Set-valued dynamic treatment regimes for competing outcomes

Eric B. Laber

Department of Statistics, North Carolina State University

JSM, Montreal, QC, August 5, 2013

Acknowledgments

- Zhiqiang Tan
- Jamie Robins
- Dan Lizotte
- Brad Ferguson
- NCSU and UNC DTR Groups

Dynamic treatment regimes

- Motivation: treatment of chronic illness
  - Ex. ADHD, HIV/AIDS, cancer, depression, schizophrenia, drug and alcohol addiction, bipolar disorder, etc.
  - Treatments adapt to evolving health status of patient
  - Requires balancing immediate and long-term outcomes
- Dynamic treatment regimes (DTRs)
  - Operationalize decision process as sequence of decision rules
    - One decision rule for each intervention time
    - Intervention times often outcome-dependent
    - Decision rule maps up-to-date patient history to a recommended treatment
  - Maximize expectation of a single clinical outcome

Estimation of DTRs

- Many options for estimating DTRs from observational or randomized studies
  - Q- and A-learning (Murphy 2003, 2005; Robins 2004)
  - Regret regression (Henderson et al., 2010)
  - Dynamical systems models (Rivera et al., 2007; Navarro-Barrientos et al., 2010, 2011)
  - Policy search
    - Augmented value maximization (Zhang 2012, 2013ab)
    - Outcome weighted learning (Zhao et al., 2012, 2013ab)
    - Marginal structural mean models (Orellana et al., 2010)
- All aim optimize the mean of a single scalar outcome by which the 'goodness' of competing DTRs are measured

Competing outcomes

- Specifying a single scalar outcome sometimes oversimplifies the goal of clinical decision making
  - Need to balance several potentially competing outcomes, e.g., symptom relief and side-effect burden
  - Patients may have heterogeneous and evolving preferences across competing outcomes
  - In some populations, patients do not know or cannot communicate their preferences

Motivating example

- CATIE Schizophrenia study (Stroup et. al., 2003)
  - 18 month SMART study with two main phases
  - Aim to compare typical and atypical antipsychotics
  - Trade-off between efficacy and side-effect burden
    - Drugs known to be efficacious also known to have severe negative side-effects (e.g., olanzapine; Brier et al., 2005)
    - Separate arms for efficacy and tolerability at second stage
- Sequential treatment of severe schizophrenia
  - Preferences vary widely across patients (Kinter, 2009)
    - Ex. Joe may be willing to tolerate severe weight gain, lethargy, and reduced social functioning to gain symptom relief.

      Joan cannot function at work if she is lethargic or experiences clouded thinking and is thus willing to take a less effective treatment if it has low side-effect burden.
  - Preference elicitation is difficult in this population
  - Preferences may evolve over time (Strauss et al., 2011)

Composite outcomes

- A natural approach is to combine competing outcomes into a single composite outcome
  - A single composite outcome for all patients (e.g., Wang et al., 2012)
    - Elicited by panel of experts
    - Requires patient preference homogeneity (e.g., all patients are ambivalent about trading 1 side-effect unit for three units of symptom relief)
    - Preferences do not change over time
  - Estimate the optimal DTR across all convex combinations of competing outcomes simultaneously (e.g., Lizotte et al., 2012)
    - 'True preferences' must be expressible as convex combinations
    - Patient preferences cannot change over time

Setup: Two-stage binary treatments

- For simplicity, we consider the two-stage binary treatment setting, this is not essential (see Laber et al., 2013)
- Observe $\{(X_{1i}, A_{1i}, X_{2i}, A_{2i}, E_i, S_i)\}_{i=1}^n$ comprising n *iid* patient trajectories :
  - $X_t \in \mathbb{R}^{p_t}$: subject covariate vector prior to tth txt
  - $A_t \in \mathcal{A}_t = \{-1, 1\}$: treatment received at time t
  - $E \in \mathbb{R}$: first outcome, coded so that higher is better
  - $S \in \mathbb{R}$: second outcome, coded so that higher is better
  - Define: $H_1 = X_1$ and $H_2 = (X_1^\intercal, A_1, X_2^\intercal)^\intercal$, let $\mathcal{H}_t$ denote the domain of H_t
- Focus on estimands

