Ranking and Selection with Covariates

Haihui SHEN

Department of Management Sciences
College of Business
City University of Hong Kong

Joint work with L. Jeff Hong (CityU) and Xiaowei Zhang (HKUST)

The Ninth POMS-HK International Conference

Contents

Ranking and Selection (R&S)

- Conventional R&S aims to find the alternative with the largest mean performance (μ) from a group of k alternatives, i.e.,

$$i^* := \arg\max_{1 \le i \le k} \mu_i.$$

- The unknown μ_i can only be learned through noisy samples, from either computer simulation or real experiments.

Ranking and Selection (R&S)

- Conventional R&S aims to find the alternative with the largest mean performance (μ) from a group of k alternatives, i.e.,

$$i^* := \arg\max_{1 \le i \le k} \mu_i.$$

- The unknown μ_i can only be learned through noisy samples, from either computer simulation or real experiments.

Ranking and Selection (R&S)

- Conventional R&S aims to find the alternative with the largest mean performance (μ) from a group of k alternatives, i.e.,

$$i^* := \arg\max_{1 \le i \le k} \mu_i.$$

- The unknown μ_i can only be learned through noisy samples, from either computer simulation or real experiments.

- Unless the sample size of each alternative goes to infinity, there is no guarantee that we indeed find i^*.

- We are satisfied if the probability of being correct is larger than some specified value (e.g., $\geq 95\%$).

Covariates

- Covariates ($\boldsymbol{X}$) are also known as personalized information, side information, auxiliary quantities or contextual variables.
- In many cases, the performance of alternative i is $\mu_i(\boldsymbol{x})$, given $\boldsymbol{X} = \boldsymbol{x}$.

Covariates

- Covariates ($\boldsymbol{X}$) are also known as personalized information, side information, auxiliary quantities or contextual variables.
- In many cases, the performance of alternative i is $\mu_i(\boldsymbol{x})$, given $\boldsymbol{X} = \boldsymbol{x}$.
- Examples:
  1. Healthcare: *Personalized medicine*.
  2. Marketing: *Personalized recommendation/promotion*.

Covariates

- Covariates ($\boldsymbol{X}$) are also known as personalized information, side information, auxiliary quantities or contextual variables.
- In many cases, the performance of alternative i is $\mu_i(\boldsymbol{x})$, given $\boldsymbol{X} = \boldsymbol{x}$.
- Examples:
  1. Healthcare: *Personalized medicine*.
  2. Marketing: *Personalized recommendation/promotion*.
- Covariates allow decisions to be made at individual level.

Ranking and Selection with Covariates (R&S-C)

- Conventional R&S + Covariates $\Rightarrow$ R&S-C.

Ranking and Selection with Covariates (R&S-C)

- Conventional R&S + Covariates $\Rightarrow$ R&S-C.
- Overview of R&S-C in simulation:
  - Performance of an alternative depends on some observable random covariates;
  - Performance can be sampled (learnt) through simulation;
  - The best alternative is a *function* of the covariates;
  - A selection procedure is required to produce a decision rule (i.e., an estimator of the *function*).

Ranking and Selection with Covariates (R&S-C)

- Conventional R&S + Covariates $\Rightarrow$ R&S-C.
- Overview of R&S-C in simulation:
  - Performance of an alternative depends on some observable random covariates;
  - Performance can be sampled (learnt) through simulation;
  - The best alternative is a *function* of the covariates;
  - A selection procedure is required to produce a decision rule (i.e., an estimator of the *function*).
- Why is R&S-C necessary?

Related Literature

- Traditional R&S:
  - Frequentist approaches: Dudewicz and Dalal (1975), Rinott (1978), Kim and Nelson (2001), Hong (2006), etc.
  - Bayesian approaches: Chen et al. (1997), Chick and Inoue (2001), Frazier et al. (2008), Chick and Frazier (2012), etc.
- Multi-armed bandit (MAB) with covariates:
  - Parametric bandits: Auer (2002), Rusmevichientong and Tsitsiklis (2010), Goldenshluger and Zeevi (2013), etc.
  - Non-parametric bandits: Rigollet and Zeevi (2010), Perchet and Rigollet (2013), Slivkins (2014), etc.
- R&S with covariates:
  - Not yet defined and studied.
  - Our work serves as an attempt to fill in the gap.

Problem

- k alternatives;
- Random covariates $\boldsymbol{X}_{\mathsf{c}} = (X_1, \ldots, X_d)^\top$ with support $\Theta_{\mathsf{c}} \subseteq \mathbb{R}^d$;

Problem

- k alternatives;
- Random covariates $\boldsymbol{X}_{\mathsf{c}} = (X_1, \ldots, X_d)^\top$ with support $\Theta_{\mathsf{c}} \subseteq \mathbb{R}^d$;
- Let $\boldsymbol{X} := (1, \boldsymbol{X}_{\mathsf{c}}^\top)^\top$ be the augmented covariates with support $\Theta := \{1\} \times \Theta_{\mathsf{c}}$;

Problem

- k alternatives;
- Random covariates $\boldsymbol{X}_{\mathsf{c}} = (X_1, \ldots, X_d)^\top$ with support $\Theta_{\mathsf{c}} \subseteq \mathbb{R}^d$;
- Let $\boldsymbol{X} := (1, \boldsymbol{X}_{\mathsf{c}}^\top)^\top$ be the augmented covariates with support $\Theta := \{1\} \times \Theta_{\mathsf{c}}$;
- $\mu_i(\boldsymbol{X})$ is the mean performance of alternative i, for $i = 1, 2, \ldots, k$.

Problem

- k alternatives;
- Random covariates $\boldsymbol{X}_{\mathsf{c}} = (X_1, \ldots, X_d)^\top$ with support $\Theta_{\mathsf{c}} \subseteq \mathbb{R}^d$;
- Let $\boldsymbol{X} := (1, \boldsymbol{X}_{\mathsf{c}}^\top)^\top$ be the augmented covariates with support $\Theta := \{1\} \times \Theta_{\mathsf{c}}$;
- $\mu_i(\boldsymbol{X})$ is the mean performance of alternative i, for $i = 1, 2, \ldots, k$.
- R&S-C hopes to find

$$i^*(\boldsymbol{x}) := \arg\max_{1 \le i \le k} \{\mu_i(\boldsymbol{X}) | \boldsymbol{X} = \boldsymbol{x}\}.$$

Why R&S-C Instead of R&S?

