Machine Learning for Dynamic Discrete Choice

Vira Semenova

(joint with Victor Chernozhukov and Whitney Newey)

MIT

January 4, 2020

Introduction

- Many covariates in modern high-dimensional data sets create challenges in computation and inference
- Machine learning (ML) methods is a way to reduce the dimension and relax parametric assumptions
- Existing work with ML methods: average treatment effects, static models of strategic interaction (Chernozhukov et al. (2017a), Chernozhukov et al. (2017b))
- This talk: introduce ML methods for estimation of value function

Motivating example: Rust (1987)+static heterogeneity

- Rust (1987)+bus characteristics:

  per-period utility function $u(s; a; \epsilon) = \begin{cases} -R + \epsilon(0), & a = 0 \\ -\mu \cdot s + \epsilon(1), & a = 1 \end{cases}$

  - s is mileage, $a \in \{1, 0\}$, $a = 0$ is bus replacement
  - future mileage s_{t+1} depends on (x_t, a_t) where state $x_t = (s_t, \text{characteristics})$
  - $\epsilon(0), \epsilon(1)$ unobserved shock
  - $\theta = (R, \mu)$ structural parameter

Motivating example: Rust (1987)+static heterogeneity

- Rust (1987)+bus characteristics:

  per-period utility function $u(s; a; \epsilon) = \begin{cases} -R + \epsilon(0), & a = 0 \\ -\mu \cdot s + \epsilon(1), & a = 1 \end{cases}$

  - s is mileage, $a \in \{1, 0\}$, $a = 0$ is bus replacement
  - future mileage s_{t+1} depends on (x_t, a_t) where state $x_t = (s_t, \text{characteristics})$
  - $\epsilon(0), \epsilon(1)$ unobserved shock
  - $\theta = (R, \mu)$ structural parameter
- Goal: estimate and conduct inference about expected value function $\mathbb{E}V(x)$
- Main challenge: $V(x)$ depends on
  - transition density $f(s_{t+1}|x_t, a_t)$
  - conditional choice probability (CCP) $p(a = 1|x)$

  that are high-dimensional objects

Main challenge: bias carries over from stage 1 to stage 2

- Need modern regularized methods to estimate transition density and CCP
- Regularization bias converges slower than root-N
- Bias carries over from stage 1 to stage 2
- Biased value function estimator cannot be used for inference based on standard asymptotic theory

Main idea: orthogonality

- Adjust second-stage moment to make it insensitive with respect to first-stage bias
- Previous work: adjustment relied on the closed-form expression of a moment (Neyman (1959), Newey (1994), Chernozhukov et al. (2017a))
- Our contributions:
  - Introduce implicit orthogonalization that does not require closed form
  - Develop asymptotic theory for value function that allows for high-dimensional state space

Literature review

Two-stage models with orthogonality Neyman (1959), Newey (1994), Robins and Rotnitzky (1995), Belloni et al. (2013), Belloni et al. (2016), Chernozhukov et al. (2017a), Chernozhukov et al. (2017b), Ichimura and Newey (2018)

DDC models Rust (1987), Aguirregabiria and Mira (2002), Aguirregabiria and Mira (2007), Bajari et al. (2007), Bajari et al. (2010)

Outline

- Value function: example of average welfare
- Value function: asymptotic theory for general case

Example: intro and key facts

- Goal: conduct inference about $\mathbb{E}V(x)$ assuming agent plays optimally
  - Subgoal: remove the impact of estimation of $f(x'|x, a)$ and $p(x)$ on $V(x; p; f)$

Example: intro and key facts

- Goal: conduct inference about $\mathbb{E}V(x)$ assuming agent plays optimally
  - Subgoal: remove the impact of estimation of $f(x'|x, a)$ and $p(x)$ on $V(x; p; f)$
- Bellman equation:

$$V(x; p; f) = \mathbb{E}_\epsilon \max_{a \in \mathcal{A}} \left[u(x, a) + \epsilon(a) + \beta \mathbb{E}_{x,a}[V(x'; p; f)|x, a] \right],$$

where x (x') is the current (future) state

Example: intro and key facts

- Goal: conduct inference about $\mathbb{E}V(x)$ assuming agent plays optimally
  - Subgoal: remove the impact of estimation of $f(x'|x,a)$ and $p(x)$ on $V(x;p;f)$
- Bellman equation:

$$V(x;p;f) = \mathbb{E}_\epsilon \max_{a\in\mathcal{A}} \left[u(x,a) + \epsilon(a) + \beta\mathbb{E}_{x,a}[V(x';p;f)|x,a]\right],$$

  where x (x') is the current (future) state

- Recursive property of the value function (e.g., Aguirregabiria and Mira (2002)):

$$V(x;p;f) = \tilde{U}(x;p) + \beta\mathbb{E}_{x'|x}[V(x';p;f)|x],$$

  where $\tilde{U}(x;p)$ is current ex-ante expected utility ,e.g.

