Optimal Two-Stage Adaptive Enrichment Designs, using Sparse Linear Programming

Michael Rosenblum
Johns Hopkins Bloomberg School of Public Health

Joint work with Xingyuan (Ethan) Fang and Han Liu
Princeton University

Working paper giving full results available here:
https://goo.gl/xXcw5C

October 23, 2015

Problem Motivation

Goal: Testing Treatment Effects in Two Subpopulations and the Overall Population in Adaptive Enrichment Designs

Example:

- Treating resistant HIV. Recent HIV drugs (maraviroc, raltegravir) have shown stronger benefit in those with lower phenotypic sensitivity to background therapy.

We assume two, predefined, subpopulations that partition the overall population.

Multiple Testing Problem: Null Hypotheses Definition

Define three treatment effects of interest:

- Δ_1: Mean Treatment Effect for Subpopulation 1 (i.e., difference between population mean of the primary outcome under treatment and under control)
- Δ_2: Mean Treatment Effect for Subpopulation 2
- $\Delta_C = p_1\Delta_1 + (1 - p_1)\Delta_2$: Mean Treatment Effect for Combined Population

Multiple Testing Problem: Null Hypotheses Definition

Define three treatment effects of interest:

- Δ_1: Mean Treatment Effect for Subpopulation 1 (i.e., difference between population mean of the primary outcome under treatment and under control)
- Δ_2: Mean Treatment Effect for Subpopulation 2
- $\Delta_C = p_1\Delta_1 + (1 - p_1)\Delta_2$: Mean Treatment Effect for Combined Population

Goal: construct adaptive enrichment design D and multiple testing procedure M for:

- $H_{01} : \Delta_1 \leq 0$,
- $H_{02} : \Delta_2 \leq 0$,
- $H_{0C} : p_1\Delta_1 + (1 - p_1)\Delta_2 \leq 0$,

that strongly controls familywise Type I error rate, and is optimal in sense defined below.

General Two-Stage Adaptive Enrichment Design

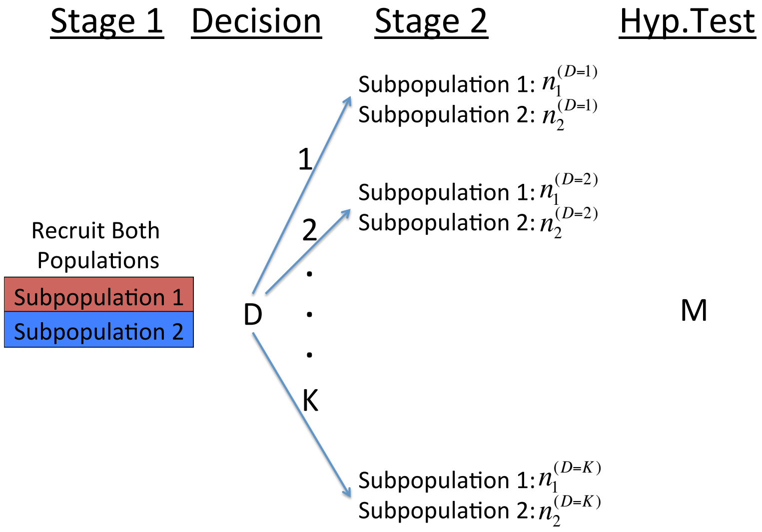

Example Two-Stage Adaptive Enrichment Design

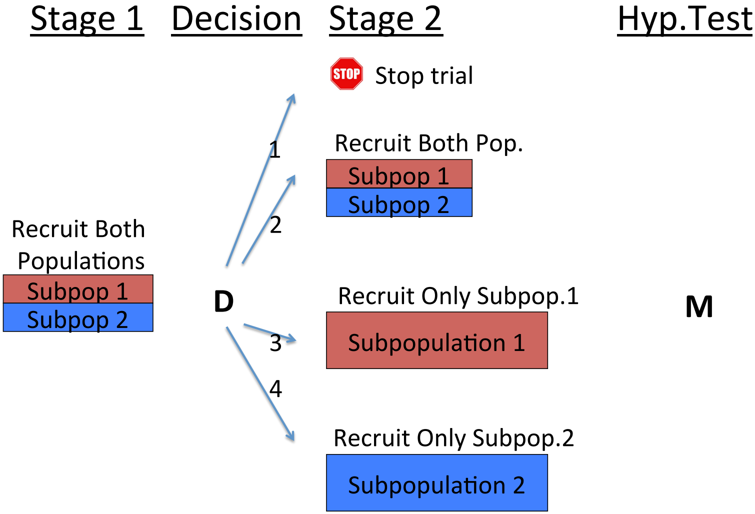

Example Two-Stage Adaptive Enrichment Design

$n =$ total sample size if both subpopulations enrolled in stage 2.

