Regret Analysis in Threshold Policy Design

Federico Crippa (Northwestern)
discussed by Toru Kitagawa (Brown)

This paper

- This paper is about **treatment choice/policy learning**
- A growing and active research area in the current econometrics, with many potential applications

This paper

- This paper is about **treatment choice/policy learning**
- A growing and active research area in the current econometrics, with many potential applications

- **Who should be treated?**: Want to learn desirable targeting (treatment assignment) over heterogeneous individuals. Manski (04 Ecta), Kitagawa and Tetenov (18 Ecta), Mbakop and Tabord-Meehan (21, Ecta), Athey & Wager, (21 Ecta)

This paper

- This paper is about **treatment choice/policy learning**
- A growing and active research area in the current econometrics, with many potential applications

- **Who should be treated?**: Want to learn desirable targeting (treatment assignment) over heterogeneous individuals. Manski (04 Ecta), Kitagawa and Tetenov (18 Ecta), Mbakop and Tabord-Meehan (21, Ecta), Athey & Wager, (21 Ecta)

- Targeting: the central policy question in various contexts.
    - Who should get active labor programs?
    - Who should get conditional cash transfers?
    - Who should be tax audited?
    - Who should be health-insured?

Outline of my talk

1. Setting and review of EWM
2. Proposal: Smooth Welfare Maximization (SWM)
3. Main takeaways
4. Comments and Questions

Statistical Treatment Choice

Suppose we have a Randomized Controlled Trial (RCT) sample with

- $X_i \in \mathcal{X}$ - pre-treatment observed covariates (past earnings, education, etc.)
- $D_i \in \{0, 1\}$ - randomized treatment (job training program)
- $Z_i \in \mathbb{R}$ - post-treatment observed outcome (earnings)

Statistical Treatment Choice

Suppose we have a Randomized Controlled Trial (RCT) sample with

- $X_i \in \mathcal{X}$ - pre-treatment observed covariates (past earnings, education, etc.)
- $D_i \in \{0,1\}$ - randomized treatment (job training program)
- $Z_i \in \mathbb{R}$ - post-treatment observed outcome (earnings)

- An assignment rule π is a map $\mathcal{X} \to \{0,1\}$.
- Let (Y_1, Y_0) be the potential outcomes. Additive welfare:

$$W(G) = E\left[Y_1 \cdot \pi(X) + Y_0 \cdot (1 - \pi(X))\right].$$

- **Goal:** Using the data (already collected), obtain statistical treatment rule (policy) $\hat{\pi}$ that performs well in terms of $W(\hat{\pi})$.

Empirical Welfare Maximization

EWM approach (Kitagawa and Tetenov 18 Ecta): maximize an empirical welfare,

$$\hat{\pi}_{EWM} = arg \max_{\pi \in \Pi} \widehat{W}(\pi),$$

where Π is a (constrained) class of feasible policies (partitions of $\mathcal{X}$).

Empirical Welfare Maximization

EWM approach (Kitagawa and Tetenov 18 Ecta): maximize an empirical welfare,

$$\hat{\pi}_{EWM} = arg \max_{\pi \in \Pi} \widehat{W}(\pi),$$

where Π is a (constrained) class of feasible policies (partitions of $\mathcal{X}$).

$$\widehat{W}(\pi) = \frac{1}{n} \sum_{i=1}^{n} \left[\frac{Y_i D_i}{e(X_i)} \cdot \pi(X_i) + \frac{Y_i(1-D_i)}{1-e(X_i)} \cdot (1-\pi(X_i)) \right]$$

where $e(x) = P(D = 1 \mid X = x)$ is the (RCT) propensity score.

Why constrained class of $\mathcal{G}$?

- **Policy-makers in reality face many constraints**:
- Has to be simple enough, e.g., single-index, decision tree, etc.
- Redistributive concerns, e.g., assignment of a social program has to be monotonic in one's wealth

Why constrained class of $\mathcal{G}$?

