Treatment choice with many covariate values

Aleksey Tetenov (University of Bristol)

Cemmap masterclass
Statistical decision theory for treatment choice and prediction
May 30-31, 2017

Stoye (2009), Proposition 4: If covariate X is continuously distributed, minimax regret is constant and does not decrease with sample size.

This result changes if

1. Stronger assumptions on how average treatment response $E[Y_t|X]$ varies with X (Stoye, 2012)

2. The set of feasible treatment rules is restricted

Source: Kitagawa and Tetenov (2017), "Who Should be Treated? Empirical Welfare Maximization Methods for Treatment Choice" Cemmap working paper CWP24/17

Regret is evaluated relative to the best implementable treatment rule.

Stoye (2009, Proposition 4) assumes that any treatment allocation is feasible, including arbitrarily complex treatment rules.

This is an unreasonable benchmark for public policies. Constraints frequently restrict the complexity and other characteristics of feasible treatment rules.

- Treatment rules are often publicly communicated to individuals and need to be understandable and transparent
- Monotonicity of treatment rules in some covariates if desirable (e.g., cannot treat the rich but not the poor)
- Some treatments may be capacity-constrained
- Other aggregate constraints (e.g., aggregate proportion treated cannot vary with race)

Setup

A Randomized Controlled Trial (RCT) sample

- $X_i \in \mathcal{X}$ - pre-treatment observed covariates
- $D_i \in \{0, 1\}$ - randomized treatment
- $Y_i \in \mathbb{R}$ - treatment outcome
- $Y_{0,i}, Y_{1,i}$ - potential outcomes
- $e(x) \in [\kappa, 1 - \kappa]$ - the probability of being randomized to treatment 1 in the experiment

We consider a restricted set of treatment rules $\mathcal{G}$. Each G specifies which subset of the population will be treated (after analyzing the experimental data)

- $X \in G$ will be assigned to treatment 1
- $X \notin G$ will be assigned to treatment 0

(excludes randomized/fractional treatment rules)

$\hat{G} \in \mathcal{G}$ treatment rule as a function of the sample

Empirical Welfare Maximization

- Estimate the policy directly by maximizing empirical welfare

$$\hat{G}_{EWM} = arg \max_{G \in \mathcal{G}} W_n(G),$$

- Sample analogue

$$W_n(G) \equiv \frac{1}{n} \sum_{i=1}^{n} \left[\frac{Y_i D_i}{e(X_i)} \cdot 1\left\{X_i \in G\right\} + \frac{Y_i(1 - D_i)}{1 - e(X_i)} \cdot 1\left\{X_i \notin G\right\} \right]$$

consistently estimates the population welfare of policy G,

$$W(G) = E\left[Y_1 \cdot 1\left\{X \in G\right\} + Y_0 \cdot 1\left\{X_i \notin G\right\}\right].$$

- EWM treatment rule: $\widehat{G}_{EWM} \equiv \arg\max_{G \in \mathcal{G}} W_n(G)$

Empirical Illustration

- National Job Training Partnership Act (JTPA) Study (Bloom et al, 1997)
- Sample: 11,204 adult applicants
- Propensity score = 2/3 (probability of treatment)
- Outcome $Y = D(Y_1 - cost) + (1 - D)Y_0$:
  - Total individual earnings in the 30-month period following treatment assignment
  - Total earnings minus $774 (average cost of each treatment assignment, taking into account variable take-up)
- Covariates X: Years of education, pre-program earnings

- Average treatment effect: $1,157
- 95% CI: ($513, $1,801)

Parametric plug-in treatment rule:
estimate $E(Y_1|X)$ and $E(Y_0|X)$ by OLS.
Assign treatment 1 if $X'\beta_1 > X'\beta_0$

No cost: treat everyone, est. gain $1,157
With $774 cost: treat 96%, est. gain $466 (per population member)

