ within stratum with no interactions
  - ~400 per group for moderate (OR: 1.7) treatment effect

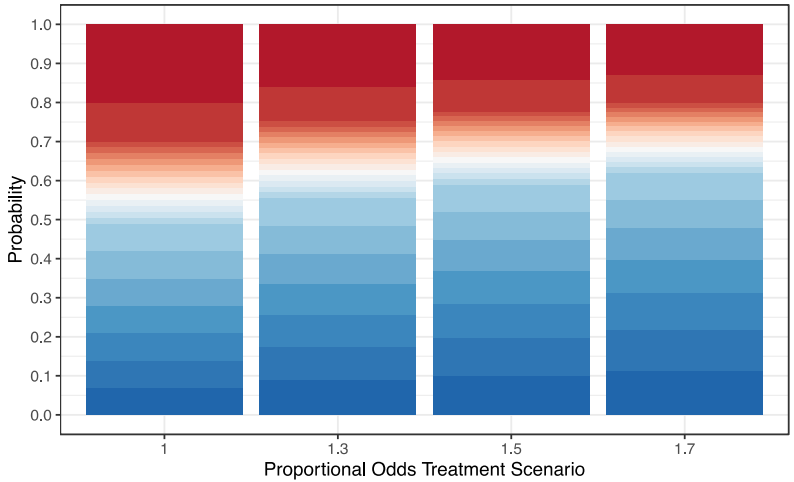

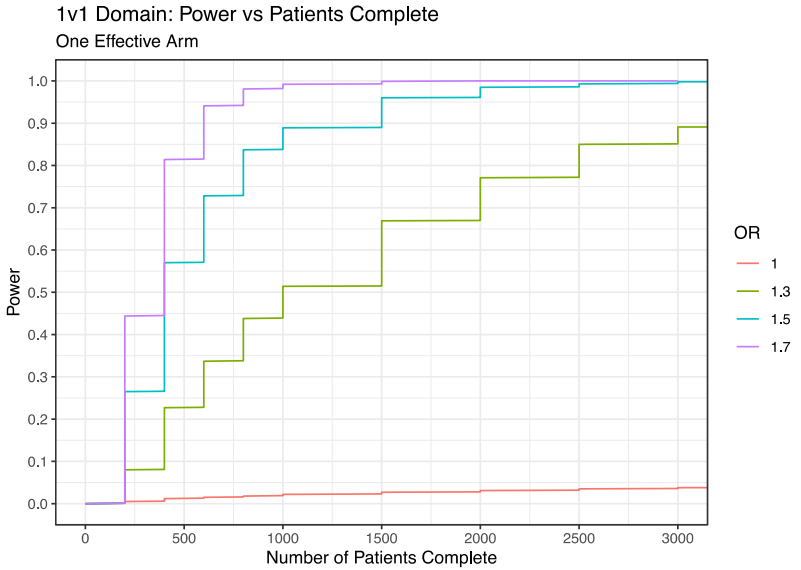

Ok, but …

- EHR data quality
- Institutional commitment
- Ethics
- Statistics and design
- Reporting and dissemination of results
- Funding
- Oversight
- Integration with other clinical research programs

A comment on eligibility …

- Sites can decline to participate in any particular domain or intervention
- Eligibility can also ‘blink’ (temporary inavailability)
- Patients can be ineligible for any particular intervention or domain
- Both patient and site eligibility, by time, is tracked in the model
  - ‘Controls’ are only those who ‘could’ have received an intervention …

A comment on RAR and contemporaneousness of controls …

- Principally, patients who receive a given intervention are compared to patients who contemporaneously serve as controls

- But, relative proportions change over time …
  - Time (by month) included in the model

A comment on suitability for registration …

- Conceptually, the trial platform is simply ‘hosting’ multiple parallel questions
  - Comparative effectiveness questions
  - Registration trial questions

- Any single domain can run as a free-standing question …

- Thus, if necessary for a registration trial …
  - Alpha error control can be specified
  - Placebo (or combination of placebo) can be specified
  - Bounds on RAR can be specified
  - Limits on select ‘co-randomization’ can be specified with other domains

Conclusions

- This pandemic forces us to do 2 things simultaneously
  - Do
  - Learn
- These activities are intertwined: we must ‘Learn While Doing’
- Unfortunately, ‘practice’ traditionally separated from ‘research’
  - The two enterprises must lean in to each other
  - Use ‘learning designs’ that accommodate ‘doing’ at the same time
- Global adaptive platform trials have potential as LWD instruments