- Recall:

R&S		R&S-C
$i^\dagger := \arg\max_{1\le i\le k} \mu_i$	vs	$i^*(\boldsymbol{x}) := \arg\max_{1\le i\le k} \{\mu_i(\boldsymbol{X}) \mid \boldsymbol{X} = \boldsymbol{x}\}$

Why R&S-C Instead of R&S?

- Recall:

R&S		R&S-C
$i^\dagger := \arg\max_{1\le i\le k} \mu_i$	vs	$i^*(\boldsymbol{x}) := \arg\max_{1\le i\le k}\{\mu_i(\boldsymbol{X})\|\boldsymbol{X}=\boldsymbol{x}\}$

- If we have to use R&S...
  1. Let $\mu_i = \mathbb{E}[\mu_i(\boldsymbol{X})]$, if we are risk-neutral with respect to $\boldsymbol{X}$.

Why R&S-C Instead of R&S?

- Recall:

R&S

$$i^\dagger := \arg\max_{1\le i\le k} \mu_i$$

vs

R&S-C

$$i^*(\boldsymbol{x}) := \arg\max_{1\le i\le k} \{\mu_i(\boldsymbol{X})|\boldsymbol{X}=\boldsymbol{x}\}$$

- If we have to use R&S...
  1. Let $\mu_i = \mathbb{E}[\mu_i(\boldsymbol{X})]$, if we are risk-neutral with respect to $\boldsymbol{X}$.

     However, $\mathbb{E}[\mu_{i^\dagger}(\boldsymbol{X})] \le \mathbb{E}[\mu_{i^*(\boldsymbol{X})}(\boldsymbol{X})]$.

Why R&S-C Instead of R&S?

- Recall:

R&S

$$i^\dagger := \arg\max_{1\le i\le k} \mu_i$$

vs

R&S-C

$$i^*(\boldsymbol{x}) := \arg\max_{1\le i\le k} \{\mu_i(\boldsymbol{X})|\boldsymbol{X}=\boldsymbol{x}\}$$

- If we have to use R&S...
  1. Let $\mu_i = \mathbb{E}[\mu_i(\boldsymbol{X})]$, if we are risk-neutral with respect to $\boldsymbol{X}$.

     However, $\mathbb{E}[\mu_{i^\dagger}(\boldsymbol{X})] \le \mathbb{E}[\mu_{i^*(\boldsymbol{X})}(\boldsymbol{X})]$.
  2. Solve the problem online, i.e., upon observing a coming customer with $\boldsymbol{X}=\boldsymbol{x}_1$, start to solve $\arg\max_{1\le i\le k}\mu_i(\boldsymbol{x}_1)$; next time $\boldsymbol{X}=\boldsymbol{x}_2$, ...

Why R&S-C Instead of R&S?

- Recall:

R&S

$$i^\dagger := \arg\max_{1\le i\le k} \mu_i$$

vs

R&S-C

$$i^*(\boldsymbol{x}) := \arg\max_{1\le i\le k} \{\mu_i(\boldsymbol{X})|\boldsymbol{X}=\boldsymbol{x}\}$$

- If we have to use R&S...
  1. Let $\mu_i = \mathbb{E}[\mu_i(\boldsymbol{X})]$, if we are risk-neutral with respect to $\boldsymbol{X}$.

     However, $\mathbb{E}[\mu_{i^\dagger}(\boldsymbol{X})] \le \mathbb{E}[\mu_{i^*(\boldsymbol{X})}(\boldsymbol{X})]$.

  2. Solve the problem online, i.e., upon observing a coming customer with $\boldsymbol{X}=\boldsymbol{x}_1$, start to solve $\arg\max_{1\le i\le k}\mu_i(\boldsymbol{x}_1)$; next time $\boldsymbol{X}=\boldsymbol{x}_2$, ...

     Drawbacks:
     – Simulation models should always be available at hand.
     – Decision is not timely if simulation is time consuming.

Offline Learning and Online Application

- With the goal of delivering $i^*(\boldsymbol{x})$ for R&S-C, we propose to:
  1. Obtain a decision rule $\widehat{i^*}(\boldsymbol{x})$, which is an estimator of $i^*(\boldsymbol{x})$, from offline simulation;
  2. Apply $\widehat{i^*}(\boldsymbol{x})$ online to guide the actual selection in *real time*, i.e., upon observing a coming customer with $\boldsymbol{X} = \boldsymbol{x}_1$, report $\widehat{i^*}(\boldsymbol{x}_1)$ immediately; next time $\boldsymbol{X} = \boldsymbol{x}_2$, report $\widehat{i^*}(\boldsymbol{x}_2)$ ...

Offline Learning and Online Application

- With the goal of delivering $i^*(\boldsymbol{x})$ for R&S-C, we propose to:
  1. Obtain a decision rule $\widehat{i^*}(\boldsymbol{x})$, which is an estimator of $i^*(\boldsymbol{x})$, from offline simulation;
  2. Apply $\widehat{i^*}(\boldsymbol{x})$ online to guide the actual selection in *real time*, i.e., upon observing a coming customer with $\boldsymbol{X} = \boldsymbol{x}_1$, report $\widehat{i^*}(\boldsymbol{x}_1)$ immediately; next time $\boldsymbol{X} = \boldsymbol{x}_2$, report $\widehat{i^*}(\boldsymbol{x}_2)$ ...

- It is impractical to get $\widehat{i^*}(\boldsymbol{x})$ from enumerating all possible values of $\boldsymbol{X}$ and conducting conventional R&S at every point.

Offline Learning and Online Application

- With the goal of delivering $i^*(\boldsymbol{x})$ for R&S-C, we propose to:
  1. Obtain a decision rule $\widehat{i^*}(\boldsymbol{x})$, which is an estimator of $i^*(\boldsymbol{x})$, from offline simulation;
  2. Apply $\widehat{i^*}(\boldsymbol{x})$ online to guide the actual selection in *real time*, i.e., upon observing a coming customer with $\boldsymbol{X} = \boldsymbol{x}_1$, report $\widehat{i^*}(\boldsymbol{x}_1)$ immediately; next time $\boldsymbol{X} = \boldsymbol{x}_2$, report $\widehat{i^*}(\boldsymbol{x}_2)$ ...

- It is impractical to get $\widehat{i^*}(\boldsymbol{x})$ from enumerating all possible values of $\boldsymbol{X}$ and conducting conventional R&S at every point.

- Assuming $\mu_i(\boldsymbol{X})$ is linear in $\boldsymbol{X}$, we try to design some selection procedures to produce $\widehat{i^*}(\boldsymbol{x})$ using offline simulation.