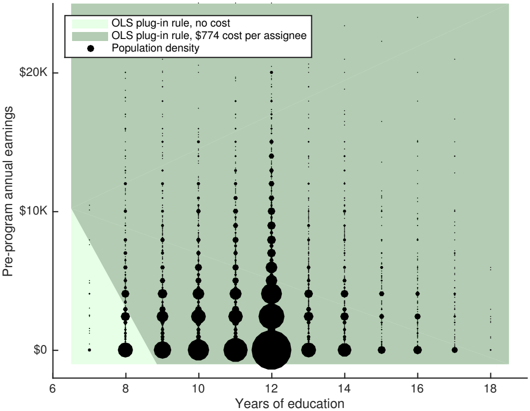

EWM linear rule: maximizes the sample analog of welfare among linear decision rules $\hat{G} = 1\{X'\beta \geq 0\}$

No cost: treat 90%, est. gain $1,408. 95% CI: ($592, $2,225)
$774 cost: treat 90%, est. gain $712. 95% CI: (-$107, $1,532)

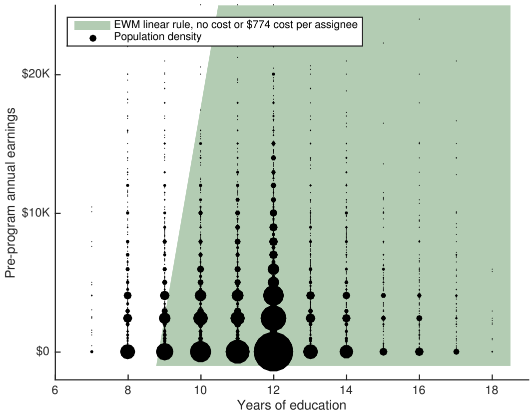

EWM quadrant rule: select best min or max threshold for each covariate $\hat{G} = 1\{x_1 > (<)t_1, x_2 > (<)t_2\}$

No cost: treat 93%, est. gain \$1,277. 95% CI: (\$519, \$2,034)
\$774 cost: treat 83%, est. gain \$687. 95% CI: (-\$71, \$1,445)

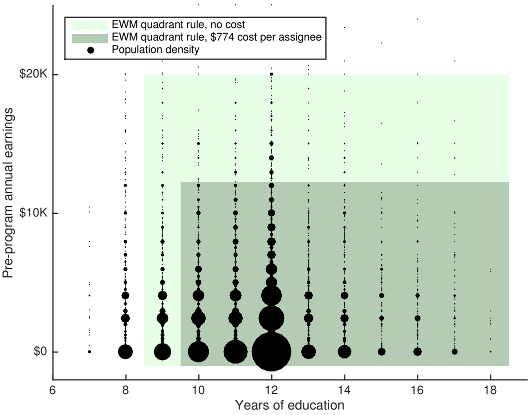

Non-parametric plug-in rule: bivariate kernel reg of $Y_1|X$ and $Y_0|X$ (ROT bandwidth).

No cost: treat 82%, est. gain \$1,867
\$774 cost: treat 69%, est. gain \$1,257

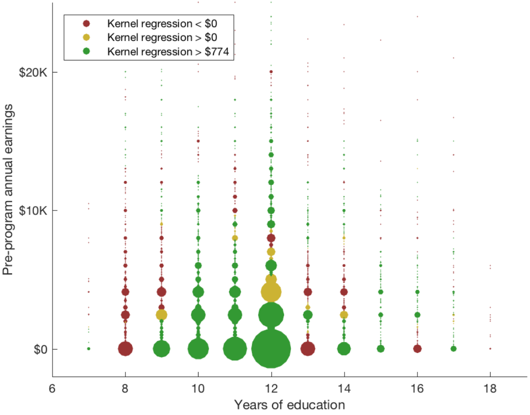

Welfare Criterion

Object of interest: policy with the highest utilitarian (additive) welfare

Outcome variable Y should reflect social preferences, so it may need to

- give different weight to different individuals
- non-linearly transform outcomes
- aggregate multi-dimensional outcomes
- subtract treatment costs from outcomes