  $\tilde{U}(x;p) = -R\cdot p(0|x) - \mu\cdot s\cdot p(1|x) + \mathbb{E}[\sum_{a\in\mathcal{A}}\epsilon(a)p(a|x)|x]$

Example: intro and key facts

- Goal: conduct inference about $\mathbb{E}V(x)$ assuming agent plays optimally
  - Subgoal: remove the impact of estimation of $f(x'|x,a)$ and $p(x)$ on $V(x;p;f)$
- Bellman equation:

$$V(x;p;f) = \mathbb{E}_\epsilon \max_{a\in\mathcal{A}} \left[u(x,a) + \epsilon(a) + \beta\mathbb{E}_{x,a}[V(x';p;f)|x,a]\right],$$

  where x (x') is the current (future) state

- Recursive property of the value function (e.g., Aguirregabiria and Mira (2002)):

$$V(x;p;f) = \tilde{U}(x;p) + \beta\mathbb{E}_{x'|x}[V(x';p;f)|x],$$

  where $\tilde{U}(x;p)$ is current ex-ante expected utility ,e.g.
  $\tilde{U}(x;p) = -R\cdot p(0|x) - \mu\cdot s\cdot p(1|x) + \mathbb{E}[\sum_{a\in\mathcal{A}}\epsilon(a)p(a|x)|x]$

- Abstract away from computation of $V(x;p;f)$

Example: distribution of naive estimator

1^{st} stage Estimate $(\hat{f}(x'|x,a))_{a\in\mathcal{A}}$ and $\hat{p}(x)$ (auxiliary sample)

2^{nd} stage Compute $\frac{1}{N}\sum_{i=1}^{N} V(x_i;\hat{p};\hat{f})$ (main sample)

Bias of naive estimator in the literature:

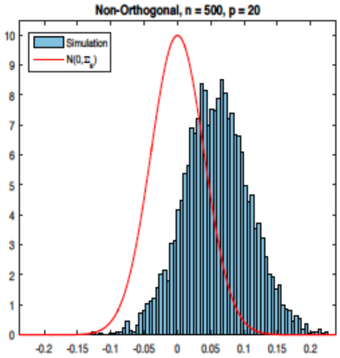

(a) Double Machine Learning (DML), Chernozhukov et al. (2017a)

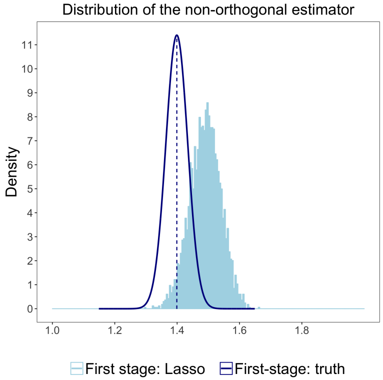

(b) DML for Set-Identified Models, Semenova (2018)

Example: bias due to CCP

- True value $f_0(x'|x, a)$ is known. Taylor expansion around p_0:

$$\mathbb{E}[V(x;p) - V(x;p_0)] \approx \partial_p \mathbb{E} V(x;p_0)[p(x) - p_0(x)]$$

Example: bias due to CCP

- True value $f_0(x'|x,a)$ is known. Taylor expansion around p_0:

$$\mathbb{E}[V(x;p) - V(x;p_0)] \approx \partial_p \mathbb{E} V(x;p_0)[p(x) - p_0(x)]$$

- Differentiate Bellman equation:

$$\begin{aligned} \partial_p V(x;p_0) &= \partial_p \mathbb{E}_\epsilon \big[\max_{a \in \mathcal{A}} u(x,a) + \epsilon(a) + \beta \mathbb{E}_{x'}[V(x')|x,a] \big] \\ &= \beta \mathbb{E}_\epsilon \partial_p [V(x';p)|x, a^*(\epsilon)], \end{aligned}$$

where $a^*(\epsilon)$ is the optimal action given x, ϵ.