Linear Model

Assumption 1 (A1)

For each $i = 1, \ldots, k$ *and* $\ell = 1, 2, \ldots$, *conditionally on* $\boldsymbol{X} = \boldsymbol{x}$,

$$\mu_i(\boldsymbol{x}) = \boldsymbol{x}^\top \boldsymbol{\beta}_i,$$
$$Y_{i\ell}(\boldsymbol{x}) = \mu_i(\boldsymbol{x}) + \epsilon_{i\ell}(\boldsymbol{x}),$$

where $\boldsymbol{\beta}_i = (\beta_{i0}, \beta_{i1}, \ldots, \beta_{id})^\top \in \mathbb{R}^{d+1}$ *is a vector of unknown parameters,* $Y_{i\ell}(\boldsymbol{x})$ *is the simulated sample, and* $\epsilon_{i\ell}(\boldsymbol{x})$ *is random simulation error which satisfies:*

(i) $\epsilon_{i\ell}(\boldsymbol{x}) \sim \mathcal{N}(0, \sigma_i^2(\boldsymbol{x}))$;

(ii) $\epsilon_{i\ell}(\boldsymbol{x})$ *is independent of* $\epsilon_{i'\ell'}(\boldsymbol{x}')$ *for any* $(i, \ell, \boldsymbol{x}) \neq (i', \ell', \boldsymbol{x}')$.

Fixed Design

- Under linear model, to obtain good $\widehat{i^*}(\boldsymbol{x})$, it suffices to estimate $\boldsymbol{\beta}_i$ to certain accuracy.
- $\boldsymbol{\beta}_i$ is estimated from running simulation.

Fixed Design

- Under linear model, to obtain good $\widehat{i^*}(\boldsymbol{x})$, it suffices to estimate $\boldsymbol{\beta}_i$ to certain accuracy.
- $\boldsymbol{\beta}_i$ is estimated from running simulation.
- So our problem becomes several linear regression problems, where the design matrix can be freely chosen by ourselves.

Fixed Design

- Under linear model, to obtain good $\widehat{i^*}(\boldsymbol{x})$, it suffices to estimate $\boldsymbol{\beta}_i$ to certain accuracy.
- $\boldsymbol{\beta}_i$ is estimated from running simulation.
- So our problem becomes several linear regression problems, where the design matrix can be freely chosen by ourselves.
- Choose $m \geq d+1$ design points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m \in \Theta$.
- We only run simulation at these chosen points.

Fixed Design

- Under linear model, to obtain good $\widehat{i^*}(\boldsymbol{x})$, it suffices to estimate $\boldsymbol{\beta}_i$ to certain accuracy.
- $\boldsymbol{\beta}_i$ is estimated from running simulation.
- So our problem becomes several linear regression problems, where the design matrix can be freely chosen by ourselves.
- Choose $m \geq d+1$ design points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m \in \Theta$.
- We only run simulation at these chosen points.

Assumption 2 (A2)

$\mathcal{X}^\top \mathcal{X}$ is nonsingular, where $\mathcal{X} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m)^\top \in \mathbb{R}^{m \times (d+1)}$.

Indifference Zone (IZ)

- It can be arbitrarily difficult to distinguish two alternatives.

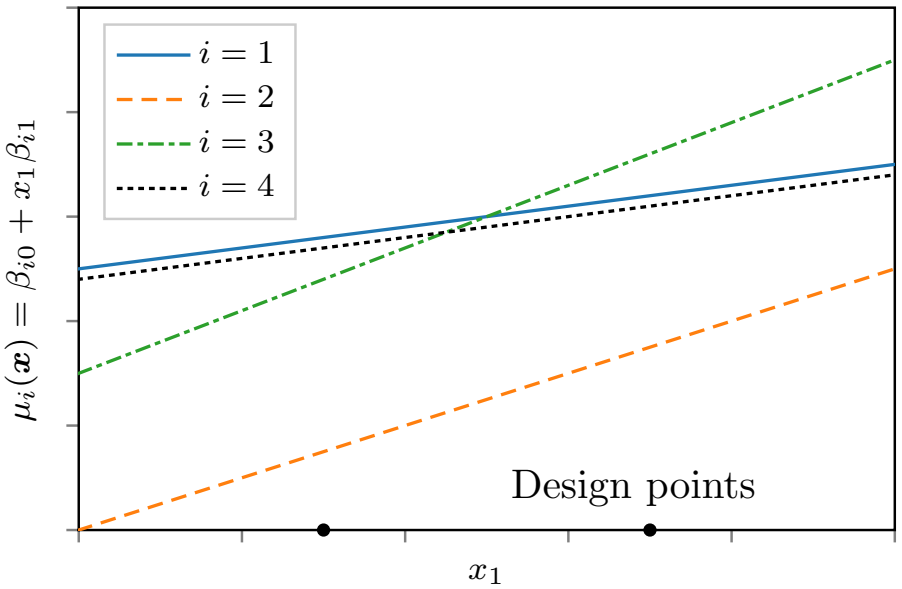

$k = 4$ alternatives,
$d = 1$.

Indifference Zone (IZ)

- It can be arbitrarily difficult to distinguish two alternatives.

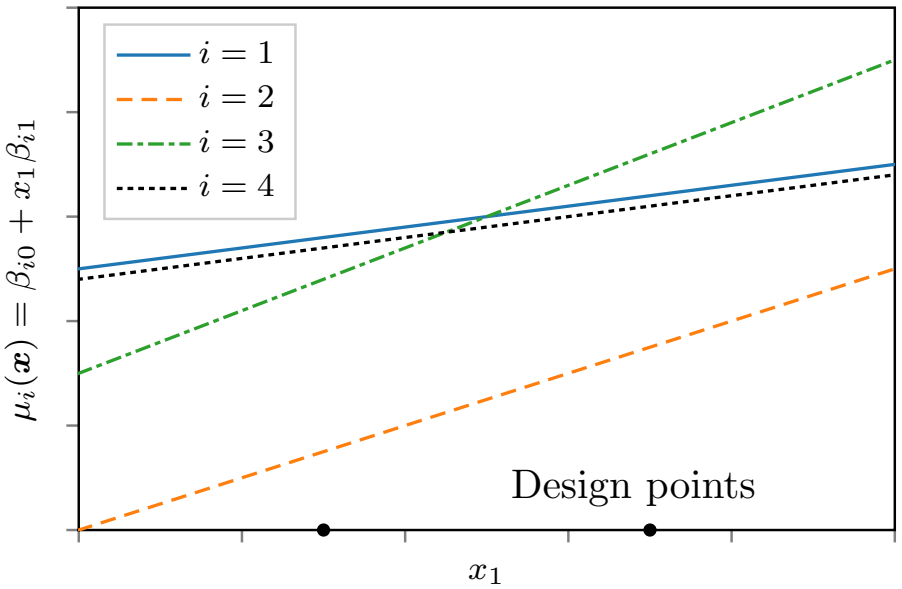

$k = 4$ alternatives,
$d = 1$.

