Topics in Causal Inference and Policy Learning with Applications to Precision Medicine

PhD defense

Pan Zhao

September 4, 2024

PreMeDICaL, Inria & Université de Montpellier

Advisors: Julie Josse (Inria) & Antoine Chambaz (Université Paris Cité)

Policy learning

Learning a treatment assignment policy is pivotal across various domains, for instance:

- individualized treatment rule in precision medicine
- personalized advertising in marketing
- educational/training programs in public policy

Basic causal setup:[1]

- data $O = (X, A, Y) \sim P$ with covariates $X \in \mathcal{X}$, treatment A and outcome Y
- complete data $\mathbb{O} = (X, A, Y(0), Y(1)) \sim \mathbb{P}$ w/ potential outcomes $Y(0), Y(1)$
- policy $d : \mathcal{X} \to \mathcal{A} = \{0, 1\}$

[1] Athey, S., & Wager, S. (2021). Policy learning with observational data. Econometrica, 89(1), 133-161.

Main approaches

(A) Heterogeneous treatment effects estimation:

$$x \mapsto \mathrm{CATE}_{\mathbb{P}}(x) = E_{\mathbb{P}}[Y(1) - Y(0) \mid X = x]$$
$$\rightsquigarrow \quad d^{\mathrm{opt}}(x) = I\{\mathrm{CATE}_{\mathbb{P}}(x) > 0\}$$

(B) Direct policy search:
define value function $d \mapsto V_{\mathbb{P}}(d) = E_{\mathbb{P}}[Y(1)d(X) + Y(0)(1 - d(X))]$

$$\begin{aligned} d^{\mathrm{opt}} &= \arg\max_{d \in \mathcal{D}} V_{\mathbb{P}}(d) \\ &= \arg\max_{d \in \mathcal{D}} E_{\mathbb{P}}[(Y(1) - Y(0))d(X) + Y(0)] \\ &= \arg\max_{d \in \mathcal{D}} E[\mathrm{CATE}_{\mathbb{P}}(X)d(X)] \end{aligned}$$

Possibly subject to application-specific constraints, such as budget, fairness, simplicity

Direct policy search - identification

Under consistency, unconfoundedness and positivity:

- inverse probability weighting (IPW):

$$V_{\mathbb{P}}(d) = E_P\left[\frac{Y\,I\{A = d(X)\}}{Pr_P(A = d(X) \mid X)}\right]$$

- outcome regression (OR):

$$V_{\mathbb{P}}(d) = E_P\left\{E_P[Y \mid A = d(X), X]\right\}$$

- Augmented IPW (AIPW):

$$V_{\mathbb{P}}(d) = E\left\{\frac{I\{A = d(X)\}}{Pr_P(A \mid X)}\left(Y - E_P[Y \mid A = d(X), X]\right) + E_P[Y \mid A = d(X), X]\right\}$$

Consistency, excess risk bound, (minimax) regret bound etc. can be established[2]

[2] Zhao, Y., Zeng, D., Rush, A. J., & Kosorok, M. R. (2012). Estimating individualized treatment rules using outcome weighted learning. Journal of the American Statistical Association, 107(499), 1106-1118.

Main Contributions

Research articles & projects

Publication and preprints:

- A Semiparametric Instrumented Difference-in-Differences Approach to Policy Learning. Major revision at Biometrika. IMS Hannan Graduate Student Award
- Positivity-free Policy Learning with Observational Data. Proceedings of The 27th International Conference on Artificial Intelligence and Statistics, PMLR 238:1918-1926, 2024.
- Efficient and robust transfer learning of optimal individualized treatment regimes with right-censored survival data. R & R at Journal of Machine Learning Research.
- Learning, Evaluating and Analysising An Individualized Decision Support Rule with Application to Early Intervention in Intensive Care Unit. In preparation.