- **Policy-makers in reality face many constraints**:
- Has to be simple enough, e.g., single-index, decision tree, etc.
- Redistributive concerns, e.g., assignment of a social program has to be monotonic in one's wealth
- This paper: a **threshold assignment policy** given a single index $g(X) \in \mathbb{R}$.

$$\pi(X) = 1\{g(X) > t\}$$

Why constrained class of $\mathcal{G}$?

- **Policy-makers in reality face many constraints**:
- Has to be simple enough, e.g., single-index, decision tree, etc.
- Redistributive concerns, e.g., assignment of a social program has to be monotonic in one's wealth
- This paper: a **threshold assignment policy** given a single index $g(X) \in \mathbb{R}$.

$$\pi(X) = 1\{g(X) > t\}$$

- Present in many contexts:
  - Clinical guidelines
  - Public healthcare insurance
  - Anti-poverty program
  - Sanctions on driving based on blood alcohol content.

Advantage/Drawback of EWM approach

- Desirable statistical properties: Under minimal conditions, **Average welfare regret** converges to zero at a minimax optimal rate: for a universal constant $C > 0$,

$$E\left[\sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}_{EWM})\right] \leq C\sqrt{\frac{\text{complexity}(\Pi)}{n}}$$

- Regret $R(\pi) = \sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}_{EWM})$ measures the loss of welfare

Advantage/Drawback of EWM approach

- Desirable statistical properties: Under minimal conditions, **Average welfare regret** converges to zero at a minimax optimal rate: for a universal constant $C > 0$,

$$E\left[\sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}_{EWM})\right] \leq C\sqrt{\frac{\text{complexity}(\Pi)}{n}}$$

- Regret $R(\pi) = \sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}_{EWM})$ measures the loss of welfare
- Drawback of EWM: computationally demanding, similar to Manski's maximum score
- Threshold policies involve only one parameter, so computation is not an issue here

Crippa (2024)

Proposal of a Smooth Welfare Maximization (SWM):

- Replace the indicator function in

$$\widehat{W}(t) = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{Y_i D_i}{e(X_i)}\cdot 1\{g(X_i) > t\} + \frac{Y_i(1-D_i)}{1-e(X_i)}\cdot 1\{g(X_i) \leq t\}\right]$$

$$= \frac{1}{n}\sum_{i=1}^{n}\left[\left(\frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}\right)\cdot 1\{g(X_i) \leq t\}\right] + \frac{1}{n}\sum_{i=1}^{n}\frac{Y_i(1-D_i)}{1-e(X_i)},$$

with **smooth cdf** (e.g., logit), $\Lambda\left(\frac{f(X_i)-t}{\sigma_n}\right)$ with bandwidth, $\sigma_n \to 0$

Crippa (2024)

Proposal of a Smooth Welfare Maximization (SWM):

- Replace the indicator function in

$$\widehat{W}(t) = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{Y_i D_i}{e(X_i)}\cdot 1\{g(X_i) > t\} + \frac{Y_i(1-D_i)}{1-e(X_i)}\cdot 1\{g(X_i)\leq t\}\right]$$
$$= \frac{1}{n}\sum_{i=1}^{n}\left[\left(\frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}\right)\cdot 1\{g(X_i)\leq t\}\right] + \frac{1}{n}\sum_{i=1}^{n}\frac{Y_i(1-D_i)}{1-e(X_i)},$$

  with **smooth cdf** (e.g., logit), $\Lambda\left(\frac{f(X_i)-t}{\sigma_n}\right)$ with bandwidth, $\sigma_n \to 0$
- This idea is a reminiscent of smooth maximum score of Horowitz (92, Ecta)

Crippa (2024)

Proposal of a Smooth Welfare Maximization (SWM):

- Replace the indicator function in

$$\widehat{W}(t) = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{Y_i D_i}{e(X_i)}\cdot 1\{g(X_i) > t\} + \frac{Y_i(1-D_i)}{1-e(X_i)}\cdot 1\{g(X_i) \le t\}\right]$$

$$= \frac{1}{n}\sum_{i=1}^{n}\left[\left(\frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}\right)\cdot 1\{g(X_i) \le t\}\right] + \frac{1}{n}\sum_{i=1}^{n}\frac{Y_i(1-D_i)}{1-e(X_i)},$$

  with **smooth cdf** (e.g., logit), $\Lambda\left(\frac{f(X_i)-t}{\sigma_n}\right)$ with bandwidth, $\sigma_n \to 0$