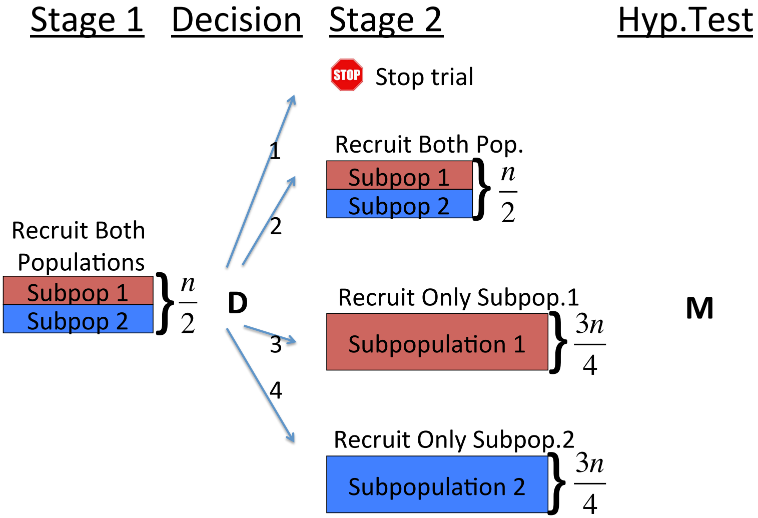

Two-Stage Adaptive Enrichment Design

Assume known variances and normally distributed outcomes; subpopulation cumulative sample sizes and z-statistics and are then sufficient statistics for $\Delta_1, \Delta_2, \Delta_C$.

Two-Stage Adaptive Enrichment Design

Assume known variances and normally distributed outcomes; subpopulation cumulative sample sizes and z-statistics and are then sufficient statistics for $\Delta_1, \Delta_2, \Delta_C$.

- $Z_s^{(1)}$: z-statistic for subpopulation s at end of stage 1;
- $Z_s^{(F)}$: Final, cumulative z-statistic for subpop. s at end of stage 2.

Two-Stage Adaptive Enrichment Design

Assume known variances and normally distributed outcomes; subpopulation cumulative sample sizes and z-statistics and are then sufficient statistics for $\Delta_1, \Delta_2, \Delta_C$.

- $Z_s^{(1)}$: z-statistic for subpopulation s at end of stage 1;
- $Z_s^{(F)}$: Final, cumulative z-statistic for subpop. s at end of stage 2.

Decision rule D is map from $\mathbf{Z}^{(1)} = (Z_1^{(1)}, Z_2^{(1)})$ to possible decisions $\mathcal{D}$.
Multiple testing procedure M is map from $\mathbf{Z}^{(F)} = (Z_1^{(F)}, Z_2^{(F)})$ and decision D to set of null hypotheses rejected (if any).

Two-Stage Adaptive Enrichment Design

Assume known variances and normally distributed outcomes; subpopulation cumulative sample sizes and z-statistics and are then sufficient statistics for $\Delta_1, \Delta_2, \Delta_C$.

- $Z_s^{(1)}$: z-statistic for subpopulation s at end of stage 1;
- $Z_s^{(F)}$: Final, cumulative z-statistic for subpop. s at end of stage 2.

Decision rule D is map from $\mathbf{Z}^{(1)} = (Z_1^{(1)}, Z_2^{(1)})$ to possible decisions $\mathcal{D}$.
Multiple testing procedure M is map from $\mathbf{Z}^{(F)} = (Z_1^{(F)}, Z_2^{(F)})$ and decision D to set of null hypotheses rejected (if any).
Power at alternative Δ_1, Δ_2 to reject H_{01} is

$$\Pr_{(\Delta_1,\Delta_2)}[M(\mathbf{Z}^{(F)}, D(\mathbf{Z}^{(1)})) \text{ rejects } H_{01}].$$

User specifies: (i) loss function $L(D, M; \Delta_1, \Delta_2)$, e.g., total sample size; and (ii) distribution Λ on alternatives (Δ_1, Δ_2).

