Adaptive Trial Designs

Wenjing Zheng, Ph.D.
Methods Core Seminar
Center for AIDS Prevention Studies
University of California, San Francisco
Nov. 17th, 2015

Trial Design

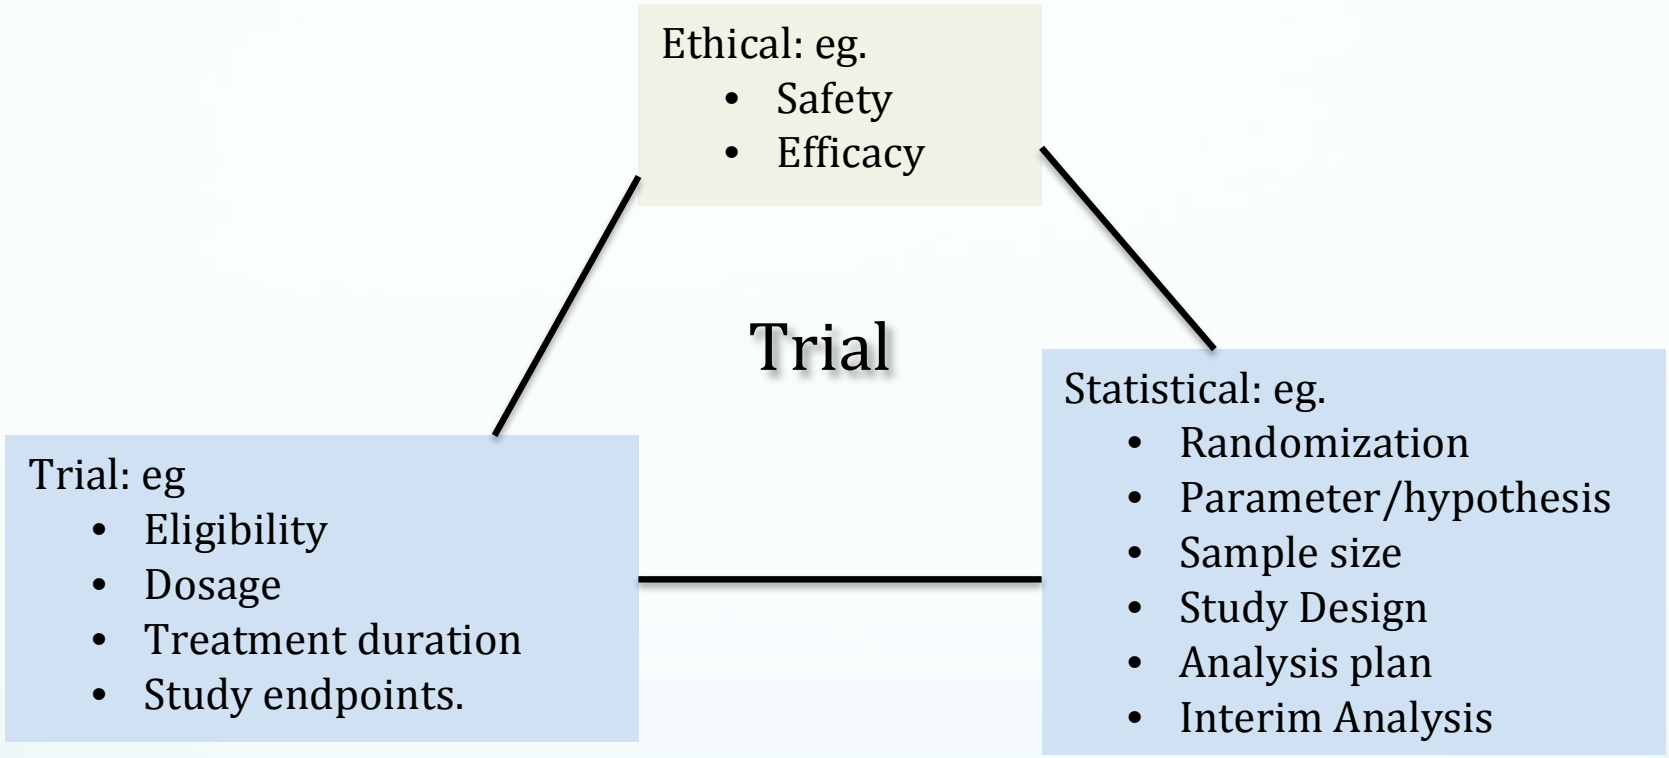

- Fixed Trial Design: All aspects/settings are set prior to start of the trial.
- Adaptive Trial Design: allows pre-specified modifications of trial aspects based on accruing data

Side bar: Adaptive interventions vs Adaptive designs

- Adaptive: uses accruing data to make decision.
- Adaptive intervention: dynamic treatment rule for a patient. Deterministically assign intervention option based on patient's accruing data.
- Adaptive design: experiment design mid-trial modification of some as trial settings based on accruing data of all subjects.

Adaptive Group Sequential Trial Designs

- Patients enrolled over time.
- At interim points, analysis of accruing data to inform:
  - sample size recalculation, stop trial early (e.g. efficacy, safety).
  - selection/enrollment from specific subpopulation (enrichment).
  - modify experimental features: treatment allocation, monitoring intensity, covariate measurement, clinical endpoints.

Adaptive Sample Size and Stopping Time

- Adaptive sample size recalculation (e.g. Tsiatis and Mehta 2003):
  - Accruing data to estimate treatment effect.
  - Modify sample size to achieve power for new alternative.
  - Or: re-estimate variances. Cluster-based trials, repeated measures (Lake et.al. 2002).

- Group-sequential testing (fixed designs):
  - Enough information to warrant early termination? (safety, efficacy).
  - Compare test statistic to a stopping rule
  - Key: Control type I error.
  - Unique a-priori alternative.

Adaptive Population Selection

- Adaptive Enrichment Designs (e.g. Rosenblum and van der Laan 2011, Simon and Simon 2013)

  - subsequently enroll only patients that would benefit most from treatment.