- This idea is a reminiscent of smooth maximum score of Horowitz (92, Ecta)
- **Is smoothing $\widehat{W}(t)$ good?**

Crippa (2024)

Proposal of a Smooth Welfare Maximization (SWM):

- Replace the indicator function in

$$\widehat{W}(t) = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{Y_iD_i}{e(X_i)}\cdot 1\{g(X_i) > t\} + \frac{Y_i(1-D_i)}{1-e(X_i)}\cdot 1\{g(X_i) \le t\}\right]$$
$$= \frac{1}{n}\sum_{i=1}^{n}\left[\left(\frac{Y_iD_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}\right)\cdot 1\{g(X_i) \le t\}\right] + \frac{1}{n}\sum_{i=1}^{n}\frac{Y_i(1-D_i)}{1-e(X_i)},$$

  with **smooth cdf** (e.g., logit), $\Lambda\left(\frac{f(X_i)-t}{\sigma_n}\right)$ with bandwidth, $\sigma_n \to 0$
- This idea is a reminiscent of smooth maximum score of Horowitz (92, Ecta)
- **Is smoothing $\widehat{W}(t)$ good?**
- **Yes**, if the underlying $W(t)$ is smooth!

Main results:

Theorems 2 & 4 in the paper

Key assumptions:

- $W(t)$ is maximized uniquely at t^*
- $W(t)$ is $s = h + 1$ times differentiable
- $\sigma_n = cn^{-\frac{1}{2h+1}}$

Then, the pointwise asymptotic distributions of the regret $R(\hat{t}) = W(t^*) - W(\hat{t})$ are

$$n^{1/3}(\hat{t} - t^*) = O_p(1),$$
$$n^{2/3}R(\hat{t}_{EWM}) = O_p(1),$$
$$\sqrt{n\sigma_n}(\hat{t} - t^*) = O_p(1),$$
$$n\sigma_n R(\hat{t}_{SWM}) = O_p(1).$$

Main results:

Theorems 2 & 4 in the paper

Key assumptions:

- $W(t)$ is maximized uniquely at t^*
- $W(t)$ is $s = h + 1$ times differentiable
- $\sigma_n = cn^{-\frac{1}{2h+1}}$

Then, the pointwise asymptotic distributions of the regret $R(\hat{t}) = W(t^*) - W(\hat{t})$ are

$$
\begin{aligned}
n^{1/3}(\hat{t} - t^*) &= O_p(1), \\
n^{2/3}R(\hat{t}_{EWM}) &= O_p(1), \\
\sqrt{n\sigma_n}(\hat{t} - t^*) &= O_p(1), \\
n\sigma_n R(\hat{t}_{SWM}) &= O_p(1).
\end{aligned}
$$

- **When $s \geq 3$, SWM has a faster regret convergence rate than EWM.**
- Monte Carlo shows the finite sample regret gain (15 - 80%) of SWM

Empirical Illustration

National Job Training Partnership Act (JTPA) Study (Bloom et al, 97)

Non-parametric estimation of CATE: bivariate kernel reg of $Y_1|X$ and $Y_0|X$ (ROT bandwidth).

Cost of treatment $774 per assignee

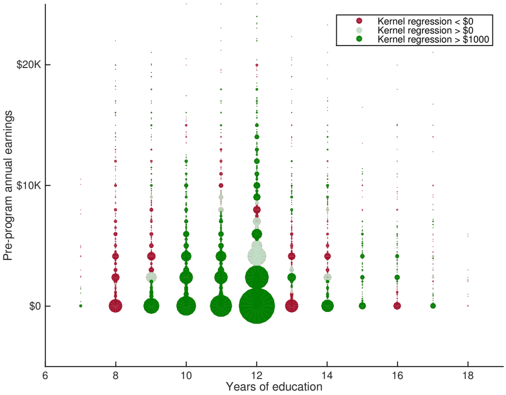

Empirical Illustration in the paper

Same as in Kitagawa and Tetenov (2018).

Data from the National Job Training Partnership Act Study.