- IZ parameter δ, represents the smallest difference between alternatives that a decision maker feels is worth detecting.

Indifference Zone (IZ)

- It can be arbitrarily difficult to distinguish two alternatives.

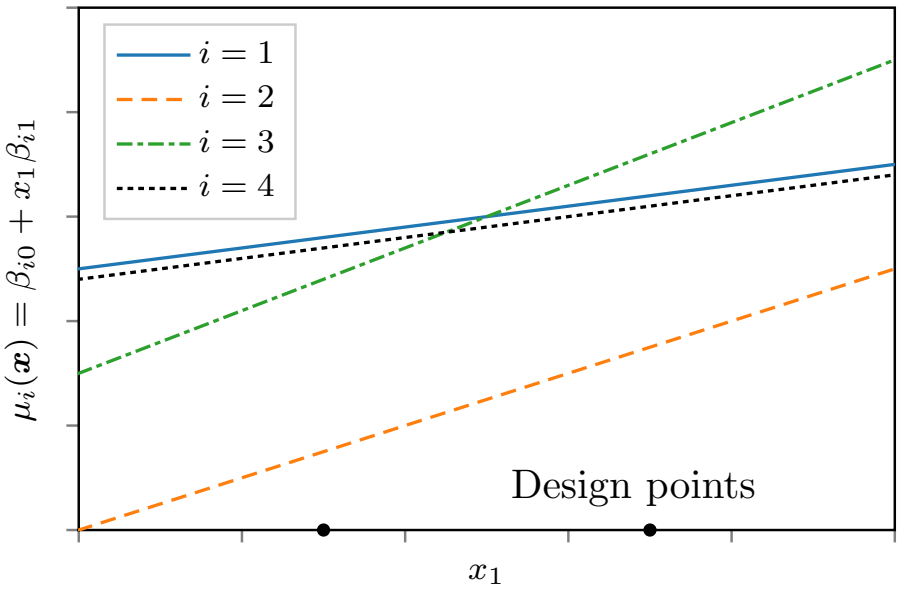

$k = 4$ alternatives,
$d = 1$.

- IZ parameter δ, represents the smallest difference between alternatives that a decision maker feels is worth detecting.
- Define a good selection (GS) conditionally on $\boldsymbol{X} = \boldsymbol{x}$ as

$$\mathrm{GS}(\boldsymbol{x}) := \left\{ \mu_{i^*(\boldsymbol{X})}(\boldsymbol{X}) - \mu_{\widehat{i^*}(\boldsymbol{X})}(\boldsymbol{X}) < \delta \middle| \boldsymbol{X} = \boldsymbol{x} \right\}.$$

Probability of Good Selection (PGS)

- Even with the IZ parameter, there is no guarantee that $\widehat{i^*}(\boldsymbol{x})$ always gives a good selection, conditionally on $\boldsymbol{X} = \boldsymbol{x}$.

- We define the *conditional* PGS as

$$\text{PGS}(\boldsymbol{x}) := \mathbb{P}\left\{ \mu_{i^*(\boldsymbol{X})}(\boldsymbol{X}) - \mu_{\widehat{i^*}(\boldsymbol{X})}(\boldsymbol{X}) < \delta \middle| \boldsymbol{X} = \boldsymbol{x} \right\},$$

  where the probability is with respect to the distribution of the samples used to produces $\widehat{i^*}(\boldsymbol{x})$.

Probability of Good Selection (PGS)

- Even with the IZ parameter, there is no guarantee that $\widehat{i^*}(\boldsymbol{x})$ always gives a good selection, conditionally on $\boldsymbol{X} = \boldsymbol{x}$.
- We define the *conditional* PGS as

$$\mathsf{PGS}(\boldsymbol{x}) := \mathbb{P}\left\{ \mu_{i^*(\boldsymbol{X})}(\boldsymbol{X}) - \mu_{\widehat{i^*}(\boldsymbol{X})}(\boldsymbol{X}) < \delta \middle| \boldsymbol{X} = \boldsymbol{x} \right\},$$

  where the probability is with respect to the distribution of the samples used to produces $\widehat{i^*}(\boldsymbol{x})$.

- Forms of *unconditional* PGS:
  - Distribution of $\boldsymbol{X}$ is known: $\mathsf{PGS}_{\mathsf{E}} := \mathbb{E}\left[\mathsf{PGS}(\boldsymbol{X})\right]$.
  - Distribution of $\boldsymbol{X}$ is unknown: $\mathsf{PGS}_{\mathsf{min}} := \min_{\boldsymbol{x}\in\Theta} \mathsf{PGS}(\boldsymbol{x})$.
  - Other forms may also be possible.

Key Task

- Design some selection procedures to produce $\widehat{i^*}(\boldsymbol{x})$, which can guarantee a particular unconditional PGS is $\geq 1-\alpha$.
- A selection procedure describes how to sample each alternative (through simulation) at the design points.

Key Task

- Design some selection procedures to produce $\widehat{i^*}(\boldsymbol{x})$, which can guarantee a particular unconditional PGS is $\geq 1-\alpha$.
- A selection procedure describes how to sample each alternative (through simulation) at the design points.
- Re-state the OLOA steps of R&S-C under A1 and A2:
  1. Carry out a selection procedure offline $\Rightarrow \widehat{\boldsymbol{\beta}}_i \Rightarrow \widehat{i^*}(\boldsymbol{x})$.
  2. Apply $\widehat{i^*}(\boldsymbol{x})$ online to guide the actual selection in *real time*.

Key Task

- Design some selection procedures to produce $\widehat{i^*}(\boldsymbol{x})$, which can guarantee a particular unconditional PGS is $\geq 1-\alpha$.
- A selection procedure describes how to sample each alternative (through simulation) at the design points.
- Re-state the OLOA steps of R&S-C under A1 and A2:
  1. Carry out a selection procedure offline $\Rightarrow \widehat{\boldsymbol{\beta}}_i \Rightarrow \widehat{i^*}(\boldsymbol{x})$.
     - A valid selection procedure should give a good enough $\widehat{i^*}(\boldsymbol{x})$.
  2. Apply $\widehat{i^*}(\boldsymbol{x})$ online to guide the actual selection in *real time*.