Stage 1	→	Interim	→	Stage 2
Recruit from entire population. Measure covariates, treatment and outcome.		Assess treatment effect on total popn and subpopn.		Recruit only from subpopn. with largest effect

- Other data-adaptive subpopn. selections:
  - most likely to adhere, most change in mediator.

- Modify population sampled: implicitly involves multiple testing

Group-Sequential Adaptive Randomization Designs

- Modify treatment allocation probabilities.
- Covariate-adaptive: modification based on balancing prognostic covariates.
- Response-Adaptive: modification based on previous patients' response. (eg. To maximize power and precision, minimize exposure to inferior rx).
- Covariate-Adjusted Response-Adaptive (CARA): allocation depends on own covariates, modification based on previous patients outcomes. ➔ address heterogeneity.
- CARA design: primary study objective + adaptation optimality criterion.

Today: CARA Randomization Design

- Our recent work on CARA Design (**Zheng** et.al. 2015, Chambaz et.al. 2012, van der Laan 2008):
  - General CARA design and analysis framework.
  - Allows general classes of optimality criterion.
  - Robust to model mis-specification.
  - Machine learning to target optimal allocation, while still have valid inference.

- The embedded conceptual framework can be applied to e.g.
  - optimal dosage finding.
  - adaptation of clinical endpoints.
  - adaptation of monitoring mechanism.
  - adaptation of covariate collection

Statistical Framework: Data Structure

- For the k-th patient, we observe pre-rx covariates W_k, treatment A_k, primary outcome Y_k.
- $O_k=(W_k, A_k, Y_k)$
- W_k and Y_k (as a function of treatment and covariate) are given to us by nature, same mechanism for everyone.
- A_k is allocated based on our k-th conditional probability $g_k(A_k=a|W_k=w)$.
- Adaptive randomization: different g_k for each patient. How do we build g_k?