Setup: DTRs vs SVDTRs

- A DTR is a pair of functions $\pi = (\pi_1, \pi_2)$ where $\pi_t : \mathcal{H}_t \to \mathcal{A}_t$
  - Under DTR π a patient presenting with $H_t = h_t$ at time t is recommended treatment $\pi_t(h_t)$
  - For scalar summary $Y = Y(E, S)$ of (E, S) the optimal DTR π^{opt} satisfies $\mathbb{E}^{\pi^{\text{opt}}} Y \geq \mathbb{E}^{\pi} Y$ for all π, where $\mathbb{E}^{\pi}$ denotes expectation under treatment assignment via π
- An SVDTR is a pair of functions $\pi = (\pi_1, \pi_2)$ where $\pi_t : \mathcal{H}_t \to 2^{\mathcal{A}_t}$
  - Under SVDTR π a patient presenting with $H_t = h_t$ at time t is offered treatment choices $\pi_t(h_t) \subseteq \mathcal{A}_t$
  - Difficult to define optimality without assumptions about patient utility; we will define 'ideal' decision rules
  - View SVDTRs as a communication tool, typically paired with graphical displays (more on this later)

Review: dynamic programming for optimal DTR

- When underlying generative distribution is known, optimal DTR π^{opt} can be found via dynamic programming
  1. Define $Q_{2Y}(h_2, a_2) \triangleq \mathbb{E}(Y|H_2 = h_2, A_2 = a_2)$
  2. $Y^* \triangleq \max_{a_2} Q_{2Y}(H_2, a_2) = Q_{2Y}(H_2, \pi_2^{\text{opt}}(H_2))$
  3. $Q_{1Y}(h_1, a_1) \triangleq \mathbb{E}(Y^*|H_1 = h_1, A_1 = a_1)$

  then $\pi_t^{\text{opt}}(h_t) = \arg\max_{a_t \in \mathcal{A}_t} Q_{tY}(h_t, a_t)$ (Bellman, 1957)
- Q-learning mimics dynamic programming but uses regression models in place of the requisite conditional expectations

Note* subscript denotes outcome used to define Q-function

SVDTRs

- Incorporate clinically significant differences Δ_E, $\Delta_S > 0$; a difference in $E(S)$ of less than $\Delta_E(\Delta_S)$ is not considered clinically significant (e.g., Friedman et al., 2010)
- Define contrast functions
  - $r_{2E}(h_2) \triangleq \mathbb{E}(E|H_2 = h_2, A_2 = 1) - \mathbb{E}(E|H_2 = h_2, A_2 = -1)$
  - $r_{2S}(h_2) \triangleq \mathbb{E}(S|H_2 = h_2, A_2 = 1) - \mathbb{E}(S|H_2 = h_2, A_2 = -1)$
- Ideal set-valued second stage decision rule

$$\pi_{2\Delta}^{\text{Ideal}}(h_2) = \begin{cases} \{\text{sgn}(r_{2E}(h_2))\}, & \text{if } |r_{2E}(h_2)| \geq \Delta_E \\ & \quad \text{and } \text{sgn}(r_{2E}(h_2))r_S(h_2) > -\Delta_S, \\ \{\text{sgn}(r_{2S}(h_2))\}, & \text{if } |r_{2S}(h_2)| \geq \Delta_S \\ & \quad \text{and } \text{sgn}(r_{2S}(h_2))r_{2E}(h_2) > -\Delta_E, \\ \{-1, 1\}, & \text{otherwise,} \end{cases}$$