Treatment: job training.

Outcome: individual earnings in the 30-month period after treatment.

Index: earnings in the year prior to the assignment.

$\hat{t}_{EWM} = 6614$ $[5334, 7893]$

$\hat{t}_{SWM} = 5924$ $[5832.7, 6411.3]$

Empirical Illustration

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

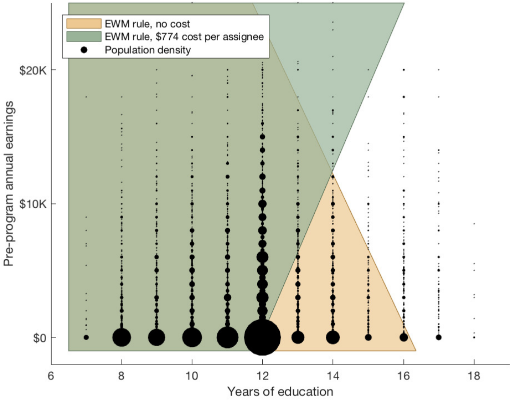

My Remarks

- Assuming smoothness of $W(t)$ is not restrictive for economics applications

My Remarks

- Assuming smoothness of $W(t)$ is not restrictive for economics applications
- However, it would be difficult to argue how many times differentiable. In theory, an optimal choice of σ_n depends on what to assume about s.

My Remarks

- Assuming smoothness of $W(t)$ is not restrictive for economics applications
- However, it would be difficult to argue how many times differentiable. In theory, an optimal choice of σ_n depends on what to assume about s.
- The shown asymptotic distributions can be useful for performing inference for t^* and $W(t^*)$, with some bias correction.

My Comments

Conceptual/practical ones:

1. **Estimating an optimal index and threshold together?** Analogy to smooth maximum score estimation, should be feasible, but not sure smoothing helps computation?

My Comments

Conceptual/practical ones:

1. **Estimating an optimal index and threshold together?** Analogy to smooth maximum score estimation, should be feasible, but not sure smoothing helps computation?
2. **Inference in treatment choice?:** Asymptotic distributions for threshold $\hat{t}$ and regret are useful for drawing inferences for them, but is inference important for policy learning?

My Comments

Conceptual/practical ones:

1. **Estimating an optimal index and threshold together?** Analogy to smooth maximum score estimation, should be feasible, but not sure smoothing helps computation?
2. **Inference in treatment choice?:** Asymptotic distributions for threshold $\hat{t}$ and regret are useful for drawing inferences for them, but is inference important for policy learning?
3. **Debiasing in regret?** The asymptotic distribution of regret contains bias. How to handle the bias for estimation and inference for $W(t^*)$?

My Comments

Technical ones:

4. **Rate improvement by margin condition** Classification/treatment choice literature discussed faster regret convergence by controlling the size of population with $CATE(X) = 0$. Comparison of the current smoothness assumption with margin condition?
5. **Relaxing the uniqueness condition of t^*?** If maximizing a welfare criterion is the first order problem, consistency of $\hat{t}$ is less important.
6. **When does smoothing harm if any?** If the order of smoothness is misspecified, what happens? E.g., the true objective function is discontinuous at t^*, smoothing can slow down the regret convergence?
7. **Adaptive methods?**: The EWM treatment choice is adaptive to the margin conditions. Any method that can be adaptive to the smoothness of $W(t)$?

Conclusion

- EWM uses only the empirical distribution and it is a tuning-parameter free approach with good robustness, but leaves potential welfare gain on the table if $W(t)$ is smooth.
- As novel contributions to the literature, SWM exploits it to improve the welfare performance, and I think it is sensible for economics applications
- However, choice of tuning parameter is a delicate problem. Ideal to be adapative to s

Conclusion

- EWM uses only the empirical distribution and it is a tuning-parameter free approach with good robustness, but leaves potential welfare gain on the table if $W(t)$ is smooth.
- As novel contributions to the literature, SWM exploits it to improve the welfare performance, and I think it is sensible for economics applications
- However, choice of tuning parameter is a delicate problem. Ideal to be adapative to s
- **Good work, I enjoyed reading!**