- The utilitarian welfare of treatment rule G is

$$\begin{aligned} W(G) &\equiv E\left[Y_1 \cdot 1\left\{X \in G\right\} + Y_0 \cdot 1\left\{X \notin G\right\}\right] \\ &= E\left[Y_0\right] + E\left[\tau(X) 1\left\{X \in G\right\}\right], \end{aligned}$$

$\tau(X) \equiv E\left(Y_1 - Y_0 | X\right)$: the conditional treatment effect

- We can equivalently work with the welfare gain of treating subset G relative to treating no one

$$\begin{aligned} V(G) &\equiv W(G) - W(\emptyset) \\ &= E\left[\tau(X) \cdot 1\left\{X \in G\right\}\right], \end{aligned}$$

First Best treatment rule (with no constraints on G)

$$\begin{aligned} G^*_{FB} &\equiv \{x : \tau(x) \geq 0)\} \\ &\in \arg \max_{G \in \mathcal{B}(\mathcal{X})} W(G) \\ &\in \arg \max_{G \in \mathcal{B}(\mathcal{X})} V(G) \end{aligned}$$

Second Best treatment rule
maximizing welfare in a constrained class $\mathcal{G}$

$$\begin{aligned} G^* &\in \arg\max_{G \in \mathcal{G}} W(G) \\ &\in \arg\max_{G \in \mathcal{G}} V(G) \end{aligned}$$

The maximized feasible welfare

$$W^*_{\mathcal{G}} \equiv \sup_{G \in \mathcal{G}} W(G) \leq W(G^*_{FB})$$

Assumptions:

Distribution of (Y_0, Y_1, D, Y) is $P \in \mathcal{P}$.

The only assumption on the distribution of treatment response:

- **Bounded Outcomes**: $Y_1,\ Y_0 \in \left[-\frac{M}{2}, \frac{M}{2}\right]$, $M < \infty$, implying $|\tau(x)| \leq M$, $\forall x$.

Restriction on experimental design (point-identifies $\tau(x)$)

- **Strict Overlap**: There exist $\kappa > 0$, s.t. $e(x) \in [\kappa, 1-\kappa]$, $\forall x$.

Restriction on $\mathcal{G}$:

- **Complexity of Decision Sets**: $\mathcal{G}$ is a countable VC-class of subsets with finite VC-dimension: $v =$ the maximal number of points in $\mathcal{X}$ that can be shattered by $\mathcal{G}$.

Examples of VC-classes $\mathcal{G}$

Linear eligibility score:

$$\mathcal{G} = \left\{ \{x : x'\beta \geq 0\} : \beta \in \mathbb{R}^{d_x} \right\}$$

has $v = d_x + 1$.

Generalized eligibility score:

$$\mathcal{G} = \left\{ \left\{ x : \sum_{k=1}^{m} a_k f_k(x) \geq g(x) \right\} : (a_1, \ldots, a_m) \in \mathbb{R}^m \right\}$$

has $v \leq m + 1$.

Multiple index rules:

$$\mathcal{G} = \{\{x : (f_1(x_1) \leq c_1) \cap \cdots (f_K(x_K) \leq c_K)\} : (c_1, \ldots, c_m) \in \mathbb{R}^m\}$$

has $v \leq K + 1$.

Upper bound on maximum regret of EWM

Theorem 2.1: *Let $\mathcal{P}$ be a class of DGPs satisfying assumptions Bounded Outcomes and Strict Overlap. Let $\mathcal{G}$ be a VC-class of treatment choice rules. Then*

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[W_{\mathcal{G}}^* - W(\hat{G}_{EWM}) \right] \leq C_1 \frac{M}{\kappa} \sqrt{\frac{v}{n}},$$

where C_1 is a universal constant.

Remarks on rate bounds:

- This rate bound is valid whether $G_{FB}^* \in \mathcal{G}$ or not.
- Parametric plug-in with misspecified regressions does not have such second-best optimality.