SVDTRs: schematic for $\pi_{2\Delta}^{\text{Ideal}}$

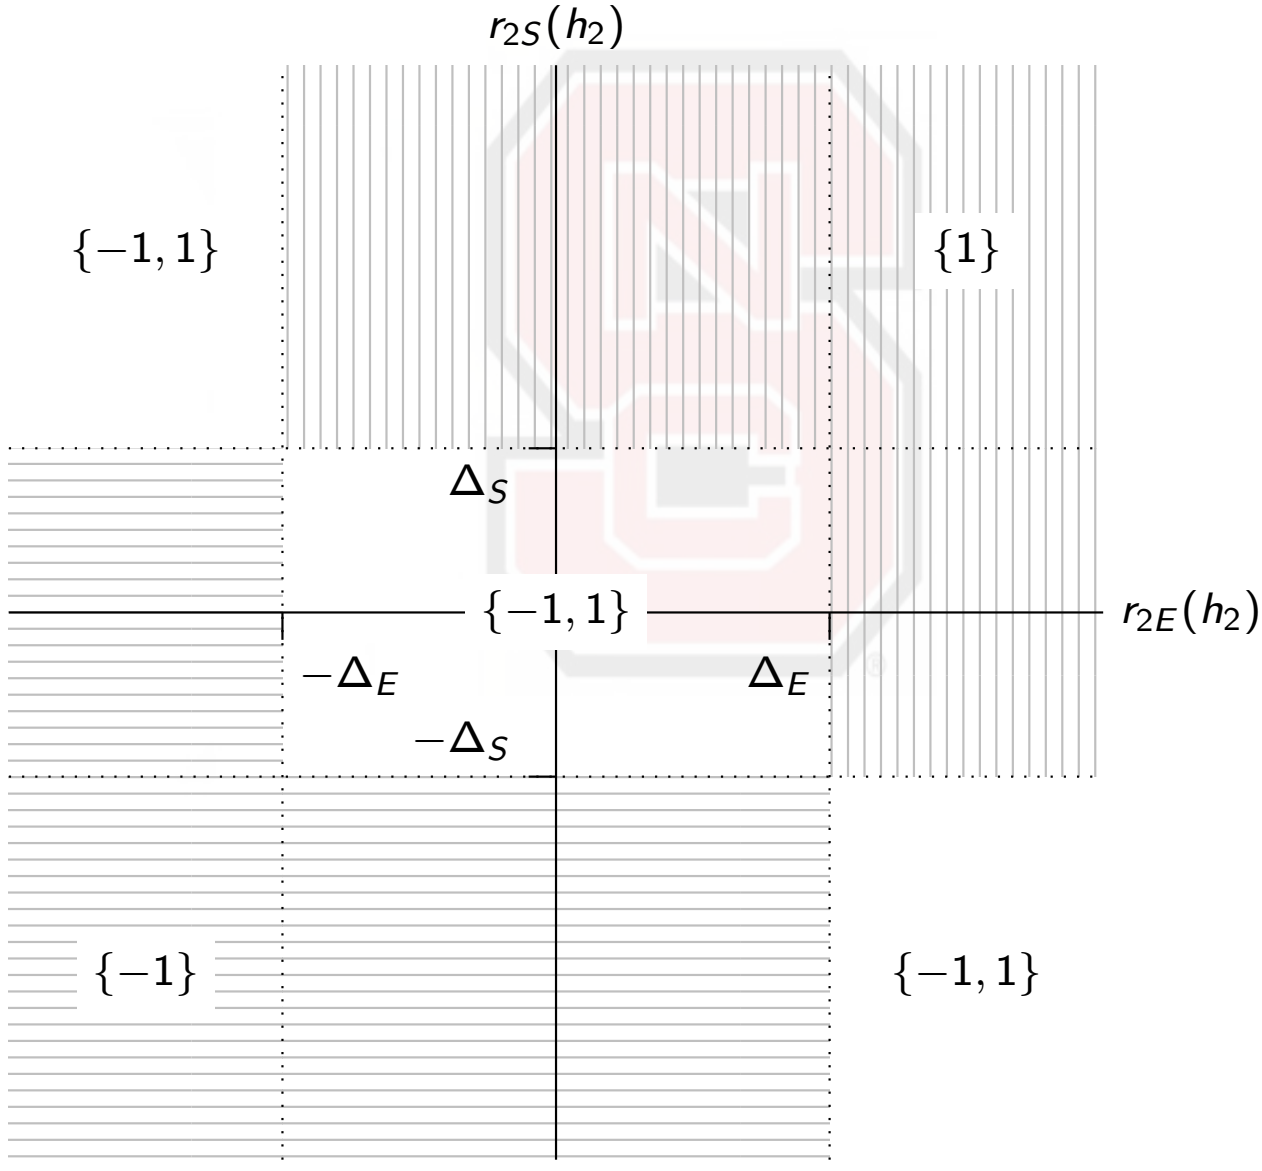

SVDTRs: defining the ideal first stage rule

- Q-learning requires a fixed second stage rule for backup step
  - Problem: $\pi_{2\Delta}^{\text{Ideal}}$ is not a decision rule
  - Idea:
    1. Define a set of 'sensible' second stage rules $\mathcal{S}(\pi_{2\Delta}^{\text{Ideal}})$
    2. For each $\tau_2 \in \mathcal{S}(\pi_{2\Delta}^{\text{Ideal}})$ derive ideal set-valued rule $\pi_{1\Delta}^{\text{Ideal}}(\cdot, \tau_2)$ assuming τ_2 will be followed at the second stage
    3. A patient presenting with $H_1 = h_1$ is offered treatments

$$\pi_{1\Delta}^{\text{Ideal}}(h_1) \triangleq \bigcup_{\tau_2 \in \mathcal{S}(\pi_{2\Delta}^{\text{Ideal}})} \pi_{1\Delta}^{\text{Ideal}}(h_1, \tau_2)$$

SVDTRs: defining a 'sensible' set

- Decision rule τ_2 is compatible with set-valued rule π_2 if $\tau_2(h_2) \in \pi_2(h_2)$ for all $h_2 \in \mathcal{H}_2$
  - Let $\mathcal{C}(\pi_2)$ denote the set of all policies compatible with π_2
  - Minimally require $\mathcal{S}(\pi_{2\Delta}^{\text{Ideal}}) \subseteq \mathcal{C}(\pi_{2\Delta}^{\text{Ideal}})$
  - $\mathcal{C}(\pi_{2\Delta}^{\text{Ideal}})$ is large and contains many unrealistic rules
