Estimating Heterogeneous Treatment Effects by Combining Weak Instruments and Observational Data

Miruna Oprescu, Nathan Kallus

Cornell University, Cornell Tech

Motivation

Online platforms, mobile health, targeted advertising, etc.:

- An abundance of (potentially confounded) observational data.
- Limited experimentation capabilities: can recommend an action/treatment but cannot enforce it due to ethical or logistical constraints.

Observational Data

- **Goal:** Estimate the conditional average treatment effect (CATE):

$$\tau(x) = \mathbb{E}[\,Y(1) \mid X = x\,] - \mathbb{E}[\,Y(0) \mid X = x\,]$$

- **Observed Data:**

$$O = \left(X_i^O, A_i^O, Y_i^O\right)_{i=1}^{n_O} \sim (X^O, A^O, Y^O(A)).$$

- **Challenge:** Without ignorability, the effect we can learn from data is biased:

$$\tau^O(x) = \mathbb{E}[\,Y \mid X = x, A = 1\,] - \mathbb{E}[\,Y \mid X = x, A = 0\,]$$

- **Bias Term:**

$$\underbrace{b(x)}_{\text{persistent bias as } n_O \to \infty} = \tau(x) - \tau^O(x)$$

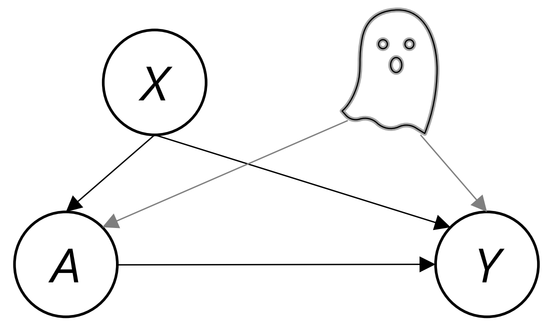

Causal graph

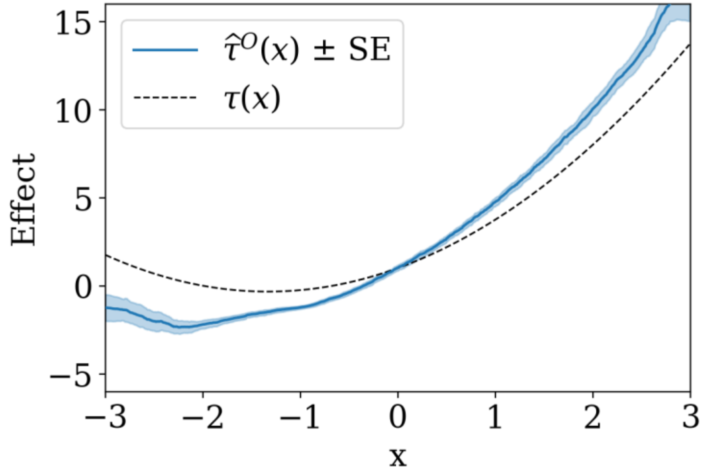

Confounded data example (Kallus et al., 2018)

Instrumental Variables (IV) Data

- **Experimental (IV) Data:**

$$E = \left(X_i^E, Z_i^E, A_i^E, Y_i^E\right)_{i=1}^{n_E} \sim \left(X^E, Z^E, A^E, Y^E(A)\right),$$

  where $Z \in \{0,1\}$ is an instrumental variable.

- **CATE Identification:**

$$\tau^E(x) = \frac{\mathbb{E}[Y \mid X = x, Z = 1] - \mathbb{E}[Y \mid X = x, Z = 0]}{\underbrace{\mathbb{E}[A \mid X = x, Z = 1] - \mathbb{E}[A \mid X = x, Z = 0]}_{\gamma(x)\ \text{(compliance factor)}}}$$

$\tau^E(x) = \tau(x)$ when $\gamma(x) \neq 0$ (+ IV assumptions).

- **Challenge:**
  - Effect is not identifiable for x when $\gamma(x) = 0$.
  - Small $\gamma(x)$ leads to high variance estimates.