Ongoing projects:

- w/ Yifan Cui (Zhejiang University): Variable Importance for Heterogeneous Treatment Effects with Survival Data and Nonparametric Inference at the Parameter Space Boundary.
- w/ Oliver Dukes & Stijn Vansteelandt (Ghent University): Orthogonal Statistical Learning for Nonparametric Instrumental Variables.
- w/ Oliver Dukes & Bo Zhang (Fred Hutch): Estimating the risk and relative vaccine efficacy of updated vaccine regimens using historical phase 3 clinical trials and immunobridging data.

Other activities

Software:

- CRAN Task View: Causal Inference
- R package `missSuperLearner`
- R implementation of all projects available on GitHub: `https://github.com/panzhaooo`

Academic visit at Ghent University w/ Oliver Dukes & Stijn Vansteelandt.

Talks:

- contributed: IDESP 2021, JDS 2022, IMS ICSDS 2023
- invited: JSM 2023, Ghent Causal Meeting, IMS APRM 2024, AISTATS 2024

Introduction to Instrumental Variable

IV setup and DAG

Basic setup:

- observed data $O = (X, Z, A, Y) \sim P$: binary instrument Z and treatment A, covariates X and outcome Y
- unmeasured confounder U
- complete data $\mathbb{O} = (X, U, Z, A(0), A(1), Y(0), Y(1)) \sim \mathbb{P}$ w/ potential outcomes: $Y = Y(1)A + Y(0)(1 - A)$, $A = A(1)Z + A(0)(1 - Z)$

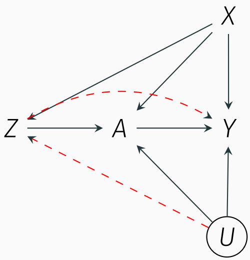

Figure 1: DAG for instrumental variable setup (red: not allowed).

Imbens & Angrist 1994

Causal assumptions for IV: under $\mathbb{P}$,

- exclusion: $Y(a) = Y(a, z)$ for $a, z \in \{0, 1\}$
- independence: $Z \perp \{Y(0), Y(1), A(1), A(0)\}$
- relevance: $E[A \mid Z = 1] > E[A \mid Z = 0]$
- monotonicity: $A(1) \geq A(0)$
  - "Always taker" $A(1) = A(0) = 1$
  - "Complier" $A(1) = 1, A(0) = 0$
  - ~~"Defier" $A(1) = 0, A(0) = 1$~~
  - "Never taker" $A(1) = A(0) = 0$

Local average treatment effect

$$\mathrm{Wald}_P \stackrel{\mathrm{def}}{=} \frac{E_P[Y \mid Z=1] - E_P[Y \mid Z=0]}{E_P[A \mid Z=1] - E_P[A \mid Z=0]} = E_{\mathbb{P}}[Y(1) - Y(0) \mid A(1) > A(0)]$$

Simple proof:

- by independence, $E_P[A \mid Z=1] - E_P[A \mid Z=0] = E_{\mathbb{P}}[A(1) - A(0)]$
- similarly,

$$\begin{aligned} & E_P[Y \mid Z=1] - E_P[Y \mid Z=0] \\ &= E_{\mathbb{P}}[(Y(1) - Y(0))A(1) + Y(0) \mid Z=1] \\ &\quad - E_{\mathbb{P}}[(Y(1) - Y(0))A(0) + Y(0) \mid Z=0] \\ &= E_{\mathbb{P}}[(Y(1) - Y(0))(A(1) - A(0))] \end{aligned}$$

- by monotonicity

$$\frac{E_{\mathbb{P}}[(Y(1) - Y(0))(A(1) - A(0))]}{E_{\mathbb{P}}[A(1) - A(0)]} = E_{\mathbb{P}}[Y(1) - Y(0) \mid A(1) > A(0)]$$

From LATE to CATE – Wang & Tchetgen Tchetgen 2018

$$\mathrm{CATE}_{\mathbb{P}}(x) = \text{def} E_{\mathbb{P}}[Y(1) - Y(0) \mid X = x]$$
$$= \frac{E_P[Y \mid Z = 1, X = x] - E_P[Y \mid Z = 0, X = x]}{E_P[A \mid Z = 1, X = x] - E_P[A \mid Z = 0, X = x]}$$