Proof: sketch

For any $\tilde{G} \in \mathcal{G}$,

$$\begin{aligned} & W(\tilde{G}) - W(\hat{G}_{EWM}) \\ \leq \; & W_n(\hat{G}_{EWM}) - W_n(\tilde{G}) + W(\tilde{G}) - W(\hat{G}_{EWM}) \\ \leq \; & \left| W_n(\hat{G}_{EWM}) - W(\hat{G}_{EWM}) \right| + \left| W_n(\tilde{G}) - W(\tilde{G}) \right| \\ \leq \; & 2 \sup_{G \in \mathcal{G}} \left| W_n(G) - W(G) \right| . \end{aligned}$$

So,

$$W^*_{\mathcal{G}} - W(\hat{G}_{EWM}) \leq 2 \sup_{G \in \mathcal{G}} \left| W_n(G) - W(G) \right|$$

Proof: sketch

$W_n(G) = E_n(f(\cdot; G))$ and $W(G) = E(f(\cdot; G))$, where

$$f(\cdot; G) = \left[\frac{Y_i D_i}{e(X_i)} 1\{X_i \in G\} + \frac{Y_i(1 - D_i)}{1 - e(X_i)} 1\{X_i \notin G\} \right]$$

Lemma A.1

If $\mathcal{G}$ is a VC-class of sets with VC-dimension v and $g(\cdot), h(\cdot)$ are two given real-valued functions of observations, then functions

$$\{f(\cdot; G) = g(\cdot) \cdot 1\{x \in G\} + h(\cdot) \cdot 1\{x \notin G\}, \quad G \in \mathcal{G}\}$$

form a VC-subgraph class with VC-dimension $\leq v$.

Using this lemma, we can apply a well-known maximal inequality for centered empirical processes to

$$\sup_{G \in \mathcal{G}} |W_n(G) - W(G)| = \sup_{G \in \mathcal{G}} |E_n(f) - E(f)|$$

Lower bound on minimax regret

Theorem 2.2: *Let $\mathcal{P}$ be a class of DGPs satisfying Bounded Outcomes and Strict Overlap. Let $\mathcal{G}$ be a VC-class of treatment choice rules. Then, for any treatment choice rule $\hat{G}$*

$$\sup_{P \in \mathcal{P}} E_{P^n}\left[W_{\mathcal{G}}^{*} - W(\hat{G})\right] \geq \frac{M}{2} e^{-2\sqrt{2}} \sqrt{\frac{v}{n}} \quad \text{for all } n \geq 16v,$$

Remarks on rate bounds:

- Both are finite-sample bounds (but not sharp).
- $\hat{G}_{EWM}$ is minimax rate optimal: no $\hat{G}$ has maximum regret converging to zero at a faster rate uniformly over $\mathcal{P}$.
- EWM is minimax rate optimal even when v grows with n.

Proof: sketch

For the lower bound, we adapt the argument in Lugosi (2002):

$$
\begin{aligned}
& \sup_{P \in \mathcal{P}} E_{P^n}\left[W_{\mathcal{G}}^* - W(G_n)\right] \\
\geq \; & \sup_{P \in \mathcal{P}^*} E_{P^n}\left[W_{\mathcal{G}}^* - W(G_n)\right] \\
\geq \; & \int_{\mathcal{P}^*} E_{P^n}\left[W_{\mathcal{G}}^* - W(G_n)\right] d\mu(P) \\
\geq \; & \int_{\mathcal{P}^*} E_{P^n}\left[W_{\mathcal{G}}^* - W(\hat{G}_{bayes})\right] d\mu(P),
\end{aligned}
$$

where $\mathcal{P}^* \subset \mathcal{P}$ is a class of DGPs that has a discrete support of X with v points and $\tau(x) = \gamma$ or $-\gamma$. For uniform prior μ, the Bayes risk can be analytically computed as a function of γ. Setting $\gamma = \sqrt{v/n}$ gives the lower bound.