Example: bias due to CCP

- True value $f_0(x'|x, a)$ is known. Taylor expansion around p_0:

$$\mathbb{E}[V(x;p) - V(x;p_0)] \approx \partial_p \mathbb{E} V(x;p_0)[p(x) - p_0(x)]$$

- Differentiate Bellman equation:

$$\partial_p V(x;p_0) = \partial_p \mathbb{E}_\epsilon \big[\max_{a \in \mathcal{A}} u(x,a) + \epsilon(a) + \beta \mathbb{E}_{x'}[V(x')|x,a] \big]$$
$$= \beta \mathbb{E}_\epsilon \partial_p [V(x';p)|x, a^*(\epsilon)],$$

  where $a^*(\epsilon)$ is the optimal action given x, ϵ.

- $\partial_p V(x;p_0)$ solves a contraction map. Unique solution $\partial_p V(x';p_0) = 0$.
- Bias of CCP has no first-order effect on the bias of value function at any state x

Example: bias due to transition density

- True value $p_0(a|x)$ is known. Let $f(x'|x) := \sum_{a \in \mathcal{A}} f(x'|x,a)p_0(a|x)$.

  $f(x'|x,\gamma)$ is a parametric submodel: $f(x'|x,\gamma_0) = f(x'|x)$

$$\mathbb{E}[V(x;\gamma) - V(x;\gamma_0)] \approx \partial_\gamma \mathbb{E} V(x;\gamma_0)[\gamma - \gamma_0]$$

Example: bias due to transition density

- True value $p_0(a|x)$ is known. Let $f(x'|x) := \sum_{a \in \mathcal{A}} f(x'|x,a)p_0(a|x)$.

  $f(x'|x,\gamma)$ is a parametric submodel: $f(x'|x,\gamma_0) = f(x'|x)$

$$\mathbb{E}[V(x;\gamma) - V(x;\gamma_0)] \approx \partial_\gamma \mathbb{E} V(x;\gamma_0)[\gamma - \gamma_0]$$

- Recall the recursive property:

$$V(x;p;\gamma) = \tilde{U}(x;p) + \beta \mathbb{E}_{x';\gamma}[V(x';p;\gamma)|x]$$

Example: bias due to transition density

- True value $p_0(a|x)$ is known. Let $f(x'|x) := \sum_{a\in\mathcal{A}} f(x'|x,a)p_0(a|x)$.
  $f(x'|x,\gamma)$ is a parametric submodel: $f(x'|x,\gamma_0) = f(x'|x)$
  $$\mathbb{E}[V(x;\gamma) - V(x;\gamma_0)] \approx \partial_\gamma \mathbb{E}V(x;\gamma_0)[\gamma - \gamma_0]$$
- Recall the recursive property:
  $$V(x;p;\gamma) = \tilde{U}(x;p) + \beta\mathbb{E}_{x';\gamma}[V(x';p;\gamma)|x]$$
- Differentiate the recursive property

$$\partial_\gamma V(x;\gamma_0) = \beta\partial_\gamma \int V(x';\gamma_0)f(x'|x;\gamma_0)dx' + \beta\mathbb{E}\partial_\gamma[V(x';\gamma_0)|x]dx'$$

Example: bias due to transition density

- True value $p_0(a|x)$ is known. Let $f(x'|x) := \sum_{a\in\mathcal{A}} f(x'|x,a)p_0(a|x)$.

  $f(x'|x,\gamma)$ is a parametric submodel: $f(x'|x,\gamma_0) = f(x'|x)$

$$\mathbb{E}[V(x;\gamma) - V(x;\gamma_0)] \approx \partial_\gamma \mathbb{E}V(x;\gamma_0)[\gamma - \gamma_0]$$

- Recall the recursive property:

$$V(x;p;\gamma) = \tilde{U}(x;p) + \beta\mathbb{E}_{x';\gamma}[V(x';p;\gamma)|x]$$

- Differentiate the recursive property

$$\partial_\gamma V(x;\gamma_0) = \beta\partial_\gamma \int V(x';\gamma_0)f(x'|x;\gamma_0)dx' + \beta\mathbb{E}\partial_\gamma[V(x';\gamma_0)|x]dx'$$

- Assuming stationarity (x and x' have same marginal dist.):

$$\partial_\gamma \mathbb{E}V(x;\gamma_0) = \frac{\beta}{1-\beta}\int V(x';\gamma_0)f(x'|x;\gamma_0)dx' = \frac{\beta}{1-\beta}\mathbb{E}V(x')S(x'|x),$$

where $S(x'|x) = \frac{\partial_\gamma f(x'|x;\gamma)}{f(x'|x;\gamma)}$ is conditional score