Constrained Bayes Optimization Problem

Problem inputs: p_1; set of possible stage 2 decisions; σ_1^2, σ_2^2; clinically meaningful min. treatment effect Δ^{min}; loss function L; distribution Λ on alternatives (Δ_1, Δ_2); $\alpha, \beta_1, \beta_2, \beta_C$.

Recall $D = D(\mathbf{Z}^{(1)})$ and $M = M(\mathbf{Z}^{(F)}, D(\mathbf{Z}^{(1)}))$.

Constrained Bayes Optimization Problem

Problem inputs: p_1; set of possible stage 2 decisions; σ_1^2, σ_2^2; clinically meaningful min. treatment effect $\Delta^{\min}$; loss function L; distribution Λ on alternatives (Δ_1, Δ_2); $\alpha, \beta_1, \beta_2, \beta_C$.

Recall $D = D(\mathbf{Z}^{(1)})$ and $M = M(\mathbf{Z}^{(F)}, D(\mathbf{Z}^{(1)}))$.

Constrained Bayes Opt. Problem: Find pair (D, M) minimizing:

$$\int E_{(\Delta_1,\Delta_2)}[L(D, M; \Delta_1, \Delta_2)]d\Lambda(\Delta_1, \Delta_2),$$

under **familywise Type I error constraints**:

$$\sup_{(\Delta_1,\Delta_2)\in\mathbb{R}^2} \Pr_{(\Delta_1,\Delta_2)}[M \text{ rejects any true null hypothesis}] \leq \alpha,$$

Constrained Bayes Optimization Problem

Problem inputs: p_1; set of possible stage 2 decisions; σ_1^2, σ_2^2; clinically meaningful min. treatment effect $\Delta^{\min}$; loss function L; distribution Λ on alternatives (Δ_1, Δ_2); $\alpha, \beta_1, \beta_2, \beta_C$.

Recall $D = D(\mathbf{Z}^{(1)})$ and $M = M(\mathbf{Z}^{(F)}, D(\mathbf{Z}^{(1)}))$.

Constrained Bayes Opt. Problem: Find pair (D, M) minimizing:

$$\int E_{(\Delta_1,\Delta_2)}[L(D, M; \Delta_1, \Delta_2)] d\Lambda(\Delta_1, \Delta_2),$$

under **familywise Type I error constraints**:

$$\sup_{(\Delta_1,\Delta_2)\in\mathbb{R}^2} \Pr_{(\Delta_1,\Delta_2)}[M \text{ rejects any true null hypothesis}] \leq \alpha,$$

and **power constraints**:

$$\begin{aligned}
\Pr_{(\Delta^{\min},0)}[M \text{ rejects } H_{01}] &\geq 1 - \beta_1. \\
\Pr_{(0,\Delta^{\min})}[M \text{ rejects } H_{02}] &\geq 1 - \beta_2. \\
\Pr_{(\Delta^{\min},\Delta^{\min})}[M \text{ rejects } H_{0C}] &\geq 1 - \beta_C.
\end{aligned}$$

Our Method to Solve Optimization Problem

Computational challenge: continuum of Type I error constraints.
(It's not enough to satisfy only at global null $(\Delta_1, \Delta_2) = (0, 0)$.)

Our Method to Solve Optimization Problem

Computational challenge: continuum of Type I error constraints. (It's not enough to satisfy only at global null $(\Delta_1, \Delta_2) = (0, 0)$.)

Our solution:

1. Discretize decision region $\mathbb{R}^2$ into small rectangles $\mathcal{R}$; for any $r \in \mathcal{R}$, enforce decision rule D makes same decision for any $(Z_1^{(1)}, Z_2^{(1)}) \in r$.
2. For each decision $d \in \{1, \ldots, K\}$, discretize rejection regions $\mathbb{R}^2$ into small rectangles $\mathcal{R}'_d$; for any $r' \in \mathcal{R}'_d$, enforce that if $D = d$, multiple testing procedure M rejects same set of hypotheses for any $(Z_1^{(F)}, Z_2^{(F)}) \in r'$.
3. Discretize Type I error constraints into fine grid on boundaries of null spaces.

Our Method to Solve Optimization Problem

Computational challenge: continuum of Type I error constraints. (It's not enough to satisfy only at global null $(\Delta_1, \Delta_2) = (0, 0)$.)