Discussion: EWM and Statistical Decision Theory

There are important open questions:

- Are EWM rules admissible?
- Finite-sample minimax regret: we know that EWM rules cannot be exactly minimax regret in some cases (when fractional/randomized treatment assignment for tie-breaking is required). Are they close to finite-sample minimax regret?
- Are there better treatment rules with the same uniform regret convergence rates?

Alternative approaches to treatment choice with covariates

Plug-in approach: uniformly estimates $\tau(x) = E(Y_1 - Y_0|X = x)$ and use treatment rule $1\{\widehat{\tau}(x) > 0\}$

- Requires assumptions on $\tau(x)$ that may not be credible
- May generate treatment rules that are not implementable

EWM approach: maximizes $\int_G \tau(x) dP_X(x)$ over a constrained set of $G \in \mathcal{G}$

- Minimal assumptions on $\tau(x)$ needed to uniformly estimate $\int_G$
- Easily incorporates constraints
- Computationally challenging

Surrogate loss functions (e.g., Support Vector Machines): maximizes $\int_G \tilde{\tau}(x) dP_X(x)$ for a more convenient $\tilde{\tau}(x) \neq \tau(x)$ s.t. $sign(\tilde{\tau}(x)) = sign(\tau(x))$

- Not well suited for constrained problems
- Computationally attractive

Computing EWM rules

EWM among policies linear in X (or its functions)

$$\begin{aligned}
\hat{G}_{EWM} &\equiv 1\left\{X'\hat{\beta} \geq 0\right\} \\
\hat{\beta} &\in \arg\max_{\beta \in \mathbf{B}} \sum_{i=1..n} g_i \cdot 1\left\{X_i'\beta \geq 0\right\}, \\
g_i &\equiv \frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}
\end{aligned}$$

Similar to the maximum score estimator.

We improve on the approach of Florios and Skouras (2008), who noticed that the problem could be substituted by an equivalent Mixed Integer Linear Programming problem.

Remark 2.1: capacity constraints

Capacity constraint: Proportion of the target population assigned to treatment 1 cannot exceed $K > 0$.

If the distribution of covariates P_X is known, restrict maximization to a subset of class $\mathcal{G}$ that satisfies the capacity constraint:

$$\mathcal{G}^K \equiv \{G \in \mathcal{G} : P_X(G) \leq K\}.$$

If P_X is unknown, we cannot guarantee that estimated policy $\hat{G}$ will satisfy the capacity constraint. To evaluate welfare, we need to specify what will happen in that case.

Remark 2.1: capacity constraints

- We assume that treatment 1 is "rationed" randomly among targeted recipients with $X \in \hat{G}$ and the resulting welfare is $W^K(G)$.
- Let $\hat{G}^K$ maximize the sample analog of the capacity-constrained welfare.

Theorem B.1: Under the same assumptions as previous theorems,

$$\sup_{P \in \mathcal{P}} E_{P^n}\left[\sup_{G \in \mathcal{G}} W^K(G) - W^K(\hat{G}^K)\right] \leq C_1 M(\kappa^{-1} + K^{-1})\sqrt{\frac{v}{n}},$$

where C_1 is a universal constant.

Remark 2.2: Target population has a different composition

EWM when target population $\neq$ sampled population.

- Suppose $E^T(Y_1 - Y_0|X) = E(Y_1 - Y_0|X) = \tau(X)$, but the distributions of X are different.
- If $G^*_{FB} \in \mathcal{G}$, G^*_{FB} is optimal for both populations.
- If $G^*_{FB} \notin \mathcal{G}$, a second best policy for the sampled population $\neq$ an optimal policy for the target population
- EWM with weighted empirical welfare: If $\rho(x) = \frac{dP^T_X/dx}{dP_X/dx}$ is known,

$$\hat{G}^T_{EWM} \equiv \arg\max_{G \in \mathcal{G}} E_n \left[\left(\frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)} \right) \rho(X) 1\{X \in G\} \right]$$

- If $\sup_x \rho(x) < \infty$, reweighting only affects the constant term of the welfare loss bounds.

