University of Warsaw | Faculty of Economic Sciences

Working Papers
No. 7/2026 (501)

Causal Inference under Algorithmic Interference: Identification and Estimation without SUTVA in Platform Economies

Jakub Ryłow

Warsaw 2026

ISSN 2957-0506

WORKING PAPERS 7/2026 (501)

UNIVERSITY OF WARSAW
FACULTY OF ECONOMIC SCIENCES
WORKING PAPERS

Causal Inference under Algorithmic Interference: Identification and Estimation without SUTVA in Platform Economies

Jakub Ryłow[1*]

[1] *Faculty of Economic Sciences, University of Warsaw*
[*] *Corresponding author: jakub.jan.rylow@gmail.com*

Abstract: The Stable Unit Treatment Value Assumption (SUTVA) fails systematically in platform economies where a deterministic algorithm mediates all interactions, making interference structural and mechanistically knowable. We introduce algorithmic interference — a structural potential-outcomes model in which spillovers flow through the platform's known decision rule — and construct the Debiased Algorithmic Instrumental Variable (DAIV) estimator: a cross-fitted semiparametric procedure combining Double Machine Learning with the IV equation implied by the algorithmic mechanism. Under local algorithmic monotonicity (LAM), both the ATE and CATE are point identified; without LAM the sharp identified set is characterised. DAIV is sqrt(n)-consistent, asymptotically normal, and semiparametrically efficient, with a formal LAM test supplied. A synthetic ride-sharing example (n = 10,000) shows that standard DML overstates the treatment effect by 52% relative to DAIV; a Hausman-type specification test strongly rejects no algorithmic interference.

Keywords: algorithmic interference, SUTVA, potential outcomes, causal identification, double machine learning, platform economies, instrumental variables, semiparametric efficiency, partial identification

JEL codes: C14, C21, C26, D47, L86

1 Introduction

Modern economies are increasingly organised through digital platforms—ride-sharing markets, online labour exchanges, e-commerce marketplaces, and social-media recommendation systems. A defining feature of these environments is that interactions between market participants are not governed by direct bilateral contact but are mediated by a deterministic algorithm: the platform's matching engine, pricing function, or recommendation policy. This architecture has a profound but largely unrecognised implication for causal identification. Every randomised controlled trial (RCT) run on a platform violates the Stable Unit Treatment Value Assumption (SUTVA) of Rubin (1974) in a manner that is structural, systematic, and knowable.

Consider a randomised experiment on a ride-sharing platform. Treatment units (drivers offered a bonus) and control units (drivers not offered the bonus) interact through the dispatch algorithm, which allocates trips on the basis of the full vector of driver states. If a treated driver accepts a trip that would otherwise have gone to a control driver, the control driver's outcome changes. This is not the general interference problem of Manski (1990) or Aronow and Samii (2017), where interactions flow through an unobserved social network of unknown structure. It is an interference problem with a *known mechanism*: the algorithm is a precisely specified, auditable function of inputs that the researcher can, in principle, recover from engineering documentation or from observing the platform's decisions over time.

This paper exploits that knowledge; our contribution is fourfold.

Contribution 1: Structural model of algorithmic interference. We introduce a formal model of algorithmic interference, in which the potential outcome of unit i under treatment vector W is written as

$$Y_i(W) = g_i(W_i, A_i(W), U_i),$$

where $A_i(W) \in \mathbb{R}^{d_A}$ is the *algorithmic exposure* of unit i—a low-dimensional sufficient statistic of the interference it experiences, computed from the platform's decision rule—and U_i is unobserved heterogeneity.

Contribution 2: Identification and sharp partial identification. We propose the condition of *local algorithmic monotonicity* (LAM): the algorithmic exposure is weakly increasing in own treatment almost surely, in analogy with the monotonicity condition of Angrist and Imbens (1994). Under LAM, both the ATE and the CATE are point identified (Theorem 2). The sharpness of this result is established formally: Proposition 4 shows that without LAM the identified set is exactly the interval $[\theta_0^{\text{DML}} - B, \theta_0^{\text{DML}} + B]$, where B is a spillover-amplitude bound computable from the known algorithm. We further provide in Proposition 5 a formal, asymptotically valid test of the LAM condition itself.

Contribution 3: The DAIV estimator with fully self-contained theory. We construct the Debiased Algorithmic Instrumental Variable (DAIV) estimator, which combines the cross-fitting strategy of Double/Debiased Machine Learning (DML; Chernozhukov et al. 2018) with the structural IV equation implied by the algorithmic model. The asymptotic normality proof (Theorem 7) is provided in full, including an explicit verification of Neyman orthogonality for the DAIV score function adapted to the interference setting. Lemma 6 verifies Assumption 7 for two canonical nuisance function classes—random forests and LASSO—under mild structural conditions on the algorithm. DAIV achieves the $\sqrt{n}$ parametric rate and is semiparametrically efficient.

Contribution 4: Illustrative numerical example with formal diagnostics. To demonstrate the practical magnitude of the bias and the behaviour of the diagnostic tests, we construct a synthetic DGP calibrated to a ride-sharing setting ($n = 10{,}000$, ten drivers per zone, oracle linear residualization). Standard DML overstates the direct bonus effect by approximately 52% relative to DAIV when the interference parameter $\gamma_0 = 0.80$. The Hausman-type test of Proposition 10 strongly rejects the null of no algorithmic interference. We emphasise that all numbers in Section 6 are generated from the stated DGP and are fully reproducible from the accompanying code.

Related literature. Our work relates to four streams. First, the partial identification literature under network interference: Manski (1990), Hudgens and Halloran (2008), Aronow and Samii (2017), Fogarty and Feller (2019). These papers treat interference as flowing through an unobserved network of unknown structure and accordingly deliver partial identification bounds. We instead exploit that the mechanism is known, which collapses the identified set to a point (Theorem 2), and we make the sharpness of this collapse precise (Proposition 4). Second, the experimental design literature for platforms: Bakshy et al. (2012), Botomaní and Bahadori (2021), Athey et al. (2018). These papers propose design-side solutions; we instead develop post-hoc estimation tools applicable to the large archive of already-completed platform RCTs. Third, the DML and semiparametric efficiency literature: Chernozhukov et al. (2018), Newey and McFadden (1994), Robins and Rotnitzky (1995). We extend the DML framework to algorithmically mediated interference, re-deriving the Neyman orthogonality condition for the resulting score and providing sufficient conditions on the nuisance classes. Fourth, the mechanism design and structural IO literature: Roth (2016), Weyl (2010). We adopt the same primitive of a known mechanism but redirect it toward causal identification rather than welfare analysis.

The remainder of the paper is organised as follows. Section 2 presents the structural model. Section 3 establishes identification, the partial identification result, and the LAM test. Section 4 introduces DAIV and proves its asymptotic properties, including nuisance-rate verification. Section 5 reports Monte Carlo results. Section 6 presents the illustrative numerical example. Section 7 concludes. Proofs are in the Appendix.

2 The Structural Model of Algorithmic Interference

2.1 Setup

Let $\mathscr{I}_n = \{1, \ldots, n\}$ be a finite population of units. Associated with each unit i is a covariate vector $X_i \in \mathscr{X} \subset \mathbb{R}^{d_X}$, a binary treatment indicator $W_i \in \{0, 1\}$, and a scalar outcome $Y_i \in \mathbb{R}$. Write $W = (W_1, \ldots, W_n)^\top \in \{0, 1\}^n$ for the treatment vector.

Definition 1 (Platform algorithm)**.** *A* platform algorithm *is a measurable function* $A : \{0, 1\}^n \times \mathscr{X}^n \to \mathbb{R}^{d_A}$*, where* $A_i(W, X) \in \mathbb{R}^{d_A}$ *denotes the* algorithmic exposure *of unit* i *under treatment vector* W *and covariates* $X = (X_1, \ldots, X_n)^\top$.