Key Task

- Design some selection procedures to produce $\widehat{i^*}(\boldsymbol{x})$, which can guarantee a particular unconditional PGS is $\geq 1-\alpha$.
- A selection procedure describes how to sample each alternative (through simulation) at the design points.
- Re-state the OLOA steps of R&S-C under A1 and A2:
  1. Carry out a selection procedure offline $\Rightarrow \widehat{\boldsymbol{\beta}}_i \Rightarrow \widehat{i^*}(\boldsymbol{x})$.
     - A valid selection procedure should give a good enough $\widehat{i^*}(\boldsymbol{x})$.
  2. Apply $\widehat{i^*}(\boldsymbol{x})$ online to guide the actual selection in *real time*.
     - $\widehat{i^*}(\boldsymbol{x})$ will be applied repeatedly for the population;
     - Overall measure of goodness of $\widehat{i^*}(\boldsymbol{x})$: PGS_E, PGS_min, etc;

Key Task

- Design some selection procedures to produce $\widehat{i^*}(\boldsymbol{x})$, which can guarantee a particular unconditional PGS is $\geq 1-\alpha$.
- A selection procedure describes how to sample each alternative (through simulation) at the design points.
- Re-state the OLOA steps of R&S-C under A1 and A2:
  1. Carry out a selection procedure offline $\Rightarrow \widehat{\boldsymbol{\beta}}_i \Rightarrow \widehat{i^*}(\boldsymbol{x})$.
     - A valid selection procedure should give a good enough $\widehat{i^*}(\boldsymbol{x})$.
  2. Apply $\widehat{i^*}(\boldsymbol{x})$ online to guide the actual selection in *real time*.
     - $\widehat{i^*}(\boldsymbol{x})$ will be applied repeatedly for the population;
     - Overall measure of goodness of $\widehat{i^*}(\boldsymbol{x})$: PGS_E, PGS_min, etc;
     - Good enough $\widehat{i^*}(\boldsymbol{x})$ means, e.g., $\text{PGS}_\text{E} \geq 1-\alpha$ or $\text{PGS}_\text{min} \geq 1-\alpha$.

Homoscedastic or Heteroscedastic Errors

- Selection procedures are designed separately depending on whether the simulation errors are *homoscedastic* (A3) or *heteroscedastic* (A4).

Assumption 3 (A3)

$\sigma_i^2(\boldsymbol{x}) \equiv \sigma_i^2 < \infty$ *for* $\boldsymbol{x} \in \Theta$ *and* $i = 1, \ldots, k$.

Assumption 4 (A4)

$\sigma_i^2(\boldsymbol{x}) < \infty$ *is a function of* $\boldsymbol{x} \in \Theta$ *and* $i = 1, \ldots, k$.

Homoscedastic or Heteroscedastic Errors

- Selection procedures are designed separately depending on whether the simulation errors are *homoscedastic* (A3) or *heteroscedastic* (A4).

Assumption 3 (A3)

$\sigma_i^2(\boldsymbol{x}) \equiv \sigma_i^2 < \infty$ *for* $\boldsymbol{x} \in \Theta$ *and* $i = 1, \ldots, k$.

Assumption 4 (A4)

$\sigma_i^2(\boldsymbol{x}) < \infty$ *is a function of* $\boldsymbol{x} \in \Theta$ *and* $i = 1, \ldots, k$.

- This analogizes the difference between the *ordinary least squares* method and the *generalized least squares* method in linear regression.

Homoscedastic Errors - Procedure FDHom

Setup. Specify δ, α, m and $\mathcal{X}$. Determine n_0, the first-stage sample size. Calculate the critical constant h (as shown in next slide).

Stage 1. For all $i = 1, \ldots, k$, take n_0 batches of observations on $\mathcal{X}$: $\boldsymbol{Y}_{i\ell} = (Y_{i\ell}(\boldsymbol{x}_1), \ldots, Y_{i\ell}(\boldsymbol{x}_m))^\top, \ell = 1, \ldots, n_0$. Let

$$\widehat{\boldsymbol{\beta}}_i(n_0) = \frac{1}{n_0}(\mathcal{X}^\top\mathcal{X})^{-1}\mathcal{X}^\top \sum_{\ell=1}^{n_0} \boldsymbol{Y}_{i\ell},$$

$$S_i^2 = \frac{1}{n_0 m - 1 - d} \sum_{\ell=1}^{n_0} (\boldsymbol{Y}_{i\ell} - \mathcal{X}\widehat{\boldsymbol{\beta}}_i(n_0))^\top (\boldsymbol{Y}_{i\ell} - \mathcal{X}\widehat{\boldsymbol{\beta}}_i(n_0)).$$

Furthermore, let $N_i = \max\left\{\left\lceil \frac{h^2 S_i^2}{\delta^2} \right\rceil, n_0\right\}$.

Stage 2. For all $i = 1, \ldots, k$, take $N_i - n_0$ batches of observations on $\mathcal{X}$ and denote them as $\boldsymbol{Y}_{i,n_0+1}, \ldots, \boldsymbol{Y}_{iN_i}$. Let $\widehat{\boldsymbol{\beta}}_i = \frac{1}{N_i}(\mathcal{X}^\top\mathcal{X})^{-1}\mathcal{X}^\top \sum_{\ell=1}^{N_i} \boldsymbol{Y}_{i\ell}$.

Selection. Return $\widehat{i^*}(\boldsymbol{x}) = \arg\max_{1\le i\le k}\{\boldsymbol{x}^\top \widehat{\boldsymbol{\beta}}_i\}$ as the decision rule.

Calculation of h

- When $\text{PGS}_{\text{E}} \geq 1-\alpha$ is designed, $h = h_{\text{E}}$, which satisfies

$$\mathbb{E}\left\{\int_0^\infty \left[\int_0^\infty \Phi\left(\frac{h_{\text{E}}}{\sqrt{(n_0 m-1-d)\left(\frac{1}{t}+\frac{1}{s}\right)\boldsymbol{X}^\top(\mathcal{X}^\top\mathcal{X})^{-1}\boldsymbol{X}}}\right) f(s)ds\right]^{k-1} f(t)dt\right\} = 1-\alpha,$$

  where $\Phi(\cdot)$ is the standard normal cdf, $f(\cdot)$ is pdf of chi-squared RV with $n_0 m - 1 - d$ degrees of freedom.