- Contrast models suggest functional form for decision rules

$$\mathcal{F} \triangleq \left\{ \tau_2 \,:\, \exists \rho \in \mathbb{R}^{p_2} \; s.t. \; \tau_2(h_2) = \text{sgn}(h_{2,2}^{\intercal} \rho) \right\}$$

- Define the set of sensible second stage rules as $\mathcal{S}(\pi_{2\Delta}^{\text{Ideal}}) \triangleq \mathcal{F} \cap \mathcal{C}(\pi_{2\Delta}^{\text{Ideal}})$

SVDTRs: ideal set-valued first stage rule

- For a fixed second stage rule τ_2 define
  - $Q_{1E}(h_1, a_1, \tau_2) \triangleq \mathbb{E}\{Q_{2E}(H_2, \tau_2(H_2))|H_1 = h_1, A_1 = a_1\}$
  - $r_{1E}(h_1, \tau_2) \triangleq Q_{1E}(h_1, 1, \tau_2) - Q_{1E}(h_1, -1, \tau_2)$
  - Q_{1S} and r_{1S} defined analogously
- Ideal set-valued first stage rule under τ_2 at the second stage:

$$\pi_{1\Delta}^{\text{Ideal}}(h_1, \tau_2) = \begin{cases} \{\text{sgn}(r_{1E}(h_1, \tau_2))\}, & \text{if } |r_{1E}(h_1, \tau_2)| \geq \Delta_E \text{ and} \\ & \text{sgn}(r_{1E}(h_1, \tau_2))r_{1S}(h_1, \tau_2) > -\Delta_S, \\ \{\text{sgn}(r_{1S}(h_1, \tau_2))\}, & \text{if} |r_{1S}(h_1, \tau_2)| \geq \Delta_S \text{ and} \\ & \text{sgn}(r_{1S}(h_1, \tau_2))r_{1E}(h_1, \tau_2) > -\Delta_E, \\ \{-1, 1\}, & \text{otherwise}, \end{cases}$$