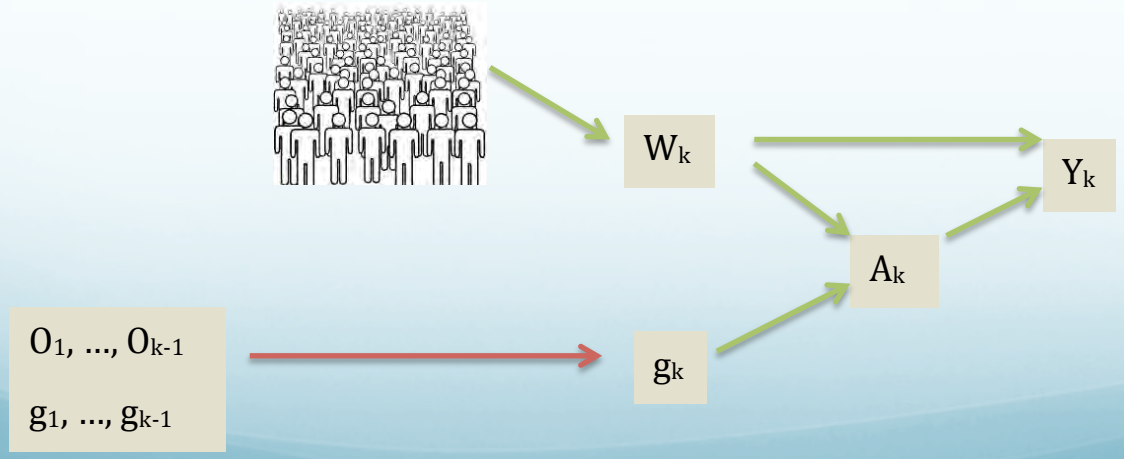

Statistical Framework: Parameter of Interest and Optimal Allocation

- Parameter of Interest: marginal treatment effect: $E_W[E(Y|A=1, W)-E(Y|A=0,W)]$.
- Optimal allocation (if we knew nature's outcome model Q_0):

  g* which minimizes the expected value of a loss function $f(g; Q_0)$, which depends on goal of adaptation
- Eg. Goal = maximize efficiency of trial:
  - then g* minimizes expected value of $f(g; Q_0)=(Y-Q_0(A,W))^2/g^2(A|W)$.
  - aka Neyman allocation.

➔

- Targeted CARA design: sequence of allocations that approximate g*.
- Targeted MLE analysis: robust estimation of the parameter of interest

Immediate additional applications

- Optimal Dose Finding:
  - Wish to find the optimal dose of an intervention.
  - Parameter of Interest: low-dimensional summary of the does-response curve.
  - Optimal dose is a function of this parameter.
  - Adaptation criterion minimizes variance for estimating this function.
- Intervention packages with multiple components
  - Compare different components.
- Longitudinal Data Structure:
  - Time-varying treatment and covariates.

Generalizing Framework to other Adaptions

Using the missing data frame work from causal inference

- Wished/ideal full data structure: underlying i.i.d. $X_1, \ldots, X_n$.
  - $X=(W, Y(a): a \text{ in } \mathcal{A})$, with BL covariates W, and outcome Y(a) under different 'design settings a in $\mathcal{A}$'.
  - e.g. $\mathcal{A}$ is the set of interventions of interest.
- Observed data: For the k-th sample, we observe $O_k=(A_k,X(A_k))=(W, A_k,Y(A_k))$, where A_k is the design setting randomly drawn from $\mathcal{A}$ according to g_k, and $Y(A_k)$ is the outcome under this design setting.
- Assume draw of design setting only depends on previous observations $O_1, \ldots, O_{k-1}$

Generalizing Framework to other adaptations

- Treatment allocation: $\mathcal{A}$ is the set of interventions of interest. Each intervention can also have multiple components.
- Monitoring Intensity: design settings in $\mathcal{A}$ can also include monitoring indicators. Can adapt how often is monitoring.
- Collect new covariates: *choice of these covariates must be either a-priori specified or informed by external sources alone. Design settings in $\mathcal{A}$ can also include collection indicators. Can adapt when to start collection, but not the choice.
- Adapt clinical endpoint: instead of VL at 1 month, may want to look at VL at 2months, and adapt follow-up time from 1 month to 2 months. Alternative endpoint must be a-priori specified. Multiple testing needed.