The key feature is that A is known to the econometrician—it is the documented, auditable decision rule of the platform.[1]

Example 2 (Ride-hailing dispatch)**.** *A dispatch algorithm assigns a trip request of type* r *to driver* i *if and only if* $i = \arg\min_{j:W_j=1 \text{ or } j\in S} \operatorname{dist}(j, r)$*, where* S *is the set of inactive drivers. The algorithmic exposure of driver* i *is* $A_i(W) = \sum_{j\neq i} W_j \cdot \mathbf{1}\{j \text{ competes with } i\}$*, the number of competing treated drivers. An increase in* W_j *for* $j \neq i$ *directly reduces driver* i*'s trip assignment probability, violating SUTVA.*

Example 3 (Marketplace ranking)**.** *An e-commerce platform ranks seller* i*'s listing by quality score* $q_i(W_i, X_i)$*. The exposure* $A_i(W) = \sum_{j\neq i} \mathbf{1}\{q_j(W_j, X_j) > q_i(W_i, X_i)\}$ *counts sellers ranked above* i*. Note that* A_i *is weakly increasing in* W_j *for any* $j \neq i$ *(more treated competitors reduces* i*'s rank), and weakly decreasing in* W_i *for fixed* W_{-i} *if treatment raises* q_i*, so LAM is satisfied from* i*'s perspective (higher* W_i *raises* q_i*, reduces* A_i*)—a case of negative exposure. The generalisation to decreasing algorithms follows by sign-flipping.*

2.2 Potential outcomes under algorithmic interference

Definition 4 (Potential outcome function)**.** *The* potential outcome *of unit* i *under treatment vector* W *is*

$$Y_i(W) = g_i(W_i, A_i(W), U_i), \tag{1}$$

where $g_i : \{0, 1\} \times \mathbb{R}^{d_A} \times \mathbb{R}^{d_U} \to \mathbb{R}$ *is measurable and* $U_i \in \mathbb{R}^{d_U}$ *represents unobserved heterogeneity.*

Equation (1) decomposes the dependence of Y_i on W into a direct channel through W_i and a spillover channel through A_i. When $A_i \equiv 0$ for all W, SUTVA holds and (1) collapses to the standard potential-outcomes model.

[1] Auditability varies across platforms and jurisdictions. In practice, approximate knowledge—recovered from engineering documentation or estimated from the platform's observed decisions—may suffice if the estimation error is small relative to the scale of interference.

2.3 The causal parameter of interest

Define the unit-level direct treatment effect holding exposure fixed:

$$\tau_i(A) = g_i(1, A, U_i) - g_i(0, A, U_i). \tag{2}$$

Fix a reference treatment vector $W^* \in \{0,1\}^n$—the all-zeros vector $\mathbf{0}$ unless stated otherwise—and let $A_i^* := A_i(W^*)$. Our primary estimand is

$$\theta_0 = \mathbb{E}[\tau_i(A_i^*)] = \mathbb{E}[g_i(1, A_i^*, U_i) - g_i(0, A_i^*, U_i)]\,. \tag{3}$$

Remark 5. *Different choices of W^* yield different estimands, each interpretable as a partial-equilibrium ATE at a specific counterfactual exposure level. The identification formula in Theorem 2 delivers θ_0 at $W^* = \mathbf{0}$ directly from the observable distribution; Remark 5 below makes this precise.*

3 Identification

3.1 Assumptions

Assumption 6 (Random assignment)**.** *$(W_i, X_i, U_i)_{i=1}^n$ are i.i.d. draws from a superpopulation distribution P, treatment is unconfounded: $W_i \perp\!\!\!\perp U_i \mid X_i$, and the propensity score $e_0(x) := P(W_i = 1 \mid X_i = x)$ is a measurable function of x only.*

Assumption 7 (Overlap)**.** *There exists $\eta > 0$ such that $\eta \le e_0(X_i) \le 1 - \eta$ a.s.*

Assumption 8 (Separability)**.** *The potential outcome satisfies $g_i(w, a, u) = \alpha(w, a, X_i) + u$ for some measurable $\alpha : \{0,1\} \times \mathbb{R}^{d_A} \times \mathscr{X} \to \mathbb{R}$, and $\mathbb{E}[U_i \mid X_i, W_i, A_i(W)] = 0$.*

Assumption 9 (Partial linearity)**.** *The function α is partially linear in treatment:*

$$\alpha(w, a, x) = \theta_0 w + h_0(a, x), \tag{4}$$

for some $\theta_0 \in \mathbb{R}$ and an unknown measurable $h_0 : \mathbb{R}^{d_A} \times \mathscr{X} \to \mathbb{R}$.

Remark 10. *Assumption 1 (unconfoundedness given X_i) does not imply the mean-independence condition in Assumption 3, because $A_i(W)$ depends on W_{-i}, which is independent of U_i by Assumption 1 but can be correlated with W_i through the joint randomisation. Assumption 3 is therefore an independent restriction ruling out sorting on anticipated spillovers. In a fully randomised experiment with known algorithm both conditions hold simultaneously.*

Assumption 11 (Local algorithmic monotonicity (LAM))**.** *For each unit i and for almost every realisation of $W_{-i} = (W_j)_{j \ne i}$,*

$$A_i(1, W_{-i}) \ge A_i(0, W_{-i}) \quad \textit{component-wise, a.s.}$$

Remark 12. *LAM says that switching unit i from control to treatment weakly increases its own algorithmic exposure. This holds in the ride-hailing setting of Example 1 (treated drivers are more available and therefore receive at least as many trips) and in marketplace experiments where treatment raises the quality score. LAM fails when the platform throttles treated units, e.g. by redistributing supply away from high-activity drivers. Section 3.4 provides a formal test for LAM that can be applied prior to estimation.*

Assumption 13 (Known algorithm and regularity)**.**

(i) The function A is known and computable by the econometrician.

(ii) The exposure $A_i(W)$ *has bounded support:* $\|A_i(W)\|_\infty \le \bar{A} < \infty$ *a.s.*

(iii) The function $\alpha(w,\cdot,x)$ *is Lipschitz-continuous in its second argument uniformly over* (w,x)*: there exists* $L_\alpha < \infty$ *such that* $|\alpha(w,a,x) - \alpha(w,a',x)| \le L_\alpha \|a - a'\|$ *for all* $(w,x) \in \{0,1\} \times \mathscr{X}$ *and all* $a, a' \in \mathbb{R}^{d_A}$.

(iv) $\kappa_0 := \mathbb{E}[\mathrm{Var}(W_i \mid X_i, A_i(W))] > 0$.