SVDTRs: Defining $\pi_{1\Delta}^{\text{Ideal}}$

- $\pi_{1\Delta}^{\text{Ideal}}(\cdot, \tau_2)$ assigns a single treatment if that treatment is expected to yield a clinically significant improvement on one or both the outcomes while not causing clinically significant loss in either outcome assuming the clinician will follow τ_2 at the second decision point.
- Recall: Ideal decision rule at the first stage

$$\pi_{1\Delta}^{\text{Ideal}}(h_1) \triangleq \bigcup_{\tau_2 \in \mathcal{S}(\pi_{2\Delta}^{\text{opt}})} \pi_{1\Delta}^{\text{Ideal}}(h_1, \tau_2)$$

SVDTR estimation

- Linear models; H_{t0}, H_{t1} summaries of H_t
  1. Regress E and S on H_{21} and H_{22} to obtain
     $\widehat{Q}_{2E}(H_2, A_2) = \widehat{\beta}_{21E}^{\mathsf{T}} H_{21} + \widehat{\beta}_{22E}^{\mathsf{T}} H_{22} A_2$
     $\widehat{Q}_{2S}(H_2, A_2) = \widehat{\beta}_{21S}^{\mathsf{T}} H_{21} + \widehat{\beta}_{22S}^{\mathsf{T}} H_{22} A_2$
  2. Define $\widehat{r}_{2E}(h_2) \triangleq \widehat{Q}_{2E}(h_2, 1) - \widehat{Q}_{2E}(h_2, -1)$
     $\widehat{r}_{2S}(h_2) \triangleq \widehat{Q}_{2S}(h_2, 1) - \widehat{Q}_{2S}(h_2, -1)$
  3. Use the plug-in estimator of $\pi_{2\Delta}^{\text{Ideal}}$ to obtain

$$\widehat{\pi}_{2\Delta}(h_2) = \begin{cases} \{\text{sgn}(\widehat{r}_{2E}(h_2))\}, & \text{if } |\widehat{r}_{2E}(h_2)| \geq \Delta_E \\ & \quad \text{and } \text{sgn}(\widehat{r}_{2E}(h_2))\widehat{r}_S(h_2) > -\Delta_S, \\ \{\text{sgn}(\widehat{r}_{2S}(h_2))\}, & \text{if } |\widehat{r}_{2S}(h_2)| \geq \Delta_S \\ & \quad \text{and } \text{sgn}(\widehat{r}_{2S}(h_2))\widehat{r}_{2E}(h_2) > -\Delta_E, \\ \{-1, 1\}, & \text{otherwise}, \end{cases}$$

SVDTR estimation cont'd

4.* Construct set $\mathcal{S}(\widehat{\pi}_{2\Delta})$

5. For each τ_2 in $\mathcal{S}(\widehat{\pi}_{2\Delta})$ regress $\widehat{Q}_{2E}(H_2, \tau_2(H_2))$ and $\widehat{Q}_{2S}(H_2, \tau_2(H_2))$ on H_{11}, H_{11} and A_1 to obtain
$\widehat{Q}_{1E}(H_1, A_1; \tau_2) = \widehat{\beta}_{11E}(\tau_2)^\intercal H_{11} + \widehat{\beta}_{12E}(\tau_2)^\intercal H_{12}$
$\widehat{Q}_{1S}(H_1, A_1; \tau_2) = \widehat{\beta}_{11S}(\tau_2)^\intercal H_{11} + \widehat{\beta}_{12S}(\tau_2)^\intercal H_{12}$