Our solution:

1. Discretize decision region $\mathbb{R}^2$ into small rectangles $\mathcal{R}$; for any $r \in \mathcal{R}$, enforce decision rule D makes same decision for any $(Z_1^{(1)}, Z_2^{(1)}) \in r$.
2. For each decision $d \in \{1, \ldots, K\}$, discretize rejection regions $\mathbb{R}^2$ into small rectangles $\mathcal{R}'_d$; for any $r' \in \mathcal{R}'_d$, enforce that if $D = d$, multiple testing procedure M rejects same set of hypotheses for any $(Z_1^{(F)}, Z_2^{(F)}) \in r'$.
3. Discretize Type I error constraints into fine grid on boundaries of null spaces.

Discretized opt. problem is not convex. However, we construct reparametrization that is sparse, linear program:

$$\max_{x} \quad \mathbf{c}^{\mathbf{T}}\mathbf{x} \qquad \text{s.t.} \quad \mathbf{A}\mathbf{x} \leq \mathbf{b}.$$

We apply advanced optimization methods to solve this.

Example

$p_1 = 1/2$, $\alpha = 0.05$, $\sigma_1^2 = \sigma_2^2$. $L =$total sample size. Prior Λ equally weighted pt. masses at (Δ_1, Δ_2) equal to $(0,0), (\Delta^{\min}, 0)$, $(0, \Delta^{\min})$, $(\Delta^{\min}, \Delta^{\min})$.

Example

$p_1 = 1/2$, $\alpha = 0.05$, $\sigma_1^2 = \sigma_2^2$. L =total sample size. Prior Λ equally weighted pt. masses at (Δ_1, Δ_2) equal to $(0,0), (\Delta^{\min}, 0)$, $(0, \Delta^{\min})$, $(\Delta^{\min}, \Delta^{\min})$. Sample size n is min s.t. standard design satisfies each power constraint at $1 - \beta_j = 0.64$ for $j \in \{1, 2, C\}$. Set each $1 - \beta_j = 0.82$.

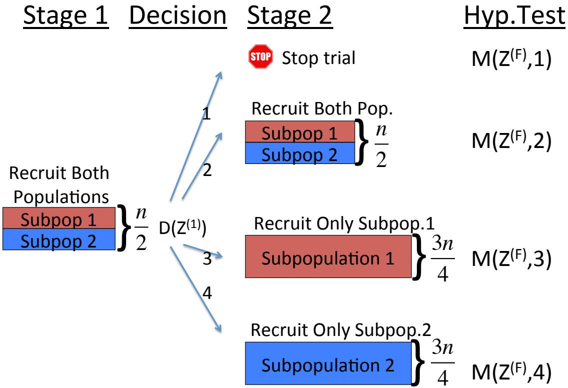

Example

$p_1 = 1/2$, $\alpha = 0.05$, $\sigma_1^2 = \sigma_2^2$. L =total sample size. Prior Λ equally weighted pt. masses at (Δ_1, Δ_2) equal to $(0,0), (\Delta^{\min}, 0)$, $(0, \Delta^{\min})$, $(\Delta^{\min}, \Delta^{\min})$. Sample size n is min s.t. standard design satisfies each power constraint at $1 - \beta_j = 0.64$ for $j \in \{1, 2, C\}$. Set each $1 - \beta_j = 0.82$.

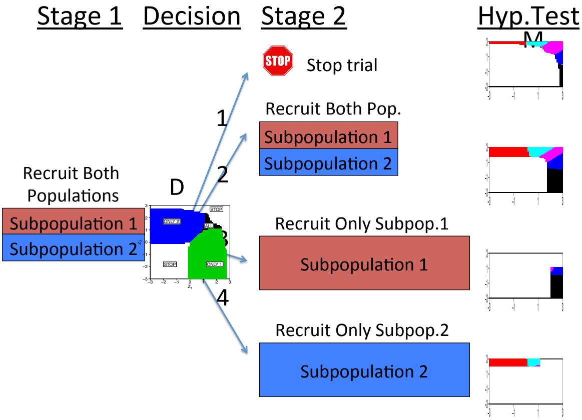

Approximately Optimal Solution

Decision Rule for Stage 2 Enrollment:

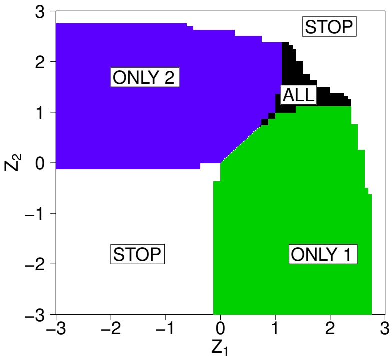

Rejection Regions under Each Decision: (in terms of $Z_1^{(F)}, Z_2^{(F)}$)