Remark 2.3: Invariance

$$W_n(G) = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{Y_i D_i}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{Y_i(1-D_i)}{1-e(X_i)} \cdot 1\{X_i \notin G\}\right]$$

- If Y is multiplied by a constant, $W_n(G)$ is multiplied by the same constant (for all G)
- If Y is replaced by $Y + c$, $W_n(G)$ changes by

$$c \cdot \frac{1}{n}\sum_{i=1}^{n}\left[\frac{D_i}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{1-D_i}{1-e(X_i)} \cdot 1\{X_i \notin G\}\right] \neq c,$$

  which in finite samples varies with G.
- Linear transformations of Y could change the proposed finite-sample treatment rule and welfare gain estimates

Remark 2.3: Invariance

- We make a simple adjustment to obtain treatment rules that are invariant to linear transformations of Y by demeaning outcomes Y_i by their sample mean:

$$Y_i^{dm} \equiv Y_i - E_n[Y_i]$$

- and maximize

$$\arg\max_{G \in \mathcal{G}} E_n \left[\frac{Y_i^{dm} D_i}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{Y_i^{dm}(1 - D_i)}{1 - e(X_i)} \cdot 1\{X_i \notin G\} \right].$$

- This modification of the EWM treatment rule has the same $\sqrt{n}$ welfare convergence rate.
- In simulations, improved performance when $E[Y]$ is far from zero.
- We use demeaned outcomes in our application.

Faster convergence with a Margin Assumption

Does EWM remain rate optimal for a smaller class of DGPs?

Correct Specification of $\mathcal{G}$: $G^*_{FB} \in \mathcal{G}$.

Assumption MA: *Margin Assumption* (Mammen & Tsybakov (99, Ann.Stat)). There exists constants $0 < \eta \leq M$ and $0 < \alpha < \infty$ such that

$$P_X\left(|\tau(X)| \leq t\right) \leq \left(\frac{t}{\eta}\right)^{\alpha} \quad \forall 0 \leq t \leq \eta.$$

Denote the class of DGPs satisfying these assumptions by $\mathcal{P}_{FB}(M, \kappa, \alpha, \eta)$.

Margin Assumption Examples

One covariate $X \sim \mathrm{Uniform}[0,1]$.

- Linear: $\tau(X) = \beta_0 + \beta_1 X$. $P(|\tau(X)| \leq t) = \frac{2}{\beta_1} t$.
  Margin $\alpha = 1$ and $\eta = \beta_1/2$.
- Discontinuous at zero: for $h > 0$

$$\tau(X) = \begin{cases} X - h & \text{for } X \leq 0 \\ X + h & \text{for } X > 0 \end{cases}$$

  Margin α can be arbitrarily large, $\alpha = +\infty$, and $\eta = h$.
- Low margin: $\tau(X) = \left(\frac{1}{2} - X\right)^3$. $P(|\tau(X)| \leq t) = 2t^{1/3}$.
  Margin $\alpha = \frac{1}{3}$, $\eta = 1/8$.

Convergence rates with a margin assumption

Theorem 2.3: *Let $\mathcal{P}_{FB}(M,\kappa,\alpha,\eta)$ be a class of DGPs satisfying Bounded Outcome, Strict Overlap, $G^*_{FB} \in \mathcal{G}$, & MA with margin coefficient $\alpha > 0$. Then,*

$$\sup_{P \in \mathcal{P}_{FB}(M,\kappa,\alpha,\eta)} E_{P^n}\left[W(G^*_{FB}) - W(\hat{G}_{EWM})\right] \leq c_3 \left(\frac{v}{n}\right)^{\frac{1+\alpha}{2+\alpha}} .$$

where c_3 is a constant that depends only on (M,κ,α).