Targeted CARA Design

- Initiate with a fixed design on first m patients.
- Given k-1 observations (O_1,..., O_{k-1}) and allocations (g_1, ..., g_{k-1}), to allocate the k-th patient:

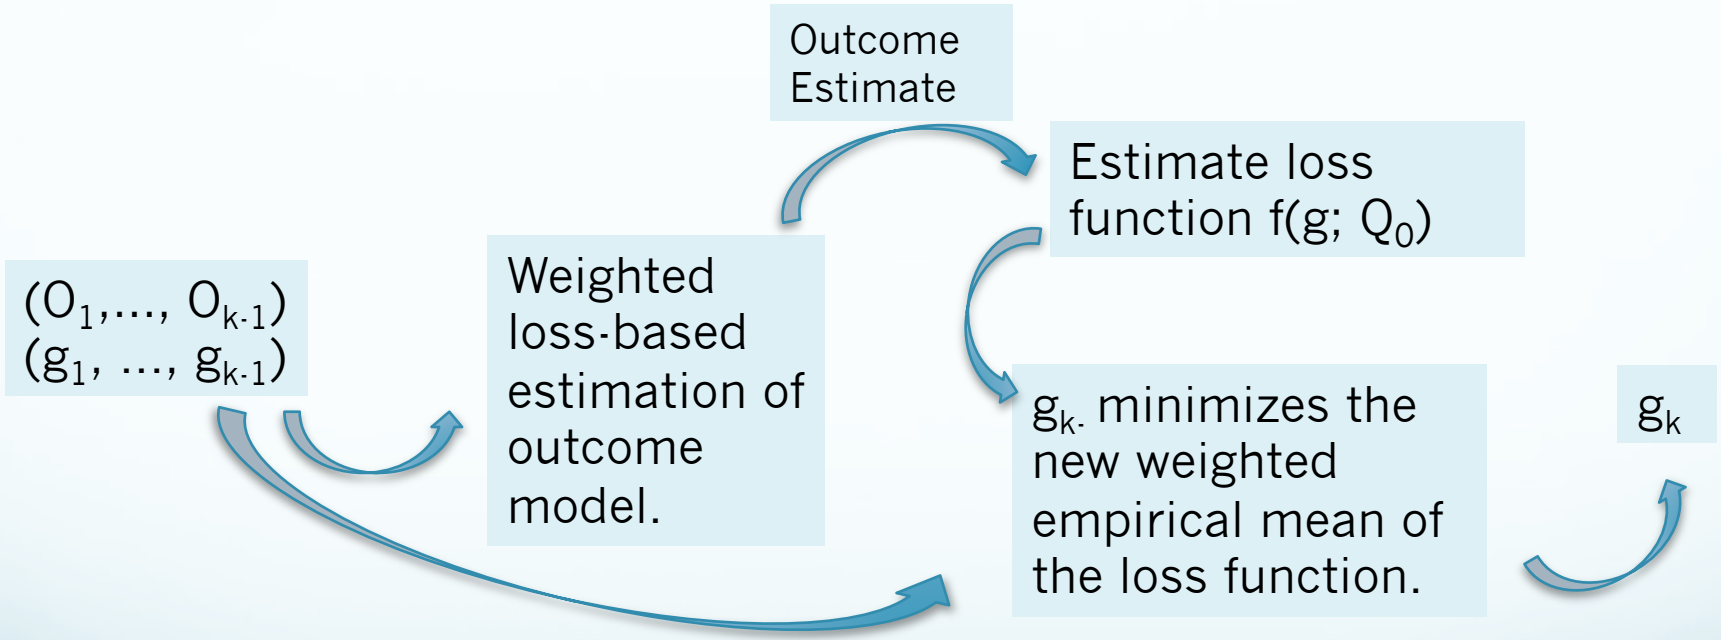

- Models can increase with k. Nonparam. or Semipara.
- Once specified how to build each g_k, we have specified the whole CARA design.