- When $\text{PGS}_{\text{min}} \geq 1-\alpha$ is designed, $h = h_{\text{min}}$, which satisfies

$$\min_{\boldsymbol{x}\in\Theta}\left\{\int_0^\infty \left[\int_0^\infty \Phi\left(\frac{h_{\text{min}}}{\sqrt{(n_0 m-1-d)\left(\frac{1}{t}+\frac{1}{s}\right)\boldsymbol{x}^\top(\mathcal{X}^\top\mathcal{X})^{-1}\boldsymbol{x}}}\right) f(s)ds\right]^{k-1} f(t)dt\right\} = 1-\alpha.$$

Statistical Validity

Theorem 1

Under A1 – A3, Procedure FDHom ensures that the target unconditional PGS is at least $1-\alpha$*, i.e.,* $PGS_E \geq 1-\alpha$ *or* $PGS_{min} \geq 1-\alpha$.

Statistical Validity

Theorem 1

Under A1 – A3, Procedure FDHom ensures that the target unconditional PGS is at least $1-\alpha$*, i.e.,* $PGS_E \geq 1-\alpha$ *or* $PGS_{min} \geq 1-\alpha$.

- We also designed a procedure for the heteroscedastic simulation error case (A4), which is called Procedure FDHet.
- The statistical validity of Procedure FDHet was also proved.

Settings

- Generic setting:
  - Covariates $X_1, \ldots, X_d$, are i.i.d. Unif $(0, 1)$ RVs.
  - Take $m = 2^d$ design points: $\{0, 0.5\} \times \cdots \times \{0, 0.5\}$.
- Benchmark problem (0):
  - $d = 3$ and $k = 5$.
  - Mean configuration: GSC, $\beta_{10} - \delta = \beta_{i0} = 0,\ \beta_{1j} = \beta_{ij} = 1$.
  - Homoscedastic errors: $\sigma_i^2(\boldsymbol{x}) \equiv \sigma_i^2$.
  - Equal variances among alternatives: $\sigma_1 = \cdots = \sigma_k = 10$.
- 8 comparing problems:
  (1) Set $k = 2$.
  (2) Set $k = 8$.
  (3) Mean configuration: Non-GSC, randomly generate all components of $\boldsymbol{\beta}_i$ from Unif $(0, 5)$, for $i = 1, \ldots, 5$.
  (4) Increasing variances among alternatives: $\sigma_1 = 5$, $\sigma_2 = 7.5$, $\sigma_3 = 10$, $\sigma_4 = 12.5$, $\sigma_5 = 15$.
  (5) Decreasing variances among alternatives.
  (6) Heteroscedastic errors: $\sigma_i(\boldsymbol{x}) = 10\boldsymbol{x}^\top \boldsymbol{\beta}_i$, for $i = 1, \ldots, 5$.
  (7) Set $d = 1$.
  (8) Set $d = 5$.

Numerical Results

- PGS_E is designed to be $\geq 95\%$ (i.e., $\alpha = 0.05$), $\delta = 1, n_0 = 50$.
- 10^4 macro replications are carried out for each procedure.
- 10^5 samples of $\boldsymbol{X}$ are used to calculate APGS_E (and APGS_min) of $\widehat{i^*}(\boldsymbol{X})$ produced by each procedure.

	Procedure FDHom				Procedure FDHet			
Problem	h_E	Sample	APGS_E	APGS_min	h_E	Sample	APGS_E	APGS_min
(0) Benchmark	3.423	46865	0.9610	0.7439	4.034	65138	0.9801	0.8080
(1) $k = 2$	2.363	8947	0.9501	0.8084	2.781	12380	0.9702	0.8517
(2) $k = 8$	3.822	93542	0.9650	0.7246	4.510	130200	0.9842	0.8052
(3) Non-GSC	3.423	46865	0.9987	0.9410	4.034	65138	0.9994	0.9615
(4) IV	3.423	52698	0.9618	0.7549	4.034	73265	0.9807	0.8147
(5) DV	3.423	52720	0.9614	0.7501	4.034	73246	0.9806	0.8114
(6) Het	3.423	58626	0.9232	0.6336	4.034	81555	**0.9846**	0.8591
(7) $d = 1$	4.612	21288	0.9593	0.7941	4.924	24266	0.9662	0.8223
(8) $d = 5$	2.141	73428	0.9656	0.7446	2.710	117630	0.9895	0.8379

Numerical Results

- PGS_{min} is designed to be $\geq 95\%$ (i.e., $\alpha = 0.05$), $\delta = 1, n_0 = 50$.
- 10^4 macro replications are carried out for each procedure.
- 10^5 samples of $\boldsymbol{X}$ are used to calculate APGS_{min} (and APGS_{E}) of $\widehat{i^*}(\boldsymbol{X})$ produced by each procedure.

	Procedure FDHom				Procedure FDHet			
Problem	h_{min}	Sample	APGS_{E}	APGS_{min}	h_{min}	Sample	APGS_{E}	APGS_{min}
(0) Benchmark	5.927	140540	0.9989	0.9594	6.990	195340	0.9997	0.9825
(1) $k = 2$	4.362	30447	0.9958	0.9466	5.132	42164	0.9987	0.9701
(2) $k = 8$	6.481	268750	0.9993	0.9642	7.651	374720	0.9999	0.9849
(3) Non-GSC	5.927	140540	1.0000	0.9958	6.990	195340	1.0000	0.9981
(4) IV	5.927	158140	0.9989	0.9574	6.990	219870	0.9998	0.9862
(5) DV	5.927	158100	0.9990	0.9617	6.990	219740	0.9998	0.9826
(6) Het	5.927	175700	0.9952	0.8999	6.990	244490	0.9999	**0.9899**
(7) $d = 1$	7.155	51161	0.9954	0.9600	7.648	58493	0.9971	0.9708
(8) $d = 5$	3.792	230220	0.9994	0.9539	4.804	369310	1.0000	0.9907

Background

- Esophageal cancer is the fourth (seventh) leading cause of cancer death among males in China (U.S.).