Theorem 2.4: *Let $\mathcal{P}_{FB}(M,\kappa,\alpha,\eta)$ be a class of DGPs satisfying Bounded Outcomes and Strict Overlap. Let $\mathcal{G}$ be a VC-class, $v \geq 2$. Then, for any treatment choice rule $\hat{G}$*

$$\sup_{P \in \mathcal{P}_{FB}(\alpha,\eta)} E_{P^n}\left[W(G^*_{FB}) - W(\hat{G})\right] \geq c_4 \left(\frac{v-1}{n}\right)^{\frac{1+\alpha}{2+\alpha}} ,$$

for all $n \geq \max\{(M/\eta)^2, 4^{2+\alpha}\}(v-1)$.

What do we learn from the margin assumption results?

These results are of theoretical value, since they do not affect estimation of EWM rules.

Pointwise regret convergence rates (holding distribution P fixed): a lot of interesting simulation examples you could come up with satisfy the margin assumption and yield a variety of pointwise convergence rates. The margin assumption explains a lot of this variation.

In some application, the margin assumption may hold uniformly in $\mathcal{P}$. For example, if it is known ex ante that $\tau(x)$ is monotonic in x and varies substantially, i.e., the absolute value of the derivative $\left|\frac{\partial \tau(x)}{dx}\right|$ is bounded away from zero.

Unknown propensity score $e(X)$

- Hybrid of EWM and regression plug-in

$$\hat{G}_{m-hybrid} \in \arg\max_{G \in \mathcal{G}} E_n \left[\hat{\tau}^m(X_i) \cdot 1\{X_i \in G\} \right]$$
$$\hat{\tau}^m(X_i) \equiv \hat{m}_1(X_i) - \hat{m}_0(X_i)$$

- Hybrid of EWM and propensity score plug-in

$$\hat{G}_{e-hybrid} \in \arg\max_{G \in \mathcal{G}} E_n \left[\hat{\tau}_i^e \cdot 1\{X_i \in G\} \right]$$
$$\hat{\tau}_i^e \equiv \left[\frac{Y_i D_i}{\hat{e}(X_i)} - \frac{Y_i(1-D_i)}{1-\hat{e}(X_i)} \right] \cdot 1\left\{ \varepsilon_n \leq \hat{e}(X_i) \leq 1 - \varepsilon_n \right\}$$

- Theorems 2.5, 2.6 establish rate upper bounds, which are the maximum of the nonparametric rate and the EWM rate
- We do not know whether these rate bounds are sharp.

EWM for non-additive social welfare functions

The EWM idea (maximizing a sample analogue of the welfare function) may be applicable to problems with social welfare functions that are not additive over $x \in \mathcal{X}$.

Examples: externalities, general equilibrium effects.

Source: Kitagawa and Tetenov (2017), "Equality-Minded Treatment Choice" Cemmap working paper CWP10/17

We extend the EWM idea to treatment choice with rank-dependent social welfare functions.

Social welfare functions

Y - individual income with distribution $F(y)$.

Two major types of social welfare functions:

1. Additively separable in individual incomes (Atkinson, 1970)

$$W(F) = \int U(y)dF(y)$$

Redistributive preferences are expressed by a concave $U(y)$.

The previously-discussed "Empirical Welfare Maximization" paper (Kitagawa and Tetenov, 2017) covers this problem, it is sufficient to replace outcomes Y_i with $U(Y_i)$.

2. Rank-dependent social welfare Mehran (1976), Weymark (1981), Yaari (1988), Ben Porath and Gilboa (1994).

$$W(F) = \int Y \cdot \omega\,(\text{Rank}(Y))\,di$$

Equality-minded: decreasing $\omega(\cdot)$, lower welfare weight is given to incomes at higher quantiles.