Remark 14 (On Condition (iv) and multivariate exposure)**.** *Condition (iv) requires that the residual variation in* W_i *after conditioning on* (X_i, A_i) *is strictly positive on average. When* $d_A = 1$*, LAM together with Assumption 2 is sufficient to guarantee (iv): a monotone, bounded, non-trivial function of* W_i *cannot carry all of the treatment variation. Formally, suppose* $d_A = 1$*. Since* $A_i(1, W_{-i}) \ge A_i(0, W_{-i})$ *a.s. and* W_{-i} *is exogenous, the conditional variance of* W_i *given* (X_i, A_i) *equals* $P(W_i = 1 \mid X_i, A_i)(1 - P(W_i = 1 \mid X_i, A_i))$*; Assumption 2 guarantees this is bounded below by* $\eta(1-\eta)$ *on the support of* (X_i, A_i)*, so* $\kappa_0 \ge \eta(1-\eta) > 0$*. When* $d_A > 1$*, component-wise monotonicity does not preclude the event* $\{A_i(W) = c \cdot \mathbf{1}_{d_A} W_i\}$ *for some constant* $c > 0$*, in which* A_i *would encode* W_i *exactly and render* $\kappa_0 = 0$*. Lemma 1 provides a set of verifiable sufficient conditions on the algorithm structure for the multivariate case.*

Lemma 15 (Non-degeneracy under multivariate exposure)**.** *Suppose* $d_A > 1$ *and Assumptions 1–5 hold. If there exists a component* $k \in \{1, \dots, d_A\}$ *such that the partial derivative* $\partial_{W_i} A_i^{(k)}(W_i, W_{-i})$ *is strictly positive with positive probability and the remaining components* $\{A_i^{(j)}\}_{j \ne k}$ *are not measurable functions of* W_i *alone (i.e., they depend non-trivially on* W_{-i}*), then Assumption 6(iv) holds.*

Proof. See Appendix C. Q.E.D. □

3.2 Identification theorem

Theorem 16 (Point identification)**.** *Under Assumptions 1, 2, 3, 4, and 6,* θ_0 *as defined in* (3) *is point identified from the observable distribution of* $(Y_i, W_i, X_i, A_i(W))$*. Specifically,*

$$\theta_0 = \frac{\mathbb{E}\left[\tilde{Y}_i \tilde{W}_i\right]}{\mathbb{E}\left[\tilde{W}_i^2\right]}, \tag{5}$$

where $\tilde{Y}_i = Y_i - \mathbb{E}[Y_i \mid X_i, A_i]$ and $\tilde{W}_i = W_i - \mathbb{E}[W_i \mid X_i, A_i]$. LAM (Assumption 5) is not required for this result; it is invoked separately in Remark 5 to connect θ_0 to the potential-outcome estimand at $W^* = \mathbf{0}$.

Proof. See Appendix A. Q.E.D. □

Remark 17. *The identification formula* (5) *delivers θ_0 at $W^* = \mathbf{0}$ from the observable distribution because under Assumption 3 and LAM, the mean potential outcome $\mathbb{E}[Y_i(w, A_i^*)]$ is identified by the partialling-out formula at the observed exposure level. Formally, $\theta_0 = \theta_0^{\text{obs}}$ where θ_0^{obs} is defined by* (5)*.*

Corollary 18 (CATE identification)**.** *Under the same assumptions, $\theta_0(x) = \mathbb{E}[\tau_i(A_i^*) \mid X_i = x]$ satisfies*

$$\theta_0(x) = \frac{\mathbb{E}[\tilde{Y}_i \tilde{W}_i \mid X_i = x]}{\mathbb{E}[\tilde{W}_i^2 \mid X_i = x]}.$$

The formula (5) is a partial linear regression of residualised outcomes on residualised treatments, where residuals are formed by conditioning on both X_i and A_i. When SUTVA holds, $A_i \equiv 0$ and (5) reduces to the DML identifying formula of Chernozhukov et al. (2018).

3.3 Partial identification without LAM

Theorem 2 requires LAM. The following proposition characterises the identified set when LAM fails and establishes that the point identification result is sharp.

Proposition 19 (Sharp partial identification without LAM)**.** *Suppose Assumptions 1–4 and 6 hold, but LAM (Assumption 5) is replaced by the weaker condition that the signed spillover $\Delta_i(W_{-i}) := A_i(1, W_{-i}) - A_i(0, W_{-i})$ satisfies $\|\Delta_i\|_\infty \leq \bar{\Delta} < \infty$ a.s. but is of unrestricted sign. Define the naive DML probability limit $\theta_0^{\text{DML}} := \operatorname{plim} \hat{\theta}^{\text{DML}}$, where $\hat{\theta}^{\text{DML}}$ applies the formula* (5) *ignoring A_i. Then the sharp identified set for θ_0 is*

$$\Theta_I = \left[\theta_0^{\text{DML}} - B,\ \theta_0^{\text{DML}} + B\right], \tag{6}$$

where the bias bound is

$$B = \frac{L_\alpha \bar{\Delta} \mathbb{E}\left[\|\tilde{W}_i\|_1\right]}{\kappa_0} > 0. \tag{7}$$

Both endpoints of Θ_I are attained: there exist data-generating processes consistent with Assumptions 1–4 and 6 that achieve $\theta_0 = \theta_0^{\text{DML}} \pm B$.

Proof. See Appendix A. Q.E.D. □

Remark 20. *Proposition 4 makes the role of LAM precise: it collapses the interval Θ_I to a point by resolving the sign ambiguity in the bias formula. Under LAM ($\Delta_i \geq 0$ a.s.), the exposure increment is non-negative, which fixes the sign of the omitted-variable term $\mathbb{E}[\delta h_i(W_i - e_0(X_i))]$ and pins down θ_0 uniquely. Without LAM the sign of Δ_i is unrestricted, and both endpoints of*

Θ_I are attained. A consistent estimator of the bound is $\hat{B}_n = \hat{L}_\alpha \hat{\bar{\Delta}}_n \hat{\mathbb{E}}[\|\tilde{W}_i\|_1]/\hat{\kappa}_n$, where each component is estimated by its sample analogue.

3.4 Testing local algorithmic monotonicity

Since LAM is a property of the known algorithm A, it can be tested using the observed treatment vector.

Proposition 21 (Testability of LAM)**.** *Define the signed exposure increment for unit i and draw W_{-i} from its empirical distribution:*

$$\hat{\Delta}_i = A_i(1, W_{-i}^{\text{obs}}) - A_i(0, W_{-i}^{\text{obs}}).$$

Let $\hat{\Delta}_i^{(k)}$ denote the k-th component. Consider the test statistic

$$T_n = \sqrt{n} \min_{k=1,\ldots,d_A} \frac{\frac{1}{n}\sum_{i=1}^{n} \hat{\Delta}_i^{(k)}}{\hat{\sigma}_k}, \tag{8}$$

where $\hat{\sigma}_k^2 = n^{-1}\sum_{i=1}^n (\hat{\Delta}_i^{(k)})^2 - \big(n^{-1}\sum_{i=1}^n \hat{\Delta}_i^{(k)}\big)^2$. Under the sharp null $H_0 : \mathbb{E}[\hat{\Delta}_i^{(k)}] = 0$ for some k, $T_n \xrightarrow{d} N(0,1)$. Under LAM, $T_n \to +\infty$ in probability. The one-sided test that rejects H_0 when $T_n > z_{1-\alpha/d_A}$ has asymptotic size α and is consistent against any fixed alternative satisfying $\min_k \mathbb{E}[\hat{\Delta}_i^{(k)}] > 0$.

Proof. See Appendix A. Q.E.D. □

Remark 22. *The test statistic T_n exploits the fact that $\hat{\Delta}_i$ is a deterministic function of $(W_{-i}^{\text{obs}}, X_i)$ given the known algorithm A, so the only source of randomness enters through the empirical distribution of W_{-i}. Under the Bernoulli assignment of Assumption 1, the CLT applies component-wise: for each k separately, $\sqrt{n}\hat{\bar{\Delta}}^{(k)} \xrightarrow{d} N(0, \sigma_k^2)$; no independence across components is assumed. For multivariate A the critical value is Bonferroni-corrected to $z_{1-\alpha/d_A}$ to maintain familywise size control across all d_A components.*

4 Estimation: The DAIV Estimator

4.1 Motivation and moment representation

Theorem 2 identifies θ_0 through the moment condition $\mathbb{E}[\psi_i(\theta_0, m_0, p_0)] = 0$, where the score is

$$\psi_i(\theta, m, p) := \big(Y_i - m(X_i, A_i) - \theta(W_i - p(X_i, A_i))\big)\big(W_i - p(X_i, A_i)\big), \tag{9}$$

and the nuisance functions are

$$m_0(x,a) := \mathbb{E}[Y_i \mid X_i = x, A_i = a], \qquad p_0(x,a) := \mathbb{E}[W_i \mid X_i = x, A_i = a].$$

Plugging in preliminary estimators $\hat{m}$, $\hat{p}$ directly into the sample analogue of (5) introduces a regularisation bias of order $\|\hat{m} - m_0\|_2 \cdot \|\hat{p} - p_0\|_2$ (Chernozhukov et al. 2018, Lemma 1) that need not vanish at the parametric rate. The DAIV estimator eliminates this bias by cross-fitting, exploiting the Neyman orthogonality of the score (9), which we verify explicitly in the proof of Theorem 7.

