Detecting