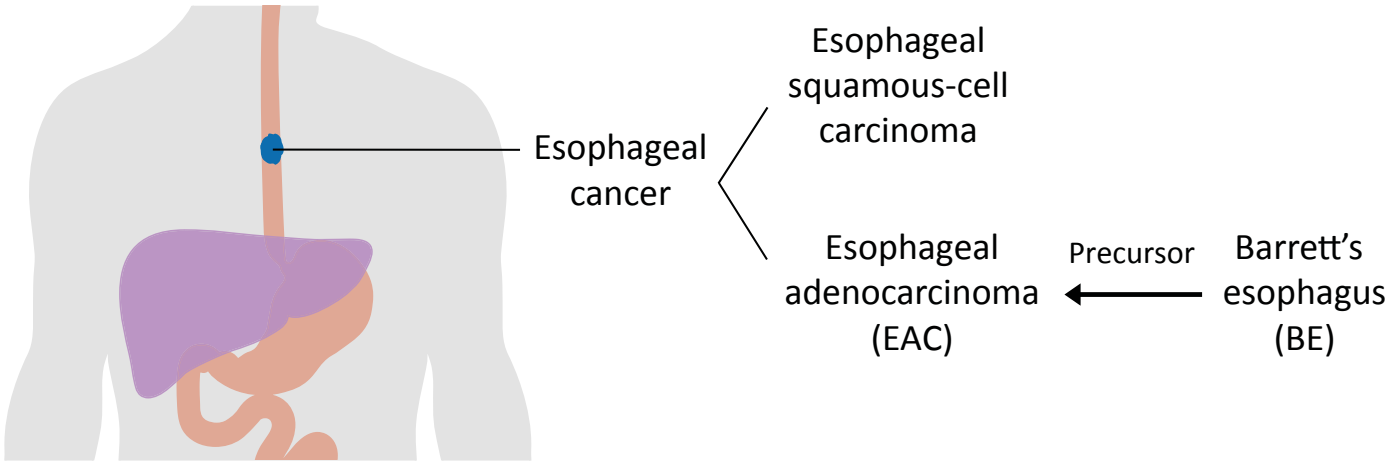

Image Source: Cancer Research UK / Wikimedia Commons.

Background

- Esophageal cancer is the fourth (seventh) leading cause of cancer death among males in China (U.S.).

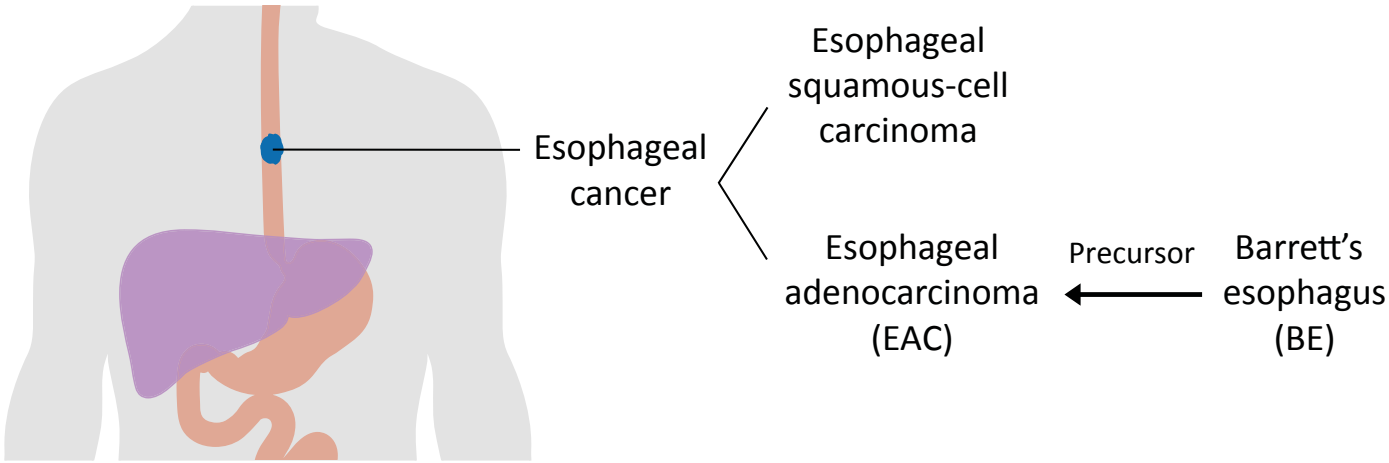

Image Source: Cancer Research UK / Wikimedia Commons.

- EAC is one sub-type of esophageal cancer, and its incidence has increased by 500% over the past 40 years (Bollschweiler et al. 2001, Hur et al. 2013).

Background

- Esophageal cancer is the fourth (seventh) leading cause of cancer death among males in China (U.S.).

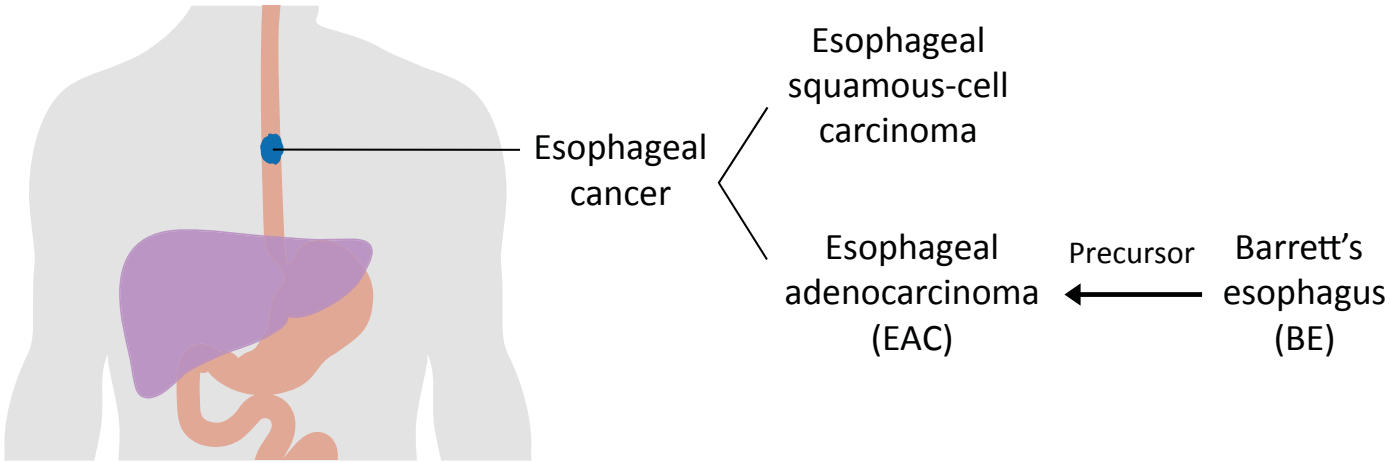

Image Source: Cancer Research UK / Wikimedia Commons.