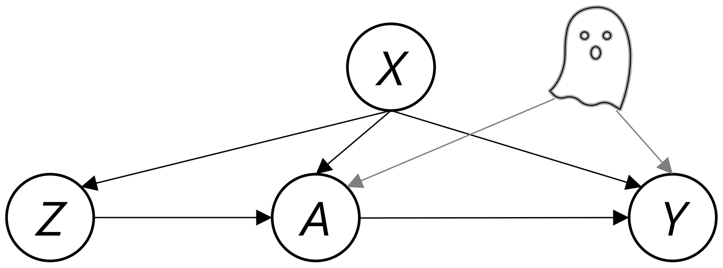

Causal graph

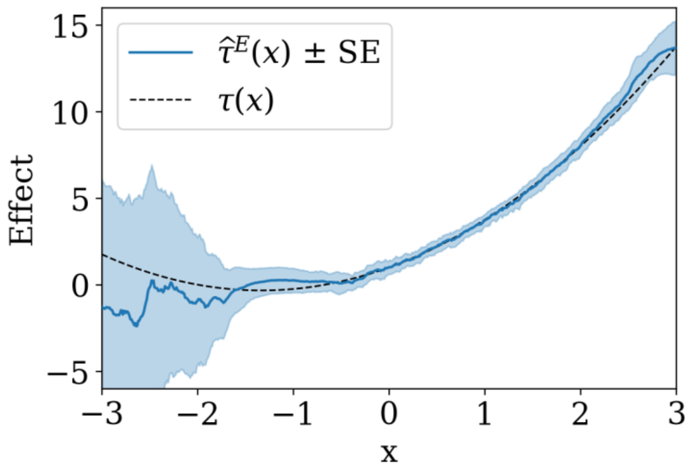

Confounded data example with IV
(based on Kallus et al., 2018)

Recap

- Relying solely on observational data results in biased estimates of $\tau(x)$.
- Using experimental (IV) data alone can yield high variance or even invalid estimates of $\tau(x)$ when the compliance factor $\gamma(x)$ is low.

Question:

Can we strategically combine the complementary strengths of both datasets to create a robust CATE estimation method for the target population?

Recap

- Relying solely on observational data results in biased estimates of $\tau(x)$.
- Using experimental (IV) data alone can yield high variance or even invalid estimates of $\tau(x)$ when the compliance factor $\gamma(x)$ is low.

This Work:

We propose a **two-stage framework** that first estimates biased CATEs from observational data and then corrects them using compliance-weighted IV samples.