The DAIV estimator can equivalently be written as the solution to the cross-fitted moment equation:

$$\hat{\theta}^{\text{DAIV}} = \left(\frac{1}{n}\sum_{\ell}\sum_{i\in I_\ell}(\tilde{W}_i^{(-\ell)})^2\right)^{-1}\left(\frac{1}{n}\sum_{\ell}\sum_{i\in I_\ell}\tilde{Y}_i^{(-\ell)}\tilde{W}_i^{(-\ell)}\right), \tag{10}$$

which is the unique solution to $n^{-1}\sum_i \psi_i(\theta, \hat{m}^{(-\ell(i))}, \hat{p}^{(-\ell(i))}) = 0$ with respect to θ.

4.2 Algorithm

Let $\{I_\ell\}_{\ell=1}^{L}$ be a partition of $\{1,\ldots,n\}$ into L folds of size approximately n/L, with complement I_ℓ^c.

Step 1: Compute algorithmic exposures. For each unit i, compute $A_i := A_i(W)$ using the known platform algorithm and the observed treatment vector W.

Step 2: Estimate nuisance functions via cross-fitting. For each fold ℓ, using only observations in I_ℓ^c, train:

$$\hat{m}^{(-\ell)} := \arg\min_{m\in\mathscr{H}_m}\sum_{i\in I_\ell^c}(Y_i - m(X_i,A_i))^2, \tag{11}$$

$$\hat{p}^{(-\ell)} := \arg\min_{p\in\mathscr{H}_p}\sum_{i\in I_\ell^c}(W_i - p(X_i,A_i))^2, \tag{12}$$

where $\mathscr{H}_m$ and $\mathscr{H}_p$ are regularised function classes.

Step 3: Compute cross-fitted residuals. For each $i \in I_\ell$, form

$$\tilde{Y}_i^{(-\ell)} = Y_i - \hat{m}^{(-\ell)}(X_i,A_i), \qquad \tilde{W}_i^{(-\ell)} = W_i - \hat{p}^{(-\ell)}(X_i,A_i).$$

Step 4: Aggregate via (10).

4.3 Asymptotic theory

Assumption 23 (Nuisance convergence rates)**.** *There exists $\delta_n \to 0$ such that*

$$\|\hat{m}^{(-\ell)} - m_0\|_2 = O_p(\delta_n), \qquad \|\hat{p}^{(-\ell)} - p_0\|_2 = O_p(\delta_n), \qquad \delta_n = o_p(n^{-1/4}).$$

Assumption 7 requires only that the product of nuisance errors vanishes faster than $n^{-1/2}$; individual rates slower than $n^{-1/4}$ are permitted. Lemma 6 below verifies this assumption for two canonical estimator classes.

Lemma 24 (Sufficient conditions for Assumption 7)**.** *Let $d = d_X + d_A$ denote the total dimension of the conditioning vector (X_i, A_i).*

(a) ***(Random forests.)*** *Suppose m_0 and p_0 lie in a Hölder class $\mathscr{C}^s(\mathscr{X} \times [0, \bar{A}]^{d_A})$ with $s > d/2$, and that the tree depth grows as $\lfloor \log n \rfloor$. Then random forests with 500 trees achieve $\|\hat{f} - f_0\|_2 = O_p(n^{-s/(2s+d)})$. Assumption 7 holds whenever $s > d/4$.*

(b) ***(LASSO with sparsity.)*** *Suppose m_0 and p_0 admit sparse linear representations in a dictionary $\Phi : \mathscr{X} \times \mathbb{R}^{d_A} \to \mathbb{R}^p$ with sparsity $s_0 \ll p$ and $s_0 \log p = o(n)$. Then the LASSO achieves $\|\hat{f} - f_0\|_2 = O_p(\sqrt{s_0 \log p / n})$. Assumption 7 holds whenever $s_0 \log p = o(\sqrt{n})$.*

Proof. See Appendix C. Q.E.D. □

Theorem 25 (Asymptotic normality of DAIV)**.** *Under Assumptions 1–4, 5, 6, and 7 and $\kappa_0 > 0$, as $n \to \infty$:*

$$\sqrt{n}\big(\hat{\theta}^{\text{DAIV}} - \theta_0\big) \xrightarrow{d} N(0, V_0)\,, \tag{13}$$

where

$$V_0 = \frac{\mathbb{E}[\psi_i^2]}{\kappa_0^2}, \qquad \psi_i = \tilde{Y}_i \tilde{W}_i - \theta_0 \tilde{W}_i^2. \tag{14}$$

Proof. See Appendix A for the complete, self-contained proof. Q.E.D. □

Corollary 26 (Consistent variance estimator)**.** *The plug-in estimator*

$$\hat{V} = \frac{n^{-1} \sum_i \hat{\psi}_i^2}{\big(n^{-1} \sum_i (\tilde{W}_i^{(-\ell(i))})^2\big)^2}, \qquad \hat{\psi}_i = \tilde{Y}_i^{(-\ell(i))} \tilde{W}_i^{(-\ell(i))} - \hat{\theta}^{\text{DAIV}}\, (\tilde{W}_i^{(-\ell(i))})^2,$$

satisfies $\hat{V} \xrightarrow{p} V_0$. An asymptotically valid $(1-\alpha)$ confidence interval for θ_0 is $\hat{\theta}^{\text{DAIV}} \pm z_{1-\alpha/2} \sqrt{\hat{V}/n}$.

Proof. See Appendix A. Q.E.D. □

4.4 Semiparametric efficiency

Proposition 27 (Semiparametric efficiency bound)**.** *The semiparametric variance bound for regular estimation of θ_0 in the model $\mathscr{M}$ defined by Assumptions 1–4 and 6 equals V_0 as defined in* (14)*. Hence $\hat{\theta}^{\mathrm{DAIV}}$ is semiparametrically efficient.*

Proof. See Appendix B. Q.E.D. □

4.5 A Hausman-type test for algorithmic interference

When the researcher is uncertain whether algorithmic interference is present, the following test compares DAIV (which conditions on A_i) to standard DML (which does not).

Proposition 28 (Hausman test for algorithmic interference)**.** *Under the null $H_0^H : \gamma_0 = 0$ (no interference, i.e., the coefficient on A_i in the outcome equation equals zero), both $\hat{\theta}^{\mathrm{DML}}$ and $\hat{\theta}^{\mathrm{DAIV}}$ are consistent and the Hausman statistic*

$$H_n = \frac{\sqrt{n}\,(\hat{\theta}^{\mathrm{DML}} - \hat{\theta}^{\mathrm{DAIV}})}{\sqrt{\hat{V}^{\mathrm{DML}} - \hat{V}^{\mathrm{DAIV}}}} \tag{15}$$

converges in distribution to $N(0,1)$ under H_0^H. Under the alternative $\gamma_0 \neq 0$, $|H_n| \to \infty$ in probability, so the test is consistent. If $\hat{V}^{\mathrm{DML}} - \hat{V}^{\mathrm{DAIV}} \leq 0$ in a given sample (which can occur in finite samples), the test is uninformative and a bootstrap version of the statistic should be used instead.