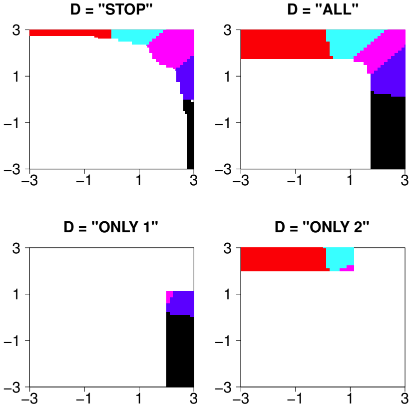

Approximately Optimal Solution

Decision Rule to Enroll Stage 2:

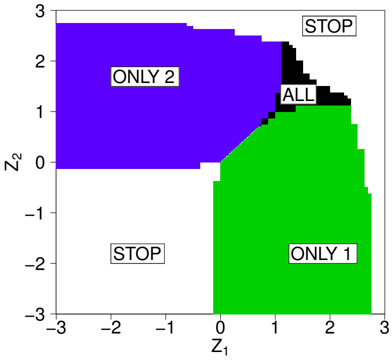

Rejection Regions for Decision:

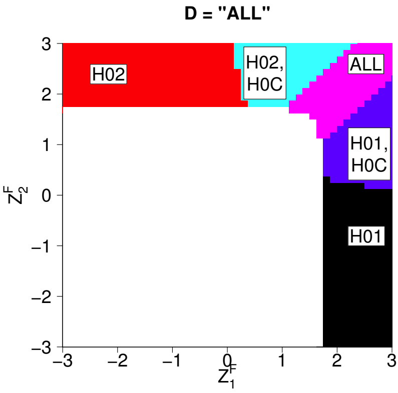

Approximately Optimal Solution

Decision Rule to Enroll Stage 2:

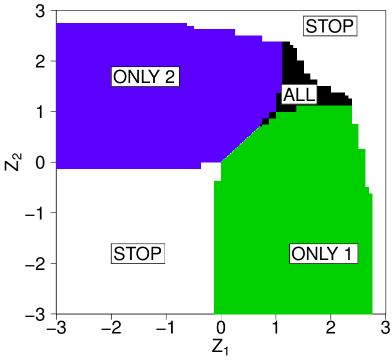

Rejection Regions for Decision:

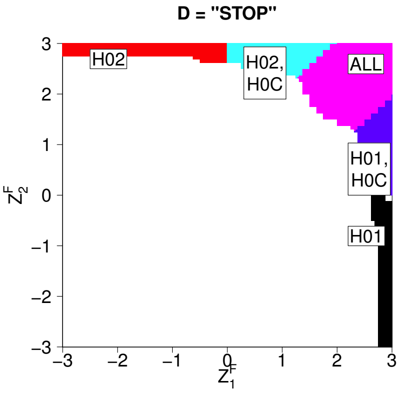

Approximately Optimal Solution

Decision Rule to Enroll Stage 2:

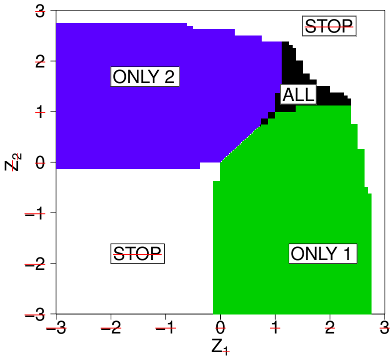

Rejection Regions for Decision:

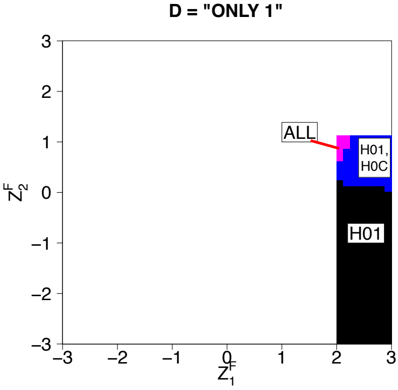

Approximately Optimal Solution

Decision Rule to Enroll Stage 2: Rejection Regions for Decision:

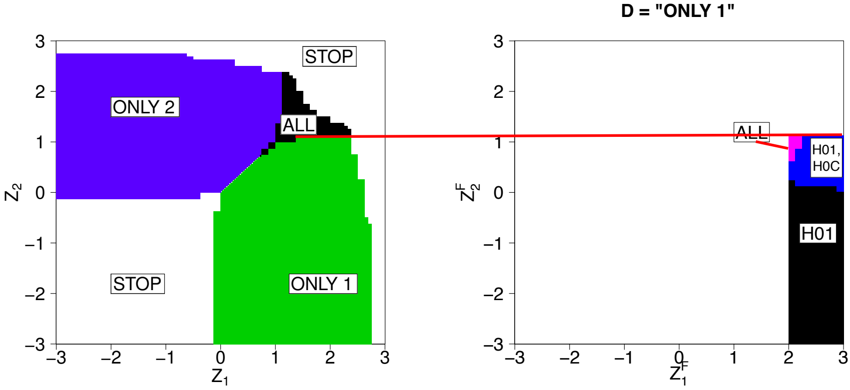

Approximately Optimal Solution

Decision Rule to Enroll Stage 2:

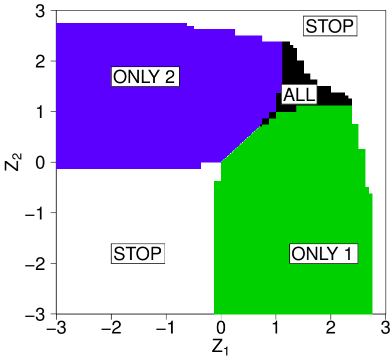

Rejection Regions for Decision:

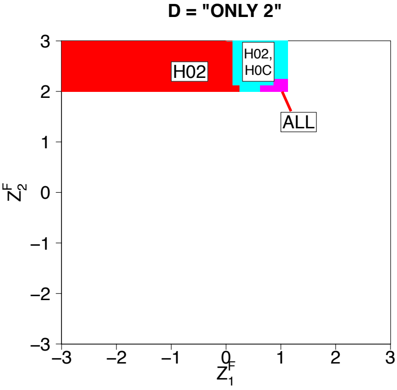

Power Comparison

Compare to adaptive enrichment design using p-value combination approach (Bauer and Köhne, 1994), with Dunnett intersection test and inverse-normal combination function. Early stopping is incorporated using O'Brien-Fleming boundaries for each intersection null hypothesis. Decision rule for stage 2:

- if combined population statistic $(Z_1^{(1)} + Z_2^{(1)})/\sqrt{2} > t_c$, enroll both subpop.
- else, enroll from each subpopulation s for which $Z_s^{(1)} > t$.

Consider $\beta = \beta_1 = \beta_2 = \beta_C$. For each power threshold $1 - \beta$, we optimized over t, t_c to minimize expected sample size under the power constraints. $n =$ total sample size if both enrolled stage 2.

Table: Minimum of $\int ESS \, d\Lambda$, as power constraint $1 - \beta$ varied.

Required Power $1 - \beta$:	70%	74%	78%	82%
Comparator	$0.97n$	$1.01n$	infeasible	infeasible
Optimal	$0.79n$	$0.84n$	$0.92n$	$1.03n$

References

Rosenblum, M., Fang, X., and Liu, H., Optimal, Two Stage, Adaptive Enrichment Designs for Randomized Trials Using Sparse Linear Programming (2014). Johns Hopkins University, Dept. of Biostatistics Working Papers. Working Paper 273. http://biostats.bepress.com/jhubiostat/paper273

Rosenblum, M., Liu, H., and Yen, E.-H. (2014), Optimal Tests of Treatment Effects for the Overall Population and Two Subpopulations in Randomized Trials, using Sparse Linear Programming, *Journal of American Statistical Association, Theory and Methods Section*, Volume 109. Issue 507. 1216-1228.

Rosenblum, M. (In Press), Adaptive Randomized Trial Designs that Cannot be Dominated by Any Standard Design at the Same Total Sample Size. *Biometrika*.

Collaborators

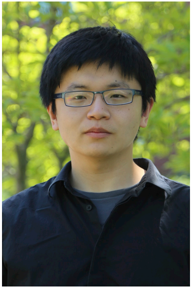

Ethan X. Fang

Han Liu

Princeton University

Thank you!

Acknowledgement:

- This research was supported by the U.S. Food and Drug Administration (HHSF223201400113C) and the Patient-Centered Outcomes Research Institute (ME-1306-03198).

- This work is solely the responsibility of the authors and does not represent the views of these agencies.