Two-Stage CATE Estimation Method

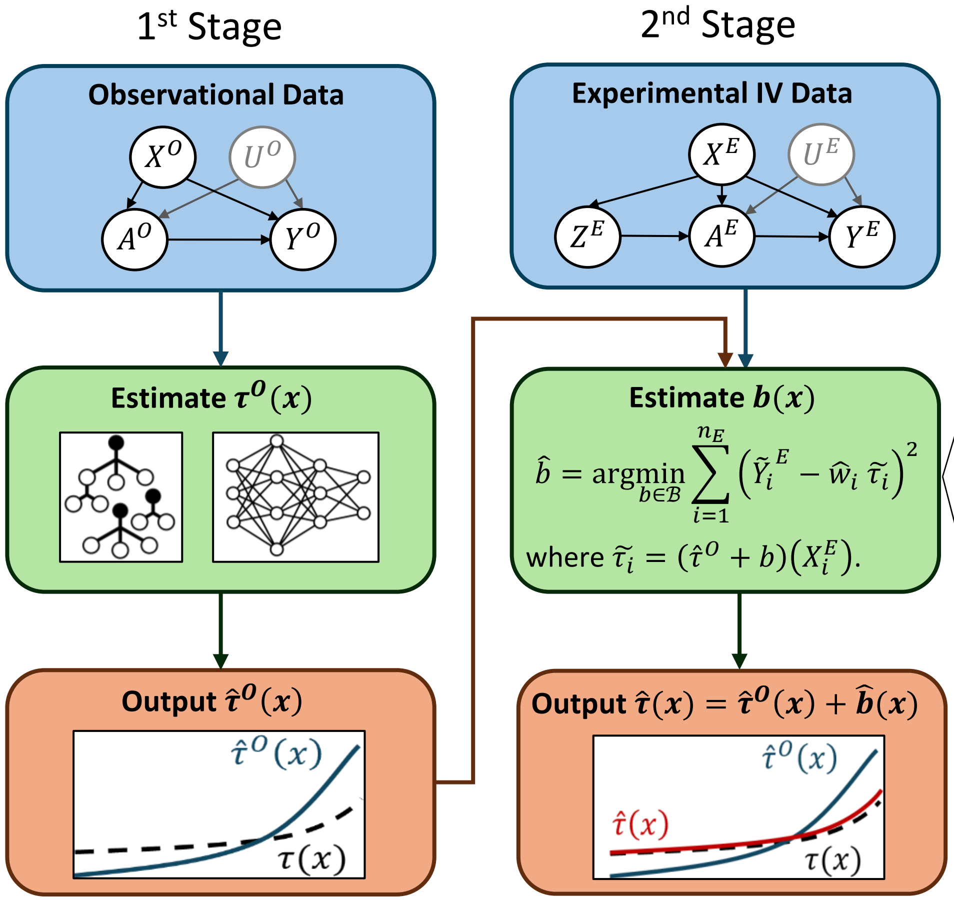

- Let:
  - $\pi_Z(x) = P(Z = 1 \mid X = x)$
  - $w(x) = \gamma(x)(1 - \pi_Z(x))\pi_Z(\mathrm{x})$
  - $\tilde{Y}^E = Y^E Z^E \big(1 - \pi_Z(X)\big) - Y^E (1 - Z^E) \pi_Z(X)$
- Note:

$$\mathbb{E}\left[\tilde{Y}^E - w(x)\big(b(x) + \tau^O(x)\big) \,\middle|\, X^E = x\right] = 0$$

for any x, regardless of the value of $\gamma(x)$.

- Learn $\hat{\pi}_Z(x)$, $\hat{\gamma}(x)$ and let:

$$\hat{b} = \operatorname*{argmin}_{b \in \mathcal{B}} \sum_{i=1}^{n_E} \left(\tilde{Y}_i^E - \hat{w}_i \, \tilde{\tau}_i\right)^2$$

where $\tilde{\tau}_i = (\hat{\tau}^O + b)\big(X_i^E\big)$

Extrapolating The Confounding Bias

To ensure generalizability, the bias class $\mathcal{B}$ must have low complexity. We consider two approaches:

1. **Parametric extrapolation:**

$$\mathcal{B} = \{\theta^T \phi(x) : \theta^T \in \mathbb{R}^d\} \text{ for a } \textit{known} \text{ mapping } \phi : \mathcal{X} \to \mathbb{R}^{\mathrm{d}} .$$

2. **Transfer learning via a common representation:**

$$\mathcal{B} = \{\theta^T \phi(x) : \theta^T \in \mathbb{R}^d\} \text{ for a } \textit{learned} \text{ mapping } \phi : \mathcal{X} \to \mathbb{R}^{\mathrm{d}}$$

   where ϕ is a shared representation between the CATEs and the bias functions:

$$\tau(x) = h^T \phi(x), \qquad \tau^O(x) = (h^O)^T \phi(x), \qquad h, h^O \in \mathbb{R}^d$$

We provide strong theoretical guarantees for both approaches.