Proof. See Appendix A. Q.E.D. □

5 Monte Carlo Simulations

5.1 Design

We generate data from the following data-generating process (DGP):

$$X_i \overset{\mathrm{iid}}{\sim} N(0, I_5), \tag{16}$$

$$W_i \mid X_i \sim \mathrm{Bernoulli}(e_0(X_i)), \quad e_0(x) = \sigma(0.5x_1), \tag{17}$$

$$A_i(W) = \frac{1}{k} \sum_{j \in Z(i)} W_j, \tag{18}$$

$$Y_i = \theta_0 W_i + \gamma_0 A_i(W) + X_i^\top \beta_0 + \varepsilon_i, \quad \varepsilon_i \sim N(0, 0.35^2), \tag{19}$$

where $\sigma(\cdot)$ is the logistic function, $Z(i)$ is the zone containing unit i (including i itself), $k = 10$ is zone size, $\beta_0 = (0.10, -0.08, 0.06, -0.05, 0.04)^\top$, and the n/k zones are assigned uniformly at random (independently of X_i). The true ATE is $\theta_0 = 0.15$; the interference strength $\gamma_0/\theta_0 \in \{0.5, 1.0, 2.0\}$. Each replication $b = 1, \dots, B$ uses random seed $1000 + b$.

Design rationale. The non-uniform propensity $e_0(X_i) \neq 0.5$ ensures that OLS and DML differ, so that the simulation cleanly isolates the contribution of conditioning on A_i beyond standard covariate adjustment. Since $Z(i)$ includes i itself, flipping W_i shifts the zone mean by exactly $1/k$, so $\mathrm{Cov}(W_i, A_i) = \mathrm{Var}(W_i)/k > 0$, which yields the analytical DML probability limit $\theta_0^{\mathrm{DML}} = \theta_0 + \gamma_0/k$, providing an analytical benchmark for each design. LAM holds with $\bar{\Delta} = 1/k$ because $A_i(1, W_{-i}) - A_i(0, W_{-i}) = 1/k > 0$ for all i.

All three estimators (OLS, DML, DAIV) use oracle linear residualization (nuisance functions fitted by OLS on the stated regressors); this removes ML approximation error and isolates the structural bias. We set $B = 2{,}000$ Monte Carlo repetitions and $n \in \{500, 1000, 2000\}$.

5.2 Results

Table 1 reports Monte Carlo bias, RMSE, and empirical coverage of nominal 95% confidence intervals. The "F-stat" column is the DAIV first-stage strength statistic $\hat{\kappa}_n^2 / \widehat{\mathrm{Var}(\tilde{W}_i^2)}$; values far above 10 indicate no weak-identification concern.

Table 1: Monte Carlo Results. True parameter $\theta_0 = 0.15$. $B = 2{,}000$ replications; propensity $e(X) = \sigma(0.5X_1)$; zone size $k = 10$.

		OLS		DML			DAIV			
γ_0	n	Bias	RMSE	Bias	RMSE	Cov.	Bias	RMSE	Cov.	F-stat
0.075	500	0.054	0.064	0.007	0.033	0.94	0.000	0.034	0.95	903
	1000	0.055	0.060	0.007	0.024	0.93	0.000	0.024	0.94	1832
	2000	0.055	0.057	0.008	0.018	0.92	0.000	0.017	0.95	3692
0.150	500	0.062	0.071	0.014	0.035	0.92	0.000	0.034	0.95	903
	1000	0.062	0.067	0.015	0.028	0.90	0.000	0.024	0.94	1832
	2000	0.062	0.065	0.015	0.022	0.84	0.000	0.017	0.95	3692
0.300	500	0.076	0.084	0.029	0.044	0.86	0.000	0.034	0.95	903
	1000	0.077	0.081	0.030	0.038	0.75	0.000	0.024	0.94	1832
	2000	0.077	0.079	0.030	0.034	0.54	0.000	0.017	0.95	3692

Source: Authors' calculations. OLS conditions on W_i only. DML residualizes on X_i. DAIV residualizes on (X_i, A_i). All nuisance functions fitted by oracle OLS. Coverage is for nominal 95% intervals using the sandwich variance of Corollary 8. F-stat is the first-stage strength statistic; $\hat{\kappa}_n \approx 0.226$ throughout.

OLS exhibits bias driven by confounding through the propensity $e_0(X_i) \neq 0.5$; DML removes this confounding but inherits the interference bias γ_0/k because it does not condition on A_i. Both biases are asymptotically non-negligible and grow proportionally with γ_0. DML coverage collapses from 0.92 to 0.54 as γ_0 increases from 0.075 to 0.300—a stark illustration that ignoring algorithmic interference invalidates standard inference even when standard unconfoundedness holds. DAIV bias is numerically zero across all designs, coverage is uniformly close to 0.95, and the first-stage F-statistic grows linearly in n, consistent with Assumption 6(iv). These results confirm analytically: plim $\hat{\theta}^{\mathrm{DML}} = \theta_0 + \gamma_0/k$, verified to three decimal places in each row.

6 Illustrative Numerical Example

This section demonstrates the bias and the behaviour of the diagnostics in a fully transparent, reproducible synthetic setting. All numbers below are generated from the stated DGP with a fixed seed; the code is included in the replication package.

6.1 Data-generating process

Setup. We draw $n = 10{,}000$ units and assign them uniformly at random to $n/k = 1{,}000$ zones of size $k = 10$. The covariate vector is $X_i \sim N(0, I_5)$, independent of zone membership, with $\beta_0 = (0.10, -0.08, 0.06, -0.05, 0.04)^\top$. Treatment is i.i.d. $W_i \sim \mathrm{Bernoulli}(0.5)$.

Platform algorithm. The algorithmic exposure is the zone treatment density:

$$A_i = \frac{1}{k} \sum_{j \in Z(i)} W_j, \tag{20}$$

where $Z(i)$ denotes the zone containing unit i. Note that A_i includes W_i itself, so that

$$A_i(W_i = 1, W_{-i}) - A_i(W_i = 0, W_{-i}) = \frac{1}{k} > 0 \quad \text{for all } i,$$

i.e. LAM holds with $\bar{\Delta} = 1/k$ constant.

Outcome.

$$Y_i = \theta_0 W_i + \gamma_0 A_i + X_i^\top \beta_0 + \varepsilon_i, \qquad \varepsilon_i \sim N(0, 0.35^2), \tag{21}$$

with $\theta_0 = 0.15$ (direct treatment effect) and $\gamma_0 = 0.80$ (interference). The theoretical DML probability limit is $\theta_0^{\mathrm{DML}} = \theta_0 + \gamma_0/k = 0.15 + 0.08 = 0.23$, so the asymptotic DML bias is exactly γ_0/k, providing a clean analytical check on the simulation output.

Estimation. All three estimators (OLS, DML, DAIV) use oracle linear residualization; nuisance functions are fitted by OLS on the stated regressors. This keeps the illustration uncontaminated by ML approximation error and focuses attention on the structural bias. Random seed: 2025.

6.2 Pre-estimation diagnostics

LAM test. Since $\hat{\Delta}_i = 1/k = 0.10$ is constant across all units, the test statistic T_n of Proposition 5 diverges as $\sqrt{n} \times 0.10/0 = +\infty$; any finite critical value is exceeded. The LAM condition is verified analytically from the stated algorithm (20).

First-stage strength. The cross-fitted first-stage conditional variance is $\hat{\kappa}_n = 0.225$ and the first-stage F-statistic equals 26,050, well above the standard threshold of 10. This large value reflects the strong within-zone variation in treatment that remains after partialling out A_i.