Causal assumptions: under $\mathbb{P}$,

- exclusion: $Y(a) = Y(a, z)$ for $a, z \in \{0, 1\}$
- independence: $Z \perp U \mid X$
- relevance: $Z \not\perp A \mid X$
- $Y(A) \perp \{A, Z\} \mid \{X, U\}$
- either no additive $U - Z$ interaction

  $$E_P[A \mid Z = 1, X, U] - E_P[A \mid Z = 0, X, U] = E_P[A \mid Z = 1, X] - E_P[A \mid Z = 0, X]$$

  or no additive $U - a$ interaction

  $$E_{\mathbb{P}}[Y(1) - Y(0) \mid X, U] = E_{\mathbb{P}}[Y(1) - Y(0) \mid X]$$

Regression, IPW and efficient multiply robust estimators are provided

IV for policy learning - Cui & Tchetgen Tchetgen 2018

Let $\delta_P(X) = Pr_P(A = 1 \mid Z = 1, X) - Pr_P(A = 1 \mid Z = 0, X)$

Causal assumptions: under $\mathbb{P}$,

- exclusion, independence, relevance
- no unmeasured common effect modifier:

$$Cov_P\{Pr_P(A = 1 \mid Z = 1, X, U) - Pr_P(A = 1 \mid Z = 0, X, U), E_{\mathbb{P}}[Y(1) - Y(0) \mid X, U] \mid X\} = 0$$

$\rightarrow$ identification of the optimal policy:

$$d^{\text{opt}} = \arg\max_{d \in \mathcal{D}} E_P\left[\frac{(2Z-1)(2A-1)Y\, I\{A=d(X)\}}{\delta_P(X) Pr_P(Z|X)}\right] = \arg\max_{d \in \mathcal{D}} E_P\left[\frac{Y\, I\{Z=d(X)\}}{\delta_P(X) Pr_P(Z|X)}\right]$$

- independent compliance type:

$$\delta_P(X) = Pr_P(A = 1 \mid Z = 1, X, U) - Pr_P(A = 1 \mid Z = 0, X, U)$$

$\rightarrow$ identification of the value function:

$$V_{\mathbb{P}}(d) = E_P\left[\frac{(2Z-1)(2A-1)Y\, I\{A = d(X)\}}{\delta_P(X) Pr_P(Z \mid X)}\right]$$

Introduction to Difference-in-Differences

Difference-in-Differences

Basic setup:

- two time points $T \in \{0,1\}$
- covariates X, treatments $A \in \{0,1\}$ or $(A_0, A_1) \in \{0,1\}^2$, outcomes Y or (Y_0, Y_1)
- potential outcomes $Y_t(a), t, a \in \{0,1\}$

Two observed data structures:

- repeated cross-section data: $O = (X, A, Y, T)$, with $Y = Y_T(A)$
- panel data: $O = (X, A_0, Y_0, A_1, Y_1)$, with $Y_t = Y_t(A_t), t \in \{0,1\}$

The complete and observed laws are $\mathbb{P}$ and P

DiD illustration

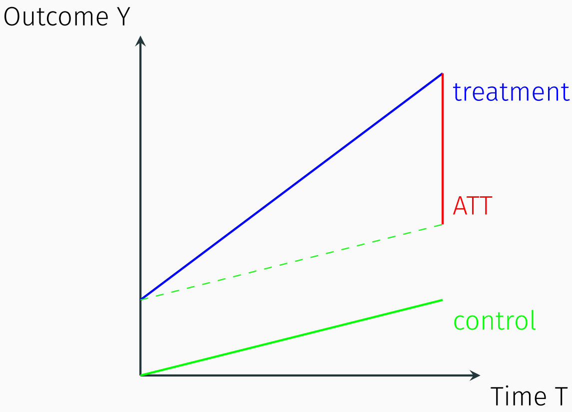

Figure 2: A simple illustration of DiD identification.