Equivalent representation:

$$W(F) = \int \Lambda(1 - F(y))dy$$

Convex, differentiable, decreasing function $\Lambda(\cdot) : [0,1] \rightarrow [0,1]$

$$\omega(r) = -\frac{d\Lambda(r)}{dr}$$

Rank-dependent welfare functions are closely linked to inequality indices

Could be expressed as

$$W(F) = \mu(F)(1 - I(F))$$

$\mu(F)$ - average income
$I(F)$ - an index of inequality (e.g. Gini when $\omega(r) = 2(1 - r)$)

Performance of a policy is summarized by the representative income: Distribution F is as good as everyone having income $Y = W(F)$.

Equality-minded treatment choice

A randomized treatment rule $\delta : \mathcal{X} \to [0,1]$ specifies the fraction of individuals with covariates X who will be treated.

It generates income distribution with CDF

$$F_{\delta}(y) \equiv \int_{\mathcal{X}} \left[(1-\delta(x))F_{Y_0|X} + \delta(x)F_{Y_1|X}\right] dP(X),$$

We would like to find δ that maximizes $W(F_{\delta})$.

Challenges:

1. A class of $\delta(\cdot)$ can be huge.
2. The value of the policy is not additive across subgroups of the population, i.e., what policy is given to one subpopulation affects what policy should be given to other subpopulations!
3. No closed-form solution for the optimal treatment rule.

Sufficiency of non-randomized treatment rules

Proposition 1:
If $W(\cdot)$ is an equality-minded welfare criterion, then for any treatment rule δ there exists a non-randomized treatment rule $\delta' = 1\{X \in G\}$ such that $W(F_{\delta'}) \geq W(F_{\delta})$.

(follows from the convexity of $\Lambda(\cdot)$ in the representation)

We index non-randomized treatment rules by their decision sets $G \in \mathcal{G}$.
$\delta(X) = 1\{X \in G\}$

Social welfare will be denoted by $W(G)$.

Empirical Welfare Maximization

We propose maximizing a sample analog of the social welfare function

$$\widehat{G} \equiv \arg\max_{G \in \mathcal{G}} \widehat{W}(G), \quad \widehat{W}(G) = \int_0^M \Lambda(0 \vee (1 - \widehat{F}_G(y)))dy$$

where $\widehat{F}_G(y)$ is the sample analog of the income CDF

$$\widehat{F}_G(y) \equiv \frac{1}{n}\sum_{i=1}^{n}\left[\frac{D_i \cdot 1\{Y_i \leq y\}}{e(X_i)} \cdot 1\{X_i \in G\} + \right. \\ \left. + \frac{(1-D_i)\cdot 1\{Y_i \leq y\}}{1-e(X_i)} \cdot 1\{X_i \notin G\}\right].$$

$e(X_i)$ is the propensity score of observation i

$\widehat{F}_G(y)$ could be normalized to a proper CDF.

Welfare regret upper bound

Proposition 2

Let $\mathcal{P}$ be a class of DGPs satisfying assumptions Bounded Outcomes and Strict Overlap. Let $\mathcal{G}$ be a VC-class of treatment choice rules. If W is an equality-minded SWF with $\Lambda(\cdot)$ that is convex, differentiable, and has a bounded derivative, then

$$\sup_{P \in \mathcal{P}} \left[\sup_{G \in \mathcal{G}} W(G) - E_{P^n} \left[W(\widehat{G}) \right] \right] \leq C \cdot |\Lambda'(0)| \frac{M}{\kappa} \sqrt{\frac{v}{n}}.$$

Welfare regret lower bound

Proposition 3 If $|\Lambda'(t^*)| > 0$ for some $t^* \in (0,1)$, then for any non-randomized treatment choice rule $\tilde{G}$,

$$\sup_{P \in \mathcal{P}} \left[\sup_{G \in \mathcal{G}} W(G) - E_{P^n} \left[W(\widehat{G}) \right] \right] \geq |\Lambda'(t^*)| M \frac{e^{-4}}{4} \sqrt{\frac{v-1}{n}}$$

for all $n >= 4(v-1)t^*$.

$1/\sqrt{n}$ is the minimax optimal uniform convergence rate over $\mathcal{P}$ in terms of welfare regret.