Targeted MLE for CARA

Estimate marginal treatment effect based on accrued n observations

1. Initial estimator of the outcome. Q_n.
2. Update Q_n based on a one-dimensional model that regresses on a specific function (based on efficient score equation) of the allocation g_n.
3. This model is fitted by weighted loss-based estimation with weights given by the allocations.
4. Obtain updated outcome estimator $Q^*_n(A,W)$ for $E(Y|A,W)$.
5. Targeted MLE for the marginal treatment effect is given by

$$\varphi_n^* \equiv \frac{1}{n}\sum_{i=1}^{n} Q_n^*(A=1,W_i) - Q_n^*(A=0,W_i)$$

This Targeted MLE solves the martingale efficient score equation that is doubly robust.

Overview of Asymptotic Results

- Theoretical Results:
  - Uniform Law of Large Numbers under CARA sampling.
  - General entropy conditions for martingale processes under CARA sampling.

- Practical Results:
  - Convergence of the adaptive allocations when complexities of the models are ‘controlled’.
  - Consistent estimation of parameter even under arbitrary model mis-specification
  - Targeted MLE asymptotically normal, if complexities of models are ‘reasonably well-controlled’.

Simulation Study

- Underlying data generation:
  - many covariates.
  - true conditional mean and conditional variance of outcome only depend on a few.
  - continuous and unbounded outcome.
- Compare Targeted MLE with outcome models:
  - Param Q: fixed parametric model that misses one relevant covariate.
  - Lasso Q: increasing LASSO, will eventually include the relevant covariates.
- Randomization Schemes:
  - mis-specified: misses one relevant covariate.
  - correctly specified.
- Parameter of Interest: marginal treatment effect.
- Goal of adaptation: maximize efficiency = minimize variance of the estimator.
- 500 runs to assess bias, variance and MSE of each Targeted MLE