Parallel trends & average treatment effect on the treated

(Conditional) parallel trends assumption:

$$E_{\mathbb{P}}[Y_1(0) - Y_0(0) \mid A = 1, X] = E_P[Y_1(0) - Y_0(0) \mid A = 0, X]$$

Simple proof:

$$\begin{aligned}
\mathrm{ATT}_{\mathbb{P}} &\overset{\text{def}}{=} E_{\mathbb{P}}[Y_1(1) - Y_1(0) \mid A = 1] \\
&= E_{\mathbb{P}}[Y_1(1) - Y_1(0) + Y_0(0) - Y_0(0) \mid A = 1] \\
&= E_{\mathbb{P}}[Y_1(1) \mid A = 1] - E_{\mathbb{P}}[Y_1(0) - Y_0(0) \mid A = 0] - E_{\mathbb{P}}[Y_0(0) \mid A = 1] \\
&= E_{\mathbb{P}}[Y_1 - Y_0 \mid A = 1] - E_{\mathbb{P}}[Y_1 - Y_0 \mid A = 0]
\end{aligned}$$

Instrumented Difference-in-Differences

Instrumented DiD setup

First introduced by Ye et al. 2022 for (conditional) ATE, also structural mean models by Vo et al. 2023

Basic setup:

- two time points $T \in \{0, 1\}$
- binary instrument Z and treatment $A \in \{0, 1\}$ or $(A_0, A_1) \in \{0, 1\}^2$
- covariates X and unmeasured confounder $U = (U_0, U_1)$
- potential outcomes $Y_t(a), t, a \in \{0, 1\}$

Two observed data structure:

- repeated cross-section data $O = (X, Z, A, Y, T)$, with $Y = Y_T(A)$
- panel data $O = (X, Z, A_0, Y_0, A_1, Y_1)$, with $Y_t = Y_t(A_t), t \in \{0, 1\}$

The complete and observed laws are $\mathbb{P}$ and P

Instrumented DiD DAG: trend scale

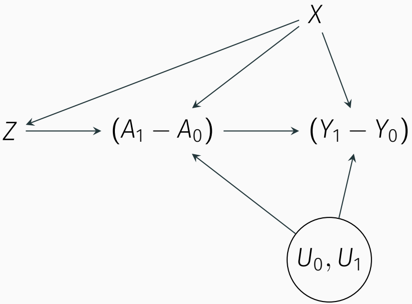

Figure 3: DAG for instrumented DiD on the trend scale.

Instrumented DiD DAG: two time points

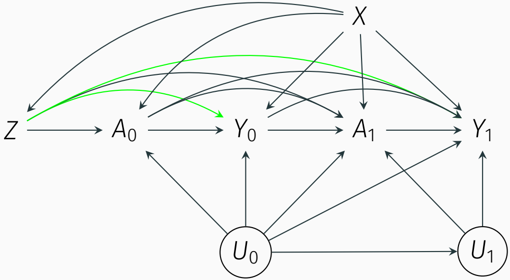

Figure 4: DAG for instrumented DiD over two time points.

Motivation & example

- IV for DiD: e.g. haphazard encouragement targeted at a subpopulation toward faster uptake of the exposure or a surrogate of such encouragement (Ye et al. 2022)
- Longitudinal randomized experiment: after a baseline period, some individuals are randomly selected to be encouraged to take the treatment regardless of treatment history
- See Ye et al. 2022 for an analysis of the effect of cigarette smoking on lung cancer mortality

Instrumented DiD to Policy Learning

Recap on policy learning

Optimal policy given by

$$d_t^{\mathrm{opt}} = \arg\max_{d \in \mathcal{D}} V_{\mathbb{P},t}(d) = \arg\max_{d \in \mathcal{D}} E_P[\mathrm{CATE}_{\mathbb{P},t}(X) d(X)].$$

- assume stable treatment effect over time
- directly maximize some functional $d \mapsto V_{\mathbb{P}}(d)$ similarly
- CATE-based approaches