Nuisance fit. The out-of-sample R^2 for the OLS fits of m_0 and p_0 equal 0.27 and 0.10, respectively. The modest $R^2_{p_0}$ reflects that A_i absorbs a share of the treatment variation but substantial residual variation in W_i remains (as guaranteed by LAM and Assumption 6(iv)).

6.3 Main results

Table 2 presents the estimates.

Table 2: Synthetic Illustration: Treatment Effect Estimates Under Zone-Level Algorithmic Interference ($n = 10{,}000$, $k = 10$, $\theta_0 = 0.15$, $\gamma_0 = 0.80$, seed 2025).

	OLS	DML	DAIV	Hausman H_n
Estimate	0.234	0.230	0.152	
Std. Error	(0.008)	(0.007)	(0.007)	
95% CI	[0.219, 0.250]	[0.216, 0.244]	[0.139, 0.166]	
Relative to DAIV	+54.0%	+51.8%	—	
DML vs DAIV				149.4^{***}

Source: Authors' calculations. OLS conditions on W_i only. DML residualizes on X_i. DAIV residualizes on (X_i, A_i). Standard errors from the sandwich formula (14). The Hausman statistic H_n tests the null of no algorithmic interference (Proposition 10); $***$ denotes $p < 0.01$. True parameter: $\theta_0 = 0.15$.

OLS and DML are both close to the theoretical probability limit of 0.23, confirming the analytical bias formula $\theta_0 + \gamma_0/k$. DAIV recovers the true value of 0.15 to within sampling error (the true value lies inside the 95% confidence interval). The Hausman statistic $H_n = 149.4$ ($p < 0.001$) strongly rejects the null of no algorithmic interference.

The bias mechanism is transparent in this DGP. DML residualizes on X_i but not on A_i; the residual $\tilde{W}_i^{\mathrm{DML}}$ therefore retains the correlation $\mathrm{Cov}(W_i, A_i \mid X_i) = \mathrm{Var}(W_i)/k$, which confounds the moment condition with the spillover term $\gamma_0 A_i$. DAIV eliminates this confound by also conditioning on A_i, leaving a residual $\tilde{W}_i^{\mathrm{DAIV}}$ that is orthogonal to the algorithmic channel.

6.4 Partial identification bounds

To illustrate the value of LAM, we compute the bound of Proposition 4 using the true Lipschitz constant $\hat{L}_\alpha = \gamma_0 = 0.80$ and $\hat{\bar{\Delta}}_n = 1/k = 0.10$. This gives

$$\hat{B}_n = 0.80 \times 0.10 \times \frac{\mathbb{E}_n[|\tilde{W}_i^{\mathrm{DML}}|]}{\hat{\kappa}_n^{\mathrm{DML}}} = 0.160,$$

so the agnostic interval without LAM is $[0.070, 0.390]$. The DAIV point estimate of 0.152 lies inside this interval. LAM collapses it to a single point, illustrating the identifying power of the monotonicity condition.

7 Conclusion

This paper has introduced a structural framework for causal inference in platform economies where SUTVA fails by design. The central insight is that the mechanism generating interference—the platform algorithm—is often known to the researcher, and that this knowledge is not merely qualitatively informative. We have shown that it enables point identification (Theorem 2), characterised exactly the cost of discarding this knowledge (Proposition 4), provided an implementable test of the key identifying condition (Proposition 5), and constructed an efficient estimator with fully verified asymptotic properties (Theorem 7 and Proposition 9).

Several extensions merit investigation. First, the paper assumes binary treatment; extension to multi-valued or continuous treatments is conceptually straightforward under the structural model but requires additional regularity for the partial-linearity assumption. Second, the paper treats the platform algorithm as static; adaptive algorithms that respond to experimental findings in real time break the fixed-A structure of Assumption 6(i) and require a dynamic extension. Third, a welfare analysis of optimal experimental design under algorithmic interference would complement the identification and estimation results here.

More broadly, hundreds of RCTs are conducted annually on major platforms whose results inform both internal business decisions and academic publication (Kohavi et al. 2020). The synthetic illustration in Section 6 shows that even at modest interference strength the DML bias can exceed 50% of the true parameter value. The tools developed here make it feasible to correct this bias at scale, requiring only that researchers observe and model the algorithmic mechanism—information that platforms increasingly make available for research purposes.

References

Angrist, J., Imbens, G., 1994. Identification and estimation of local average treatment effects. *Econometrica* 62 (3), 467–475.

Aronow, P. M., Samii, C., 2017. Estimating average causal effects under interference between units. *Annals of Applied Statistics* 11 (4), 1912–1947.

Athey, S., Eckles, D., Imbens, G. W., 2018. Exact p-values for network interference. *Journal of the American Statistical Association* 113 (521), 230–240.

Bakshy, E., Karrer, B., Adamic, L., 2012. Social influence and the diffusion of user-created content. In: Proceedings of the ACM Conference on Electronic Commerce, pp. 325–334.

Botomaní, S., Bahadori, M. T., 2021. Switchback experiments under stochastic interference. Working Paper.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., Robins, J., 2018. Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal* 21 (1), C1–C68.

Fogarty, C. B., Feller, A., 2019. Studentized sensitivity analysis for the sample average treatment effect in paired observational studies. *Journal of the American Statistical Association* 115 (531), 722–734.

Hudgens, M. G., Halloran, M. E., 2008. Toward causal inference with interference. *Journal of the American Statistical Association* 103 (482), 832–842.

Kohavi, R., Tang, D., Xu, Y., 2020. Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing. Cambridge University Press, Cambridge.

Manski, C. F., 1990. Nonparametric bounds on treatment effects. *American Economic Review, Papers and Proceedings* 80 (2), 319–323.

Newey, W. K., McFadden, D., 1994. Large sample estimation and hypothesis testing, in: Engle, R. F., McFadden, D. (Eds.), Handbook of Econometrics, Vol. 4. North-Holland, Amsterdam, pp. 2111–2245.

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

Roth, A. E., 2016. Marketplaces, markets, and market design. *American Economic Review* 108 (7), 1609–1658.

Rubin, D. B., 1974. Estimating causal effects of treatments in randomised and non-randomised studies. *Journal of Educational Psychology* 66 (5), 688–701.

Wager, S., 2018. Efficient policy learning. Technical Report, Stanford University.

Wainwright, M. J., 2019. High-Dimensional Statistics: A Non-Asymptotic Viewpoint. Cambridge University Press, Cambridge.

Weyl, E. G., 2010. A price theory of multi-sided platforms. *American Economic Review* 100 (4), 1642–1672.

Appendix A Proofs of Main Results

Proof of Theorem 2

Step 1: Algebraic decomposition. Under Assumption 3, the observed outcome satisfies

$$Y_i = \theta_0 W_i + h_0(X_i, A_i) + U_i. \tag{22}$$

Step 2: Partialling out. Taking conditional expectations of (22) with respect to (X_i, A_i) and using Assumption 1:

$$m_0(X_i, A_i) := \mathbb{E}[Y_i \mid X_i, A_i] = \theta_0\, p_0(X_i, A_i) + h_0(X_i, A_i).$$

Subtracting this from (22) gives

$$\tilde{Y}_i = \theta_0\, \tilde{W}_i + U_i. \tag{23}$$

Step 3: Moment condition. Multiplying (23) by $\tilde{W}_i$ and taking expectations:

$$\mathbb{E}[\tilde{Y}_i \tilde{W}_i] = \theta_0\, \mathbb{E}[\tilde{W}_i^2] + \mathbb{E}[U_i \tilde{W}_i]. \tag{24}$$

We now show $\mathbb{E}[U_i \tilde{W}_i] = 0$. Write $\tilde{W}_i = W_i - p_0(X_i, A_i)$. By Assumption 1, W is independent of U_i given X_i; in particular, $U_i \perp\!\!\!\perp (W_i, W_{-i}) \mid X_i$. Since $A_i = A_i(W)$ is a deterministic function of (W, X_i), we have $U_i \perp\!\!\!\perp A_i \mid X_i$, and hence $U_i \perp\!\!\!\perp (W_i, A_i) \mid X_i$. Therefore

$$\mathbb{E}[U_i \tilde{W}_i] = \mathbb{E}\big[U_i\big(W_i - p_0(X_i, A_i)\big)\big] = \mathbb{E}\big[\mathbb{E}[U_i \mid X_i, A_i] \cdot \tilde{W}_i\big] = 0,$$

where the last equality uses Assumption 3.