Example: orthogonal moment for average welfare

Recall the expression for the derivative:

$$\partial_\gamma \mathbb{E} V(x;\gamma_0) = \tfrac{\beta}{1-\beta} \mathbb{E} V(x') S(x'|x).$$

- Correction term $\frac{\beta}{1-\beta}\big(V(x';f) - \mathbb{E}_f[V(x';f)|x]\big)$
  - $V(x';f) - \mathbb{E}_f[V(x';f)|x]$ mean zero residual
  - $\frac{\beta}{1-\beta}$ multiplier, known in case of average welfare!
- Orthogonal moment function for $\mathbb{E} V(x)$:

$$g(\text{data}; p; f) := V(x;f) + \underbrace{\frac{\beta}{1-\beta}\big(V(x';f) - \mathbb{E}_f[V(x';f)|x]\big)}_{\text{correction}(f)},$$

where data $= (x, a, x')$

Example: robustness to misspecification of transition density

- Define $\Delta V(x;p) = V(x;p;f) - V(x;p;f_0)$. Recursive property implies:

$$\mathbb{E}_f[V(x';p;f)|x] = \frac{1}{\beta}(V(x;p;f) - \tilde{U}(x;p))$$
$$\mathbb{E}_f[V(x';p;f)|x] - \mathbb{E}_{f_0}[V(x';p;f_0)|x] = \frac{1}{\beta}\Delta V(x).$$

- Robustness to misspecification

$$\mathbb{E}[g(x;f) - g(x;f_0)] = \mathbb{E}[\Delta V(x) + \frac{\beta}{1-\beta}\Delta V(x') - \frac{1}{1-\beta}\Delta V(x)]$$
$$= \frac{\beta}{1-\beta}\mathbb{E}[\Delta V(x') - \Delta V(x)] = 0.$$

- Special case of linearity implication:
  - Recursive equation is linear in the density $V(x;p;f) \Rightarrow$ robustness to any order of the bias of $f(x'|x,a)$

Example: distribution of the estimator based on modified moment equation

1st stage Estimate $(\hat{f}(x'|x,a))_{a\in\mathcal{A}}$ and $\hat{p}(x)$ (auxiliary sample)

2nd stage Compute $\frac{1}{N}\sum_{i=1}^{N} g(\text{data}_i;\hat{p};\hat{f})$ (main sample)

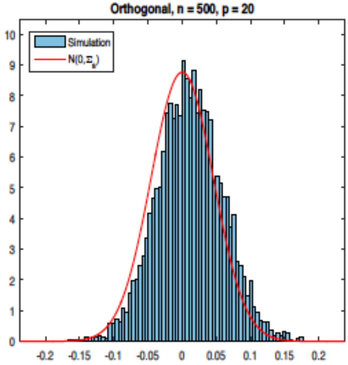

(a) Double Machine Learning (DML), Chernozhukov et al. (2017a)

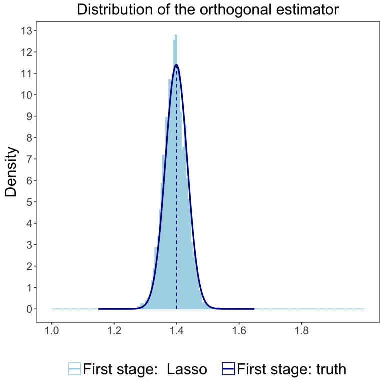

(b) DML for Set-Identified Models, Semenova (2018)

Example: summary

Goal estimate the average welfare $\mathbb{E}V(x)$ in single-agent DDC model when

- state x is high-dimensional
- $p(a|x)$ and $f(x'|x,a)$ are estimated by regularized methods (lasso, ridge, boosting, etc.)

Steps

1. $V(x;p)$ is orthogonal with respect to CCP $p(\cdot)$ at each x
2. Derived correction term for the transition density $f(x'|x)$
3. Orthogonal moment for $\mathbb{E}V(x)$ does not depend on $f(x'|x)$

Notes

- Main equation of the talk

$$V(x;p;\gamma) = \tilde{U}(x;p) + \beta\mathbb{E}_{x';\gamma}[V(x';p;\gamma)|x]$$

Outline

- Example
- Asymptotic theory for general case

General case: motivating examples

Goal is to estimate and provide confidence interval for $\mathbb{E}w(x)V(x)$.