Causal assumptions

Let $\pi_P(t, z, x) = \Pr_P(T = t, Z = z \mid X = x)$. Under $\mathbb{P}$,

- consistency: $A = A_T(Z)$ and $Y = Y_T(A)$
- positivity: $c_1 < \pi_P(t, z, x) < 1 - c_1$ for some $0 < c_1 < 1/2$
- random sampling:
  $T \perp \{A_t(z), Y_t(a) : t = 0, 1, z = 0, 1, a = 0, 1\} \mid X, Z$
- stable treatment effect over time:
  $E_{\mathbb{P}}[Y_0(1) - Y_0(0) \mid X] = E_{\mathbb{P}}[Y_1(1) - Y_1(0) \mid X]$
- trend relevance:
  $E_{\mathbb{P}}[A_1(1) - A_0(1) \mid Z = 1, X] \neq E_{\mathbb{P}}[A_1(0) - A_0(0) \mid Z = 0, X]$
- independence & exclusion restriction:
  $Z \perp \{A_t(1), A_t(0), Y_t(1) - Y_t(0), Y_1(0) - Y_0(0) : t = 0, 1\} \mid X$
- no unmeasured common effect modifier:
  $Cov_{\mathbb{P}}\{A_t(1) - A_t(0), Y_t(1) - Y_t(0) \mid X\} = 0$ for $t = 0, 1$

Identification of optimal policy

For $C \in \{A, Y\}$, let $\mu_{P,C}(t, z, x) = E_P[C \mid T = t, Z = z, X = x]$, and $\delta_{P,C}(x) = \mu_C(1, 1, x) - \mu_C(0, 1, x) - \mu_C(1, 0, x) + \mu_C(0, 0, x)$

- CATE-based approach:

$$d^{\mathrm{opt}} = \arg\max_{d \in \mathcal{D}} E_P \left[\frac{\delta_{P,Y}(X)}{\delta_{P,A}(X)} d(X) \right]$$

- novel IPW formula 1:

$$d^{\mathrm{opt}} = \arg\max_{d \in \mathcal{D}} E_P \left[\frac{(2Z-1)(2T-1)(2A-1)YI\{A = d(X)\}}{\pi_P(T, Z, X)\delta_{P,A}(X)} \right]$$

- novel IPW formula 2:

$$d^{\mathrm{opt}} = \arg\max_{d \in \mathcal{D}} E_P \left[\frac{(2T-1)YI\{Z = d(X)\}}{\pi_P(T, Z, X)\delta_{P,A}(X)} \right]$$

$\rightarrow$ simple plug-in estimators can be constructed

Semiparametric efficiency

Efficient influence function (Ye et al. 2022)

$$\Delta_P(O) = \frac{\delta_{P,Y}(X)}{\delta_{P,A}(X)} + \frac{(2Z-1)(2T-1)}{\pi_P(T,Z,X)\delta_{P,A}(X)} \left\{ Y - \mu_{P,Y}(T,Z,X) - \frac{\delta_{P,Y}(X)}{\delta_{P,A}(X)}(A - \mu_{P,A}(T,Z,X)) \right\},$$

Recall the optimization tasks:

$$\arg\max_{d \in \mathcal{D}} E_P[W_P^{(1)} I\{A = d(X)\}], \quad \arg\max_{d \in \mathcal{D}} E_P[W_P^{(2)} I\{Z = d(X)\}]$$

where

$$W_P^{(1)} = \frac{(2A-1)\delta_{P,Y}(X)}{\delta_{P,A}(X)} + \frac{(2A-1)(2Z-1)(2T-1)}{\pi_P(T,Z,X)\delta_{P,A}(X)} \left\{ Y - \mu_{P,Y}(T,Z,X) - \frac{\delta_{P,Y}(X)}{\delta_{P,A}(X)}(A - \mu_{P,A}(T,Z,X)) \right\}$$

and

$$W_P^{(2)} = \frac{(2Z-1)\delta_{P,Y}(X)}{\delta_{P,A}(X)} + \frac{2T-1}{\pi_P(T,Z,X)\delta_{P,A}(X)} \left\{ Y - \mu_{P,Y}(T,Z,X) - \frac{\delta_{P,Y}(X)}{\delta_{P,A}(X)}(A - \mu_{P,A}(T,Z,X)) \right\}$$