Step 4: Identification. Since $\mathbb{E}[\tilde{W}_i^2] = \kappa_0 > 0$ by Assumption 6(iv), solving (24) gives (5). Q.E.D.

Proof of Proposition 4

Upper bound on bias. The naive DML estimator ignores A_i and effectively conditions only on X_i. Its probability limit is

$$\theta_0^{\mathrm{DML}} = \frac{\mathbb{E}[(Y_i - \mathbb{E}[Y_i \mid X_i])(W_i - e_0(X_i))]}{\mathbb{E}[(W_i - e_0(X_i))^2]}.$$

Under Assumption 3, write $\mathbb{E}[Y_i \mid X_i] = \theta_0\, e_0(X_i) + \mathbb{E}[h_0(A_i, X_i) \mid X_i]$. Denoting $\bar{h}(X_i) = \mathbb{E}[h_0(A_i, X_i) \mid X_i]$ and $\delta h_i = h_0(A_i, X_i) - \bar{h}(X_i)$, we get

$$\theta_0^{\mathrm{DML}} - \theta_0 = \frac{\mathbb{E}[\delta h_i\, (W_i - e_0(X_i))]}{\mathbb{E}[(W_i - e_0(X_i))^2]}.$$

Now $|\delta h_i| \le L_\alpha \|A_i - \mathbb{E}[A_i \mid X_i]\|$ by Assumption 6(iii), and the variation of A_i given (X_i, W_{-i}) is bounded by $\bar{\Delta}/2$ in each component (since flipping W_i shifts the exposure by at most $\bar{\Delta}$ by assumption). Therefore, using Cauchy–Schwarz and the triangle inequality,

$$|\theta_0^{\mathrm{DML}} - \theta_0| \le \frac{L_\alpha \bar{\Delta} \mathbb{E}[\|\tilde{W}_i\|_1]}{\kappa_0} =: B.$$

Sharpness. Construct a DGP where $h_0(a,x) = L_\alpha a$ (sign chosen to push θ_0^{DML} toward the upper boundary) and $A_i(1, W_{-i}) - A_i(0, W_{-i}) = \bar{\Delta}$ a.s. Then $\delta h_i \cdot (W_i - e_0(X_i)) = L_\alpha \bar{\Delta}(W_i - e_0(X_i))^2$ a.s., so the numerator equals $L_\alpha \bar{\Delta} \kappa_0$ and the bias $\theta_0^{\mathrm{DML}} - \theta_0 = L_\alpha \bar{\Delta} \kappa_0 / \kappa_0 = B$. The lower endpoint is attained by reversing the sign of h_0. Q.E.D.

Proof of Proposition 5

Under H_0, $\mathbb{E}[\hat{\Delta}_i^{(k)}] = 0$ for some k. Since $\hat{\Delta}_i$ is a deterministic function of $(W_{-i}^{\mathrm{obs}}, X_i)$ and the observations are i.i.d. under Assumption 1, the CLT applies component-wise: for each k separately, $\sqrt{n}\bar{\hat{\Delta}}^{(k)} \xrightarrow{d} N(0, \sigma_k^2)$; no independence across components is assumed or required. The test statistic (8) evaluates the component with the smallest sample mean; under the sharp null that the k-th component has zero mean, $T_n = \sqrt{n}\bar{\hat{\Delta}}^{(k)}/\hat{\sigma}_k \xrightarrow{d} N(0,1)$ and the Bonferroni-corrected critical value $z_{1-\alpha/d_A}$ maintains familywise size control across all d_A components. Under LAM, $\mathbb{E}[\hat{\Delta}_i^{(k)}] \ge c > 0$ for all k by Assumption 5, so $T_n = \sqrt{n}\min_k \bar{\hat{\Delta}}^{(k)}/\hat{\sigma}_k \to +\infty$ in probability. Q.E.D.

Proof of Theorem 7 (Complete, self-contained)

Step 1: Neyman orthogonality. The score function (9) satisfies the Neyman orthogonality condition $\partial_\eta \mathbb{E}[\psi_i(\theta_0, \eta)]\big|_{\eta=\eta_0} = 0$, where $\eta = (m, p)$ and the Gateaux derivative in direction $(\delta m, \delta p)$ is

$$\begin{aligned}
\partial_\eta \mathbb{E}[\psi_i]\big|_{\eta_0}[\delta\eta] = \mathbb{E}[-\delta m(X_i, A_i)(W_i - p_0(X_i, A_i))] \\
+ \mathbb{E}[(Y_i - m_0(X_i, A_i) - \theta_0(W_i - p_0(X_i, A_i)))(-\delta p(X_i, A_i))] \\
+ \mathbb{E}[\theta_0\, \delta p(X_i, A_i)(W_i - p_0(X_i, A_i))].
\end{aligned}$$

The first term equals $-\mathbb{E}[\delta m \cdot \tilde{W}_i]$. By the law of iterated expectations and $\mathbb{E}[U_i \mid X_i, A_i] = 0$, $\mathbb{E}[\tilde{Y}_i \mid X_i, A_i] = \theta_0 \tilde{W}_i$, so the second and third terms cancel: $\mathbb{E}[-\tilde{Y}_i \delta p + \theta_0 \tilde{W}_i \delta p] = -\mathbb{E}[\theta_0 \tilde{W}_i \delta p] + \mathbb{E}[\theta_0 \tilde{W}_i \delta p] = 0$. The first term is also zero because $\mathbb{E}[\tilde{W}_i \mid X_i, A_i] = 0$ by definition of p_0. Hence $\partial_\eta \mathbb{E}[\psi_i]\big|_{\eta_0} = 0$.

Step 2: Taylor expansion. Write $\hat{\eta} = (\hat{m}^{(-\ell)}, \hat{p}^{(-\ell)})$. A second-order Taylor expansion of the cross-fitted sample moment around (θ_0, η_0) gives

$$0 = \frac{1}{\sqrt{n}} \sum_i \psi_i(\theta_0, \eta_0) + \frac{1}{n} \sum_i \partial_\theta \psi_i(\theta_0, \eta_0) \cdot \sqrt{n}(\hat{\theta} - \theta_0) + R_n,$$

where the remainder R_n collects all terms involving $\hat{\eta} - \eta_0$.

Step 3: Bounding the remainder. By Neyman orthogonality (Step 1), the first-order derivative $\partial_\eta \mathbb{E}[\psi_i]|_{\eta_0}$ vanishes, so the remainder R_n consists only of second-order terms. For each unit i in fold ℓ, the cross-fitting construction ensures that $\hat{\eta}^{(-\ell)}$ is trained on I_ℓ^c, which is independent of the fold-ℓ observations. The second-order contribution of unit i is bounded by $O_p(\|\hat{m} - m_0\|_2 \|\hat{p} - p_0\|_2) = O_p(\delta_n^2)$ per unit. Summing over all n units (with the independence from cross-fitting ensuring no further entropy argument is required), the total remainder satisfies

$$R_n = O_p\left(\sqrt{n}\,\delta_n^2\right) = o_p(1),$$

where the last step uses $\sqrt{n}\,\delta_n^2 \to 0$, which follows from $\delta_n = o_p(n^{-1/4})$ in Assumption 7.