6. Define $\widehat{r}_{1E}(h_1, \tau_2) \triangleq \widehat{Q}_{1E}(h_1, 1, \tau_2) - \widehat{Q}_{1E}(h_1, -1, \tau_2)$
$\widehat{r}_{1S}(h_1, \tau_2) \triangleq \widehat{Q}_{1S}(h_1, 1, \tau_2) - \widehat{Q}_{1S}(h_1, -1, \tau_2)$

7. Use the plugin-estimator of $\pi_{1\Delta}^{\text{Ideal}}$ to obtain $\widehat{\pi}_{1\Delta}(h_1, \tau_2)$ (similar to step 3)

8. Define $\widehat{\pi}_{1\Delta}(h_1) \triangleq \bigcup_{\tau_2 \in \mathcal{S}(\widehat{\pi}_{2\Delta})} \widehat{\pi}_{1\Delta}(h_1, \tau_2)$

SVDTR computation

- Constructing the set $\mathcal{S}(\widehat{\pi}_{2\Delta})$ is a seemingly difficult enumeration problem
  - While $\mathcal{S}(\widehat{\pi}_{2\Delta})$ is uncountably infinite, one can prove there exists a finite set $\mathcal{S}'(\widehat{\pi}_{2\Delta})$ so that for all h_1:

$$\bigcup_{\tau_2 \in \mathcal{S}'(\widehat{\pi}_{2\Delta})} \widehat{\pi}_{1\Delta}(h_1, \tau_2) = \bigcup_{\tau_2 \in \mathcal{S}(\widehat{\pi}_{2\Delta})} \widehat{\pi}_{1\Delta}(h_1, \tau_2)$$

  - We cast construction of $\mathcal{S}'(\widehat{\pi}_{2\Delta})$ as a linear Mixed Integer Program (MIP)
    - Computed efficiently using CPLEX
    - With CATIE data ($n \approx 1100$, $p \approx 20$) runtime is less than 1 minute on a dual core laptop (2.8GHz, 8GiB)

Communicating the SVDTR

- Patient presenting at the second stage with $H_2 = h_2$ given:
  1. $\widehat{\pi}_{2\Delta}(h_2)$
  2. Estimates $(\widehat{Q}_{2E}(h_2, 1), \widehat{Q}_{2S}(h_2, 1))$, $(\widehat{Q}_{2S}(h_2, -1), \widehat{Q}_{2E}(h_2, -1))$
- Patient presenting at the first stage with $H_1 = h_1$ given:
  1. $\widehat{\pi}_{1\Delta}(h_1)$
  2. Plot of $\widehat{Q}_{1E}(h_1, a_1, \tau_2)$ against $\widehat{Q}_{1S}(h_1, a_1, \tau_2)$ across $\tau_2 \in \mathcal{S}(\widehat{\pi}_{1\Delta})$ with separate plotting symbols/colors for $a_1 = \pm 1$

First stage plot for 'mean' CATIE subject

- Estimated SVDTR using data from the CATIE study; outcomes $E = PANSS$ and $S = BMI$ (details in Laber et al., 2013)
- Constructed plot for a subject with $h_1 = n^{-1} \sum_{i=1}^{n} H_{1i}$