Multiple robustness

Optimal policy identified by

$$\arg\max_{\mathcal{D}} E_P\left[W_P^{(1)} I\{A = d(X)\}\right] = \arg\max_{\mathcal{D}} E_P\left[W_P^{(2)} I\{Z = d(X)\}\right]$$
$$= \arg\max_{\mathcal{D}} E_P\left[\Delta_P(X) d(X)\right]$$

Under the union model $\mathcal{M}_1 \cup \mathcal{M}_2 \cup \mathcal{M}_3$

- $\mathcal{M}_1$: models for π_P and $\delta_{P,A}$ are correct
- $\mathcal{M}_2$: models for π_P and $\delta_{P,Y}/\delta_{P,A}$ are correct
- $\mathcal{M}_3$: models for $\delta_{P,Y}/\delta_{P,A}$ and $\mu_{P,C}(0,0,\cdot)$, $\mu_{P,C}(1,0,\cdot)$, $\mu_{P,C}(0,1,\cdot)$ for $C \in \{A, Y\}$ are correct

Double/Debiased machine learning

Cross-fitted estimator:

$$\hat{M}_{CF} = \frac{1}{K}\sum_{k=1}^{K} P_{n,k}\{\Delta(O; \hat{\mu}_{A,-k}, \hat{\mu}_{Y,-k}, \hat{\pi}_{-k})d(X)\},$$

1. randomly split the data into K folds;
2. for $k = 1, \ldots, K$, learn the nuisance parameters $\mu_{P,A}, \mu_{P,Y}, \pi_P$ with $\hat{\mu}_{A,-k}, \hat{\mu}_{Y,-k}, \hat{\pi}_{-k}$ using data excluding the k-th fold; then evaluate the value on the k-th fold;
3. average the value estimates from the K folds.

Asymptotic analysis of policy learning

Focus on a class of feasible policies $\mathcal{D} = \left\{x \mapsto I\{\eta^\top x > 0\} : \eta \in \mathbb{H}\right\}$

Policy estimator:

$$\hat{\eta} = \arg\max_{\eta \in \mathbb{H}} \hat{M}(\eta) = \arg\max_{\eta \in \mathbb{H}} \frac{1}{n} \sum_{i=1}^{n} \hat{\Delta}(O_i) I\{\eta^\top X_i > 0\},$$

where $\hat{\Delta}$ is the estimator of Δ_P obtained by substitution.

Under certain regularity and rate of convergence conditions:

- $\|\hat{\eta} - \eta^*\|_2 = O_p(n^{-1/3})$
- $\sqrt{n}\{M(\hat{\eta}) - M(\eta^*)\} = o_p(1)$
- $\sqrt{n}\{\hat{M}(\hat{\eta}) - M(\eta^*)\} \rightsquigarrow \mathcal{N}(0, \sigma^2)$

Extension to panel data: identification

- Analog causal assumptions for panel data
- Alternatively, Vo et al. 2023 consider sequential ignorability for structural mean model
- We also prove identification if, under $\mathbb{P}$,
  - sequential ignorability: $Y_t(a) \perp A_t \mid U, X, Z$ for $t, a = 0, 1$
  - no additive interaction of either

$$E_{\mathbb{P}}[A_1 - A_0 \mid X, U, Z = 1] - E_{\mathbb{P}}[A_1 - A_0 \mid X, U, Z = 0] = E_{\mathbb{P}}[A_1 - A_0 \mid X, Z = 1] - E_{\mathbb{P}}[A_1 - A_0 \mid X, Z = 0]$$

or

$$E_{\mathbb{P}}[Y_t(1) - Y_t(0) \mid U, X] = E_{\mathbb{P}}[Y_t(1) - Y_t(0) \mid X]$$

$\rightarrow$ CATE identified by

$$\mathrm{CATE}_{\mathbb{P}}(x) = \frac{E_P[Y_1 - Y_0 \mid X = x, Z = 1] - E_P[Y_1 - Y_0 \mid X = x, Z = 0]}{E_P[A_1 - A_0 \mid X = x, Z = 1] - E_P[A_1 - A_0 \mid X = x, Z = 0]} \stackrel{\text{def}}{=} \tau_P(x)$$