Simulation Results: Kallus et al., 2018 DGP

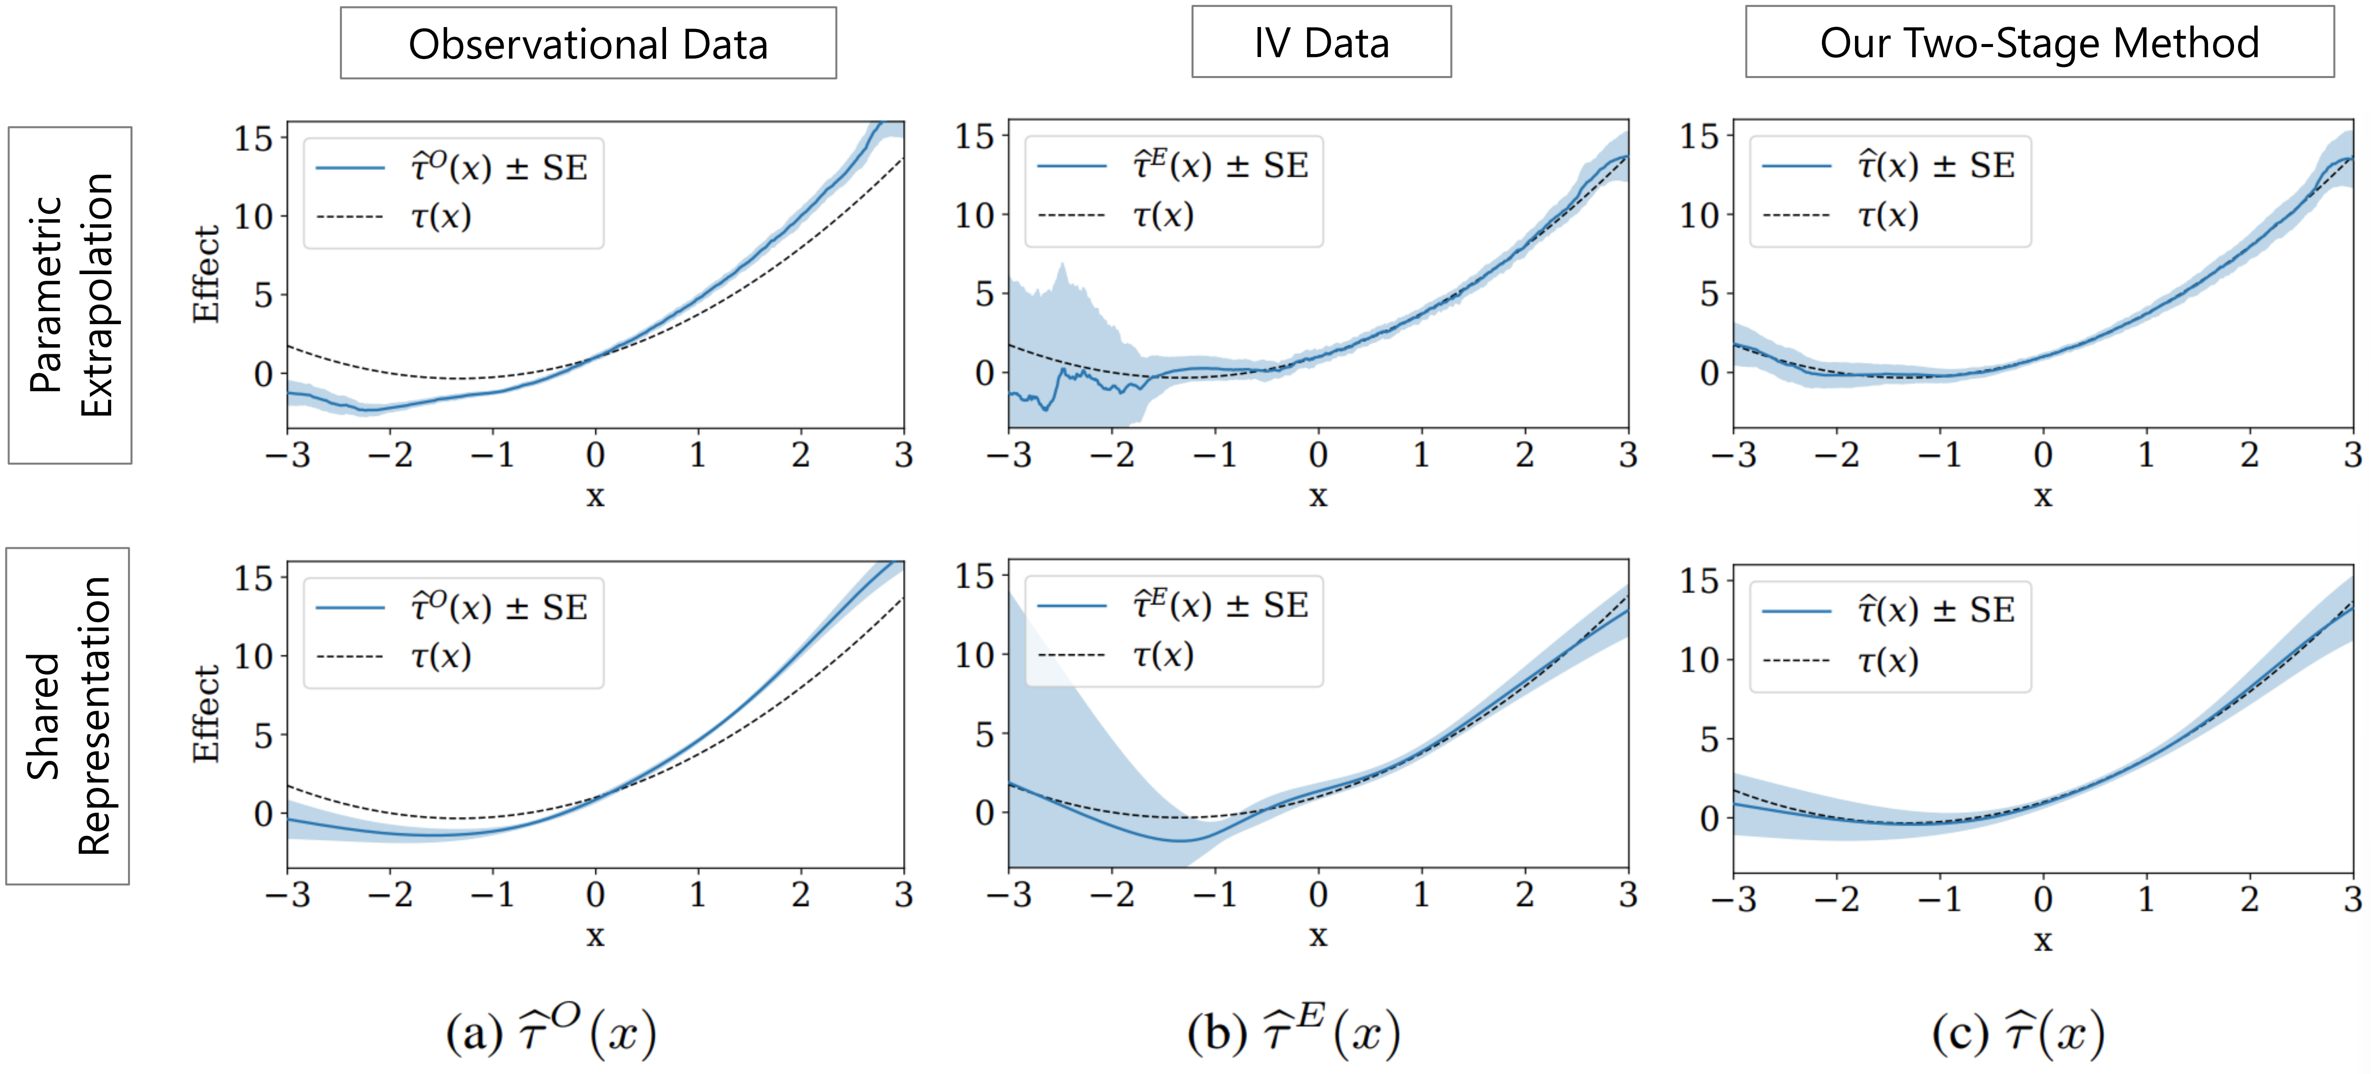

Real World Results: Effect of 401(k) on Wealth

- We study the impact of 401(k) participation on net worth by education level. 401(k) eligibility is an instrumental variable.
- We introduce non-compliance for the ≤12 years of education population and we use parametric extrapolation to estimate the CATE for this group.
- We validate against $\hat{\tau}^E$ from the high-compliance IV dataset.