Simulation Study of Targeted MLE

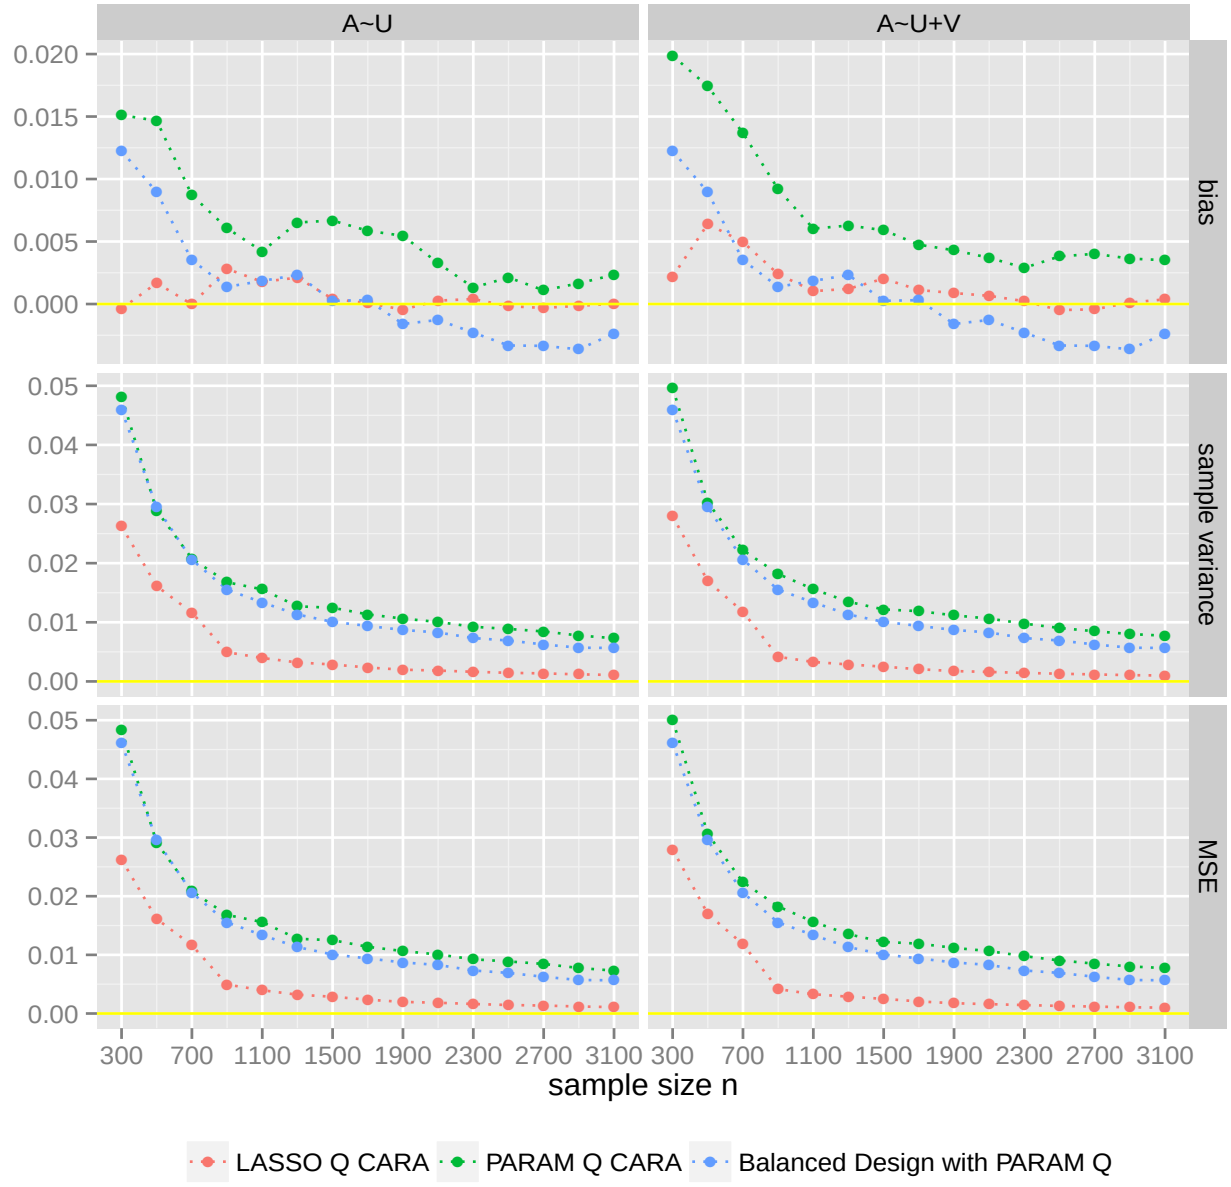

- Consistent estimator despite outcome model mis-spec.
- LASSO Q offers more aggressive bias reduction.
- Optimality criterion=mini. Variance: CARA +LASSO Q better.

Concluding Remarks

- A general design and analysis framework for Covariate-Adjusted Response-Adaptive randomization designs.
- Allows general parameters of interest, optimality criterion.
- Allows machine learning/data-adaptive estimators for outcome models.
- Robust parameter estimate under mis-specified models. Asymptotically normal estimator.
- Statistical framework can be generalized to
- Longitudinal data, multi-component intervention packages.
- Optimal dose finding.
- Modify monitoring intensity, variable collection, clinical endpoints, etc.

wenjing.zheng@ucsf.edu

This work was funded by

- The Traineeship in AIDS Prevention Studies by CAPS, UCSF.
- Chateaubriand Fellowship in STEM by the Republic of France.

Thank you!