Semiparametric efficiency

EIF given by

$$
\begin{aligned}
&\phi_{\text{panel,P}}(o) \\
&= \frac{\delta_{P,Y,1}(x) - \delta_{P,Y,0}(x)}{\delta_{P,A,1}(x) - \delta_{P,A,0}(x)} \\
&\quad - \frac{z - \pi_{P,Z}(x)}{\pi_{P,Z}(x)(1 - \pi_{P,Z}(x))(\delta_{P,A,1}(x) - \delta_{P,A,0}(x))^2} \{(y_1 - y_0)(\delta_{P,A,1}(x) - \delta_{P,A,0}(x)) \\
&\quad -(a_1 - a_0)(\delta_{P,Y,1}(x) - \delta_{P,Y,0}(x)) + \delta_{P,Y,1}(x)\delta_{P,A,0}(x) - \delta_{P,Y,0}(x)\delta_{P,A,1}(x)\} - \tau_P(x)
\end{aligned}
$$

Optimal policy:

$$
\arg\max_{\mathcal{D}} E_P\left[\frac{\delta_{P,Y,1}(X) - \delta_{P,Y,0}(X)}{\delta_{P,A,1}(X) - \delta_{P,A,0}(X)} d(X)\right] = \arg\max_{\mathcal{D}} E_P\left[\Delta_{\text{panel,P}}(X) d(X)\right],
$$

Asymptotic results can be obtained similarly

Simulation

Data-generation process:

$$X_1, X_2 \sim \mathcal{N}(0,1), U_0, U_1 \sim \mathrm{Bridge}(0.5), T \sim \mathrm{Bernoulli}(0.5) \text{ independently}$$

$$Pr(A_0 = 1 \mid Z, U, X) = \mathrm{expit}(2 - 7Z + 0.2U_0 + 2X_1),$$

$$Pr(A_1 = 1 \mid Z, U, X) = \mathrm{expit}(-1.5 + 5Z - 0.15U_1 + 1.5X_2),$$

$$(Y_0 \mid Z, U, X, A_0) \sim \mathcal{N}(\mu_0, 1), (Y_1 \mid Z, U, X, A_1) \sim \mathcal{N}(\mu_1, 1)$$

where

$$\mu_0 = 200 + 10(A_0(1.5X_1 + 2X_2 - 0.5) + 0.5U_0 + 2Z + 1.5X_1 + 2X_2)$$

$$\mu_1 = 240 + 10(A_1(1.5X_1 + 2X_2 - 0.5) + 0.5U_1 + 2Z + 2X_1 + 1.5X_2)$$

Evaluate by percentage of correct decisions (PCD) of estimated $\hat{d}(x)$

$$1 - N^{-1} \sum_{i=1}^{N} |\hat{d}(X_i) - d^{\mathrm{opt}}(X_i)|$$

Compare with standard IV methods (Cui & Tchetgen Tchetgen 2018)

Correctly specified parametric models, or random forests (`grf`)

Results

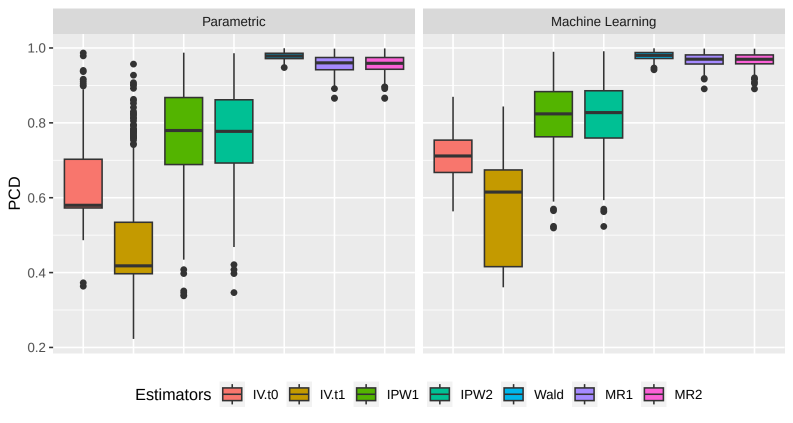

Figure 5: Results of the estimated optimal policies, using parametric models (left) or machine learning (right).