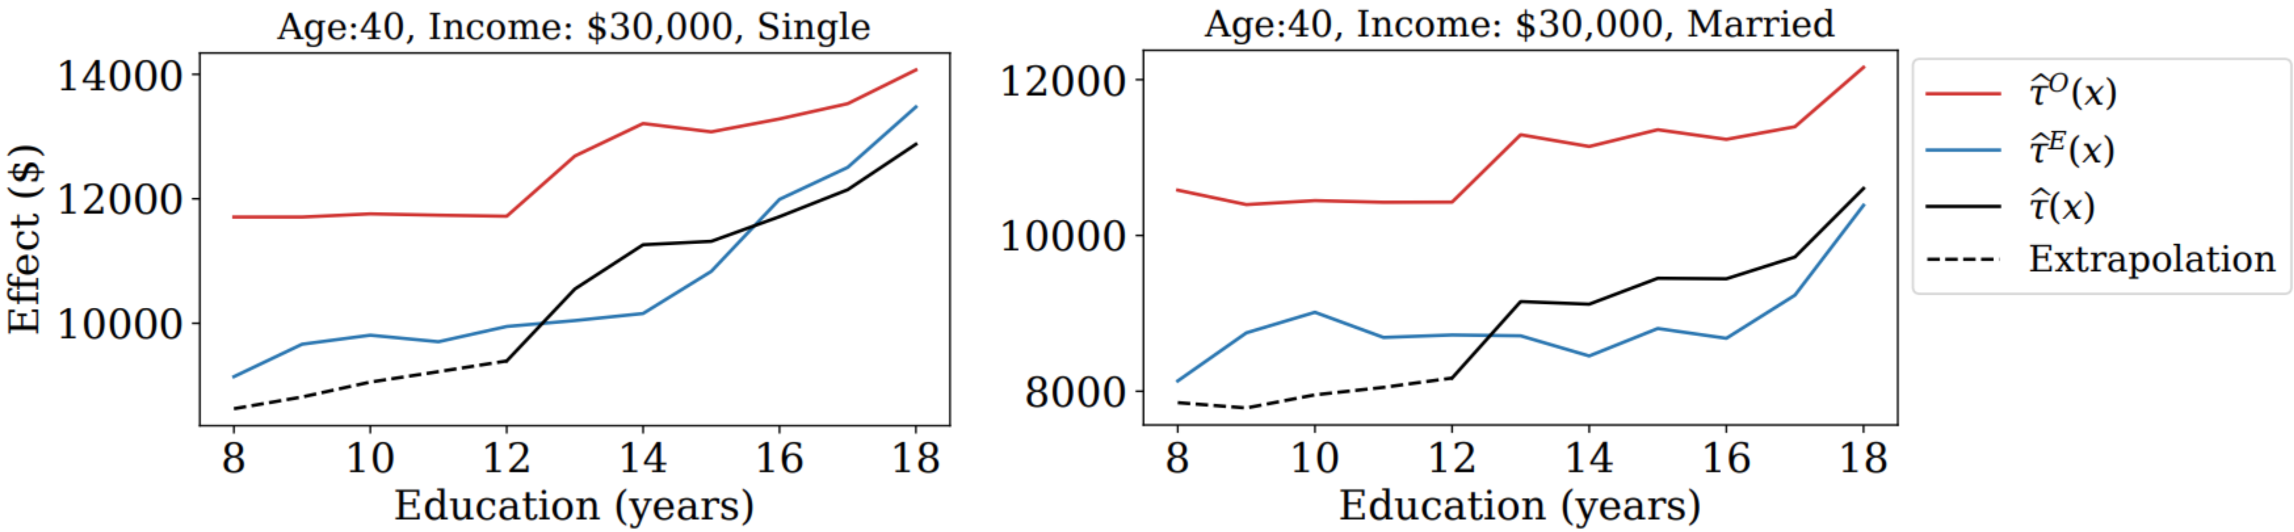

Summary of Contributions and Impact

Key Contributions:

- Introduced a **two-stage framework** combining observational and IV data to address **unobserved confounders** and **low IV compliance**.
- Two variation of our framework:
    1. **Parametric extrapolation** of the confounding bias.
    2. **Transfer learning** leveraging shared representations.
- Supported by strong theoretical guarantees for consistency.
- Validated through simulation and and real-world applications.

Broader Impact:

- Delivers robust insights for **digital platforms**, **personalized medicine**, **economics**, and beyond.