Step 4: Central limit theorem. Since $\partial_\theta \psi_i(\theta_0, \eta_0) = -\tilde{W}_i^2$ and $n^{-1} \sum_i \tilde{W}_i^2 \xrightarrow{p} \kappa_0 > 0$ by Assumption 6(iv), we obtain

$$\sqrt{n}(\hat{\theta}^{\text{DAIV}} - \theta_0) = \frac{n^{-1/2} \sum_i \psi_i(\theta_0, \eta_0)}{\kappa_0} + o_p(1).$$

The leading term is a sum of i.i.d. mean-zero random variables (mean zero by Neyman orthogonality; i.i.d. by Assumption 1 and the deterministic structure of A). The variance is $\kappa_0^{-2} \mathbb{E}[\psi_i^2] = V_0 < \infty$ by Assumptions 6(ii)–(iii). The Lindeberg–Feller CLT therefore gives $n^{-1/2} \sum_i \psi_i \xrightarrow{d} N(0, \mathbb{E}[\psi_i^2])$, and Slutsky's theorem completes the proof of (13)–(14). Q.E.D.

Proof of Corollary 8

By the law of large numbers under Assumption 1, $n^{-1} \sum_i \hat{\psi}_i^2 \xrightarrow{p} \mathbb{E}[\psi_i^2]$ and $n^{-1} \sum_i (\tilde{W}_i^{(-\ell(i))})^2 \xrightarrow{p} \kappa_0$, using that $\|\hat{\theta}^{\text{DAIV}} - \theta_0\| = O_p(n^{-1/2})$ from Theorem 7 and the boundedness condition Assumption 6(ii). The continuous mapping theorem gives $\hat{V} \xrightarrow{p} V_0$. Q.E.D.

Proof of Proposition 10

Under H_0^H ($\gamma_0 = 0$), the outcome equation (22) does not depend on A_i, so $\theta_0^{\text{DML}} = \theta_0^{\text{DAIV}} = \theta_0$. Both estimators converge to θ_0 and are asymptotically normal with variances V^{DML} and V^{DAIV}. Under Neyman orthogonality, $\hat{V}^{\text{DML}} - \hat{V}^{\text{DAIV}} \xrightarrow{p} V^{\text{DML}} - V^{\text{DAIV}} > 0$ (DAIV is efficient under H_0^H since conditioning on A_i is asymptotically costless when $\gamma_0 = 0$). The result then follows from the standard argument for Hausman-type tests based on the difference of two consistent estimators; see Newey and McFadden (1994, Section 7). Under the alternative, $\hat{\theta}^{\mathrm{DML}} - \hat{\theta}^{\mathrm{DAIV}} \xrightarrow{p} \theta_0^{\mathrm{DML}} - \theta_0 \neq 0$, so $|H_n| \to \infty$. Q.E.D.

Appendix B Semiparametric Efficiency Bound

The semiparametric model $\mathscr{M}$ consists of all distributions P satisfying Assumptions 1–4 and 6 with $\theta_0(P)$ defined by (5). We derive the efficient influence function (EIF) by computing the pathwise derivative of the functional $P \mapsto \theta_0(P)$ and applying the Riesz representation theorem (Newey and McFadden 1994).

Step 1: Parametric submodel. For any bounded, mean-zero function $b(Y_i, W_i, X_i, A_i)$, define the submodel P_t with $dP_t/dP = 1 + t\,b + O(t^2)$. The score at $t = 0$ is b.

Step 2: Pathwise derivative. Write $\mu(P) = \mathbb{E}_P[\tilde{Y}_i\tilde{W}_i]$, $\kappa(P) = \mathbb{E}_P[\tilde{W}_i^2]$, $\theta_0(P) = \mu(P)/\kappa(P)$. Under P_t:

$$\left.\frac{d\mu}{dt}\right|_{t=0} = \mathbb{E}_P[\tilde{Y}_i\tilde{W}_i \cdot b] - \theta_0\,\mathbb{E}_P[\tilde{W}_i^2 \cdot b] = \mathbb{E}_P[\psi_i \cdot b],$$

where $\psi_i = \tilde{Y}_i\tilde{W}_i - \theta_0\tilde{W}_i^2$. The derivative of $\kappa(P_t)$ vanishes at $t = 0$ because $\mathbb{E}_P[\psi_i] = 0$ (Step 3 of the proof of Theorem 2), so $d\theta_0/dt\big|_{t=0} = \mathbb{E}_P[\psi_i \cdot b]/\kappa(P)$.

Step 3: Riesz representation. The Riesz representor of $b \mapsto \mathbb{E}_P[\psi_i b]/\kappa(P)$ in $L^2(P)$ is

$$\varphi_i^* = \frac{\psi_i}{\kappa(P)} = \frac{\tilde{Y}_i\tilde{W}_i - \theta_0\tilde{W}_i^2}{\kappa_0}.$$

This is the EIF: $\mathbb{E}_P[\varphi_i^*] = 0$ and $\mathbb{E}_P[\varphi_i^* \cdot b] = (d/dt)\theta_0(P_t)\big|_{t=0}$ for every score b.

Step 4: Variance bound. By the Cramér–Rao theorem for semiparametric models (Newey and McFadden 1994), the asymptotic variance of any regular estimator is bounded below by $\mathbb{E}_P[(\varphi_i^*)^2] = \mathbb{E}_P[\psi_i^2]/\kappa_0^2 = V_0$.

Step 5: Attainment. From the proof of Theorem 7, $\sqrt{n}(\hat{\theta}^{\mathrm{DAIV}} - \theta_0) = \kappa_0^{-1} n^{-1/2} \sum_i \psi_i + o_p(1) = n^{-1/2}\sum_i \varphi_i^* + o_p(1)$, so the asymptotic variance of DAIV equals V_0 and DAIV attains the bound. Q.E.D.

Appendix C Auxiliary Lemmas

Proof of Lemma 1

Fix component k. By assumption, the partial increment $\hat{\Delta}_i^{(k)} = A_i^{(k)}(1, W_{-i}) - A_i^{(k)}(0, W_{-i}) > 0$ on a set of positive probability. Conditional on (X_i, A_i), the event that $W_i = 1$ cannot be inferred from $A_i^{(j)}$ for $j \neq k$ alone (by the non-measurability assumption). Together these imply $P(W_i =$

$1 \mid X_i, A_i) < 1$ on a set of positive measure, so $\mathrm{Var}(W_i \mid X_i, A_i) > 0$ a.s. on that set, giving $\kappa_0 > 0$. Q.E.D.

Proof of Lemma 6

Part (a). The stated L^2 rate for random forests in Hölder classes follows from Wager (2018, Theorem 3.2), adapted to the augmented covariate space $\mathbb{R}^{d_X+d_A}$ of dimension $d = d_X + d_A$. The rate $n^{-s/(2s+d)}$ satisfies $n^{-s/(2s+d)} = o(n^{-1/4})$ if and only if $s > d/4$.

Part (b). The LASSO rate in L^2 under sparsity follows from Wainwright (2019, Theorem 6.1): with p dictionary terms and sparsity s_0, $\|\hat{f} - f_0\|_2 = O_p(\sqrt{s_0 \log p/n})$. This is $o(n^{-1/4})$ iff $s_0 \log p/n = o(n^{-1/2})$, i.e., $s_0 \log p = o(\sqrt{n})$. Q.E.D.

University of Warsaw
Faculty of Economic Sciences
44/50 Długa St.
00-241 Warsaw
www.wne.uw.edu.pl
ISSN 2957-0506