Data application – Australian Longitudinal Survey

- Conducted annually since 1984, mainly on the dynamics of the youth labour market, including basic demographic, labour market and background variables, and topics related to the main labour market theme
- Card 2001: endogeneity of education might partially explain the continuing interest “in this very difficult task of uncovering the causal effect of education in labor market outcomes”
- Cross-section data from 1984 and 1985 waves (Vella 1994)

Policies	intercept	born_australia	married	uni_mem	gov_emp	age	year_expe
IV.t0	0.4442	−0.4547	0.1311	−0.1179	−0.5181	0.0080	−0.5444
IV.t1	−0.2518	−0.3103	0.2445	−0.6157	−0.1406	0.2015	−0.5840
IPW1	−0.4203	−0.0847	0.5454	−0.3941	−0.5690	0.0299	0.1969
IPW2	−0.2503	−0.0529	0.6051	−0.4384	−0.5801	0.0207	0.1980
Wald	0.5032	0.3891	0.4738	0.5755	−0.1656	−0.0772	0.0793
MR1	−0.0513	0.1341	−0.6039	0.4127	0.5861	−0.0226	−0.3168
MR2	0.5480	−0.3937	−0.4072	0.4393	0.4167	−0.0302	−0.1064

- Coefficients should be interpreted with caution
- Majority vote from Wald, MR1, MR2 estimators

Discussion

- By monotonicity assumption $A_t(1) \geq A_t(0)$ for $t = 0, 1$
  $\rightarrow$ optimal policy for compliers
- Fuzzy DiD in econometrics (De Chaisemartin & d'Haultfoeuille 2018)
- DiD on multiple time points, or continuous time
- Weak IV, continuous IV

Thank you! & Questions?

Positivity-free Policy Learning with Observational Data

Assign treatment 1 with probability

$$d(x) = \frac{\delta(x)\,\pi(x)}{\delta(x)\,\pi(x) + 1 - \pi(x)}$$

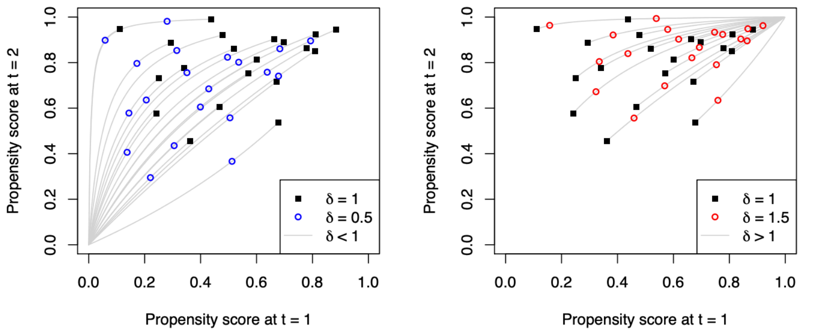

Figure 1: Observational propensity scores for $n = 20$ simulated units in a study with $T = 2$ timepoints, and their values under incremental interventions based on different δ values ($\delta \leq 1$ in the left plot, $\delta \geq 1$ in the right).

Backup slides II

Efficient and robust transfer learning of optimal individualized treatment regimes with right-censored survival data

Target super population

(unknown) **Source sampling** I_S

Complete source sample
$\{T_i(1), T_i(0), X_i, I_{S,i} = 1, I_{T,i} = 0\}_{i=1}^{n}$

Treatment assignment A
Censoring C

Observed source sample
$\{X_i, A_i, U_i, \Delta_i, I_{S,i} = 1, I_{T,i} = 0\}_{i=1}^{n}$

(known) **Target sampling** I_T

Complete target sample
$\{T_i(1), T_i(0), X_i, I_{S,i} = 0, I_{T,i} = 1\}_{i=n+1}^{n+m}$

Only observe covariates X

Observed target sample
$\{X_i, I_{S,i} = 0, I_{T,i} = 1\}_{i=n+1}^{n+m}$