Examples:

- Truncated average welfare:

$$\mathbb{E}V(x)1_{\{q_{0.25}\leq V(x)\leq q_{0.75}\}}$$

  and $\mathbb{E}w(x)V(x)$ is the linear approximation of the functional above

- Average partial effect w.r.t $x_1 \subset x$

$$\mathbb{E}\partial_{x_1} V(x) = \mathbb{E}w(x)V(x),$$

  where $w(x) = -\frac{\partial_{x_1} f(x_1|x_{-1})}{f(x_1|x_{-1})}$ and $f(x_1|x_{-1})$ is the conditional density of x_1 given x_{-1}

- Counterfactual policy effect (Stock (1989)) from changing the distribution of x to $t(x)$:

$$\mathbb{E}[V(t(x)) - V(x)] = \mathbb{E}w(x)V(x),$$

  where $w(x) = \frac{f_t(x)}{f_x(x)} - 1$, $f_t(x), f_x(x)$ are the marginal densities of $t(x), x$

General case: comparison with average welfare

- CCP has no first-order effect on $\mathbb{E}w(x)V(x)$.

General case: comparison with average welfare

- CCP has no first-order effect on $\mathbb{E}w(x)V(x)$.
- Correction for the transition density takes the form

$$\text{correction} = \beta\lambda(x)(V(x') - \mathbb{E}[V(x')|x]),$$

  - residual $V(x') - \mathbb{E}[V(x')|x]$ same as before
  - multipler $\beta\lambda(x) = \beta\sum_{k=0}^{\infty}\mathbb{E}\beta^k[w(x_{-k})|x]$
  - Recap for $w(x) = 1$: $\lambda(x) = \frac{\beta}{1-\beta}$

General case: comparison with average welfare

- CCP has no first-order effect on $\mathbb{E}w(x)V(x)$.
- Correction for the transition density takes the form
$$\text{correction} = \beta\lambda(x)(V(x') - \mathbb{E}[V(x')|x]),$$
  - residual $V(x') - \mathbb{E}[V(x')|x]$ same as before
  - multipler $\beta\lambda(x) = \beta\sum_{k=0}^{\infty}\mathbb{E}\beta^k[w(x_{-k})|x]$
  - Recap for $w(x) = 1$: $\lambda(x) = \frac{\beta}{1-\beta}$
- Double robustness requirement
$$\|\hat{\lambda}(x) - \lambda(x)\|_{2,N}\|\hat{f}(x'|x) - f(x'|x)\|_{2,N} = o(N^{-1/2})$$

General case: comparison with average welfare

- CCP has no first-order effect on $\mathbb{E}w(x)V(x)$.
- Correction for the transition density takes the form

$$\text{correction} = \beta\lambda(x)(V(x') - \mathbb{E}[V(x')|x]),$$

  - residual $V(x') - \mathbb{E}[V(x')|x]$ same as before
  - multipler $\beta\lambda(x) = \beta\sum_{k=0}^{\infty}\mathbb{E}\beta^k[w(x_{-k})|x]$
  - Recap for $w(x) = 1$: $\lambda(x) = \frac{\beta}{1-\beta}$
- Double robustness requirement

$$\|\hat{\lambda}(x) - \lambda(x)\|_{2,N}\|\hat{f}(x'|x) - f(x'|x)\|_{2,N} = o(N^{-1/2})$$

General case: relaxing stationarity

- Stationarity: marginal distribution of $(x, x'', \ldots, x^k, \ldots)$ and $(x^t, x^{t+1}, \ldots, x^{t+k}, \ldots)$ are the same. Example: mileage in Rust (1987)

General case: relaxing stationarity

- Stationarity: marginal distribution of $(x, x'', \ldots, x^k, \ldots)$ and $(x^t, x^{t+1}, \ldots, x^{t+k}, \ldots)$ are the same. Example: mileage in Rust (1987)
- Without stationarity:
  - data $= (a, x)_{t=1}^{\infty}$
  - correction $\alpha(\text{data}) = \beta \sum_{k \geq 0} \beta^k \left(V(x^{k+1}) - \mathbb{E}_f[V(x^{k+1})|x^k] \right)$
- With stationarity:
  - data $= (x, a, x')$
  - correction $\alpha(\text{data}) = \beta \frac{1}{1-\beta} \left(V(x') - \mathbb{E}_f[V(x')|x] \right)$
- Stationarity is strong; not applicable for suboptimal policies and Bajari et al. (2007) type problems