- EAC is one sub-type of esophageal cancer, and its incidence has increased by 500% over the past 40 years (Bollschweiler et al. 2001, Hur et al. 2013).
- BE is a precursor to EAC, and its management is important and attracts many attentions.

Best Treatment Regimen

- Consider 3 treatment regimens (i.e., alternatives) for BE; all regimens include standard endoscopic surveillance:
  (1) No drug;
  (2) Aspirin chemoprevention;
  (3) Statin chemoprevention.

- Consider some individual characteristics (i.e., covariates $\boldsymbol{X}$):
  X_1 – Age;
  X_2 – Risk (i.e., the annual progression rate of BE to EAC);
  X_3 – Effect of aspirin (i.e., progression reduction effect);
  X_4 – Effect of statin.

- Consider expected quality-adjusted life years (QALYs) after getting BE, as mean performance of each alternative.

- The best treatment regimen for BE is patient-specific (i.e., depending on individual covariates).

Applying R&S-C

- Use the simulation model developed and calibrated by Hur et al. (2004) and Choi et al. (2014).
- Assume the distribution of $\boldsymbol{X}$ is known and specified as follows

Covariates	Distributions	Support	Mean
X_1	Discrete	$\{55, \ldots, 80\}$	64.78
X_2	Unif $(0, 0.1)$	$[0, 0.1]$	0.05
X_3	Triangular $(0, 0.59, 1)$	$[0, 1]$	0.53
X_4	Triangular $(0, 0.62, 1)$	$[0, 1]$	0.54

Personalized Medicine via R&S-C

- Two ways of choosing treatment regimen:
  - (i) Traditional way: the best treatment is the one that works best for the average of population ($i^\dagger \equiv 3$).
  - (ii) Personalized way: the best treatment is $i^*(\boldsymbol{x})$.

    Get the estimate, $\widehat{i^*}(\boldsymbol{x})$, from conducting Procedure FDHet with target $\mathsf{PGS}_\mathsf{E} \geq 95\%, \delta = 0.2, n_0 = 100$.

Personalized Medicine via R&S-C

- Two ways of choosing treatment regimen:
  - (i) Traditional way: the best treatment is the one that works best for the average of population ($i^\dagger \equiv 3$).
  - (ii) Personalized way: the best treatment is $i^*(\boldsymbol{x})$.

    Get the estimate, $\widehat{i^*}(\boldsymbol{x})$, from conducting Procedure FDHet with target $\mathsf{PGS}_\mathsf{E} \geq 95\%, \delta = 0.2, n_0 = 100$.

- The actual achieved PGS_E (APGS_E):
  - (i) Traditional way: $\mathsf{APGS}_\mathsf{E} \approx 78.0\%$.
  - (ii) Personalized way: $\mathsf{APGS}_\mathsf{E} \approx 99.7\%$.

Personalized Medicine via R&S-C

- Distributions of expected QALYs under the two ways.
  - (a) For the *entire population* of patients.
  - (b) For a more *specific group* of patients, i.e., patients with $\boldsymbol{X} = (1, X_1, X_2, 0.9, 0.2)^\top$.

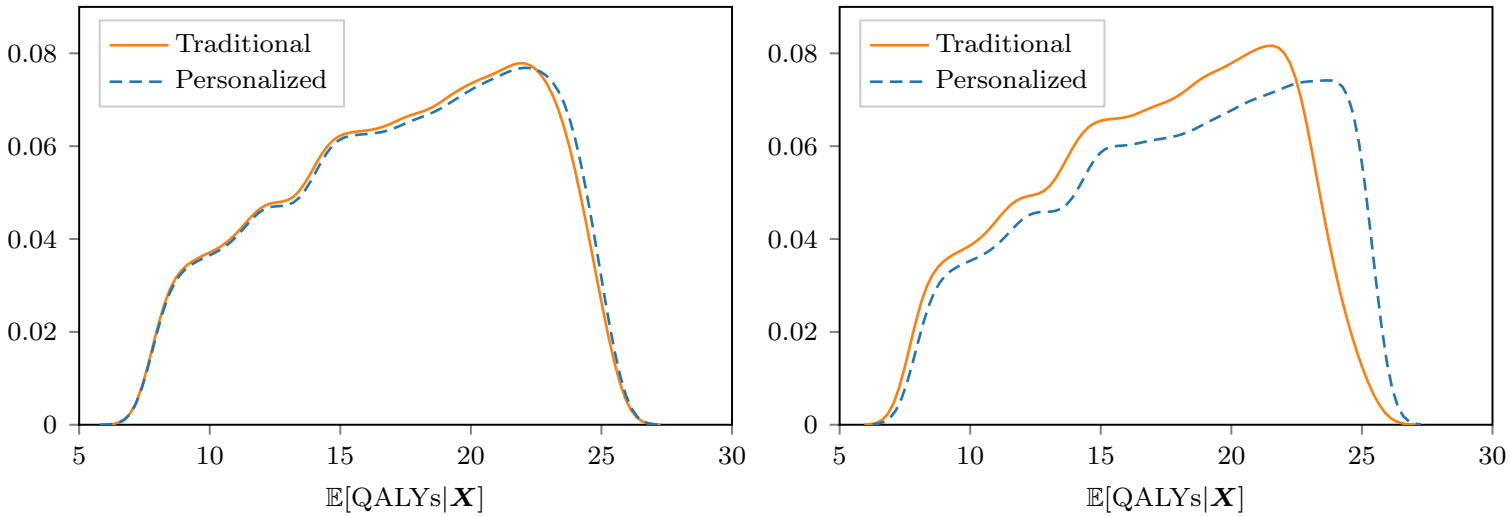

(a) (b)

  - (c) For a *specific individual* with $\boldsymbol{X} = (1, 55, 0.1, 0.9, 0.2)^\top$, expected QALYs increases by 2.43 years when personalized medicine is performed.

Conclusions

- Personalized decisions lead us to consider Ranking and Selection with Covariates.
- We use a linear model—the simplest yet most useful parametric model in practice—to capture the relationship between the response and covariates.
- There are many directions that R&S-C may be studied towards, e.g., non-parametric models, Bayesian formulations, random designs and sequential procedures.

Thank You!

Haihui SHEN
Jan. 6, 2018