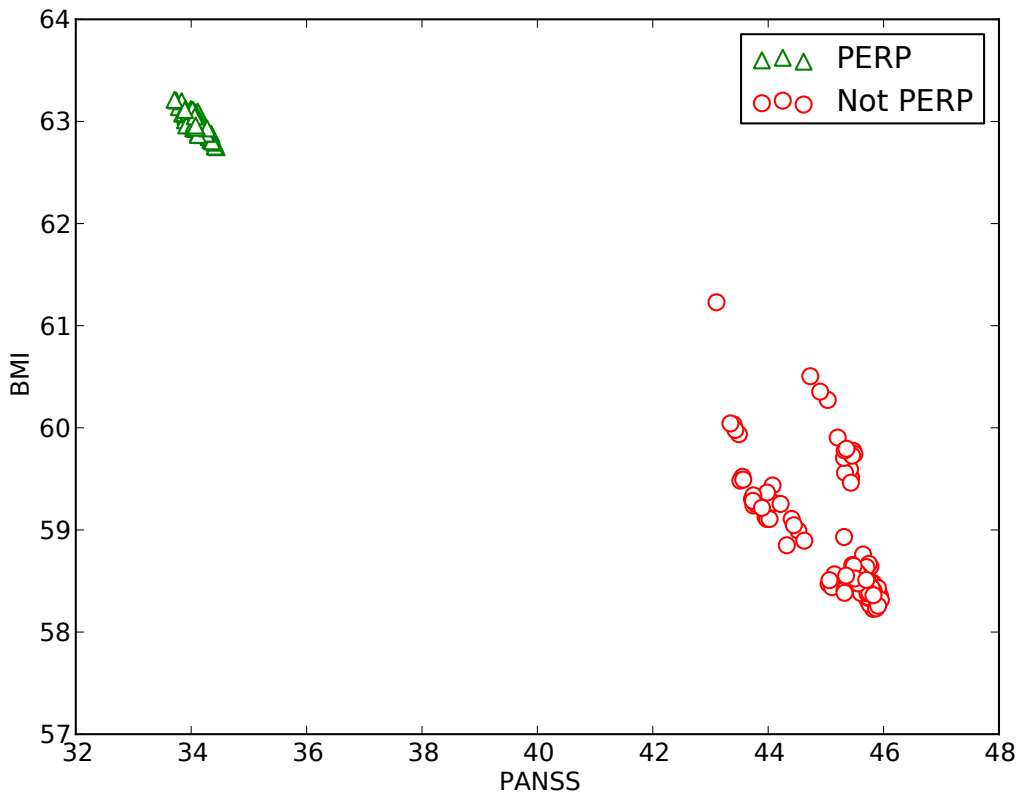

Discussion

- Introduced set-valued dynamic treatment regimes (SVDTRs)
- One approach to dealing with competing outcomes under:
  - Patient preference heterogeneity
  - Difficult to elicit preferences
  - Evolving patient preference
- Presented the two-stage binary treatment case, method extends to an arbitrary number of txts and stages (see Laber et al., 2013)
- Suggests a new framework for decision problems with both a decision maker and a treatment 'screener.' This framework generalized MDPs and may facilitate minimax and other definitions of optimality for SVDTRs.

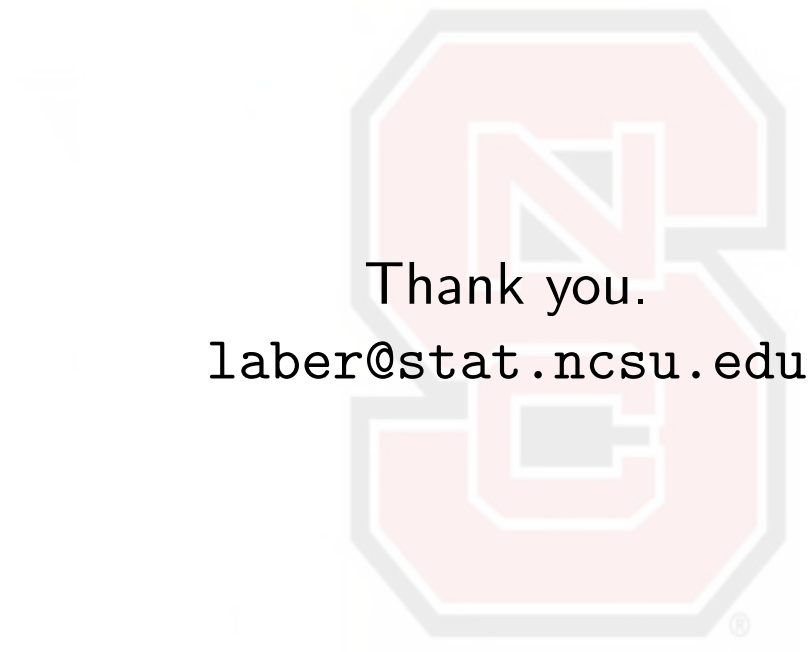

Thank you.

laber@stat.ncsu.edu