General case: weighted bootstrap

Goal is to provide confidence interval for $\mathbb{E}w(x)V(x)$. Weighted bootstrap statistic

$$\tilde{\sigma}^b = \frac{1}{N}\sum_{i=1}^N e_i g(\text{data}_i, \hat{\xi}_i, \tilde{\theta}^b)$$

- $\hat{\xi} = \{\hat{f}(x'|x,a), \hat{\lambda}(x), \hat{p}(x), \hat{\theta}\}$ is estimated on auxiliary sample
- $(e_i)_{i=1}^N$ is Exp(1) i.i.d weights drawn independently from data

General case: asymptotic results

Theorem 1. Asymptotic theory for the welfare estimator, known density

Suppose $(f_0(x'|x,a))_{a\in\mathcal{A}}$ is known and $\|\hat{p}-p\|_{N,2}=o(N^{-1/4})$. Under mild conditions,

$$\sqrt{N}\frac{1}{N}\sum_{i=1}^{N} V(\text{data}_i;\hat{p}) \approx \sqrt{N}\frac{1}{N}\sum_{i=1}^{N} V(\text{data}_i;p_0) + o_P(1)$$

Theorem 2. Asymptotic theory for the welfare estimator, general case

Suppose $\|\hat{p}-p\|_{N,2}=o(N^{-1/4})$ and $\|\hat{\lambda}-\lambda\|_{N,2}\|\hat{f}(x'|x)-f(x'|x)\|_{N,2}=o(1)$.

$$\sqrt{N}\frac{1}{N}\sum_{i=1}^{N} g(\text{data}_i;\hat{\xi}) \approx \sqrt{N}\frac{1}{N}\sum_{i=1}^{N} g(\text{data}_i;\xi_0) + o_P(1),$$

where $\xi=\{p(x),f(x'|x),\lambda(x)\}$

Notes

- std. case: directly assume bounded second-order derivatives
- this paper: derive asymptotic theory from the properties of linear integral equations

General case: summary

Goal estimate and conduct inference about $\mathbb{E}w(x)V(x)$

Key points

- $\mathbb{E}w(x)V(x)$ is orthogonal to CCP
- Correction term changes from $\beta(V(x') - \mathbb{E}[V(x')|x]$ to $\beta\lambda(x)(V(x') - \mathbb{E}[V(x')|x]$
- Robustness to any order misspecification turns into double robustness $\|\hat{\lambda}(x) - \lambda(x)\|_{2,N}\|\hat{f}(x'|x,a) - f(x'|x,a)\|_{2,N} = o(N^{-1/2})$ in expectation
- Relaxed stationarity

Results

- Derived sufficient conditions for asymptotic theory
- Provided weighted bootstrap procedure for inference

Outline

- Value function: example of average welfare
- Value function: asymptotic theory for general case

Aguirregabiria, V. and Mira, P. (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models. *Econometrica*, 70(4):1519–1543.

Aguirregabiria, V. and Mira, P. (2007). Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

Bajari, P., Hong, H., Krainer, J., and Nekipelov, D. (2010). Estimating static models of strategic interaction. *Journal of Business and Economic Statistics*, 28(4):469–482.

Belloni, A., Chernozhukov, V., Fernandez-Val, I., and Hansen, C. (2013). Program evaluation and causal inference with high-dimensional data. *arXiv preprint arXiv:1311.2645*.

Belloni, A., Chernozhukov, V., and Wei, Y. (2016). Post-selection inference for generalized linear models with many controls. *Journal of Business & Economic Statistics*, 34(4):606–619.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2017a). Double/debiased machine learning for treatment and causal parameters.

Chernozhukov, V., Escanciano, J. C., Ichimura, H., and Newey, W. (2017b). Locally robust semiparametric estimation.

Ichimura, H. and Newey, W. (2018). The influence function of semiparametric estimators.

Newey, W. (1994). The asymptotic variance of semiparametric estimators. 62:245–271.

Neyman, J. (1959). Optimal asymptotic tests of composite statistical hypotheses. *Probability and Statistics*, 213(57).

Robins, J. and Rotnitzky, A. (1995). Semiparametric efficiency in multivariate regression models with missing data. *Journal of American Statistical Association*, 90(429):122–129.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033.

Semenova (2018). Machine learning for set-identified linear models. *https://arxiv.org/abs/1712.10024*.

Stock, J. (1989). Nonparametric policy analysis. *Journal of American Statistical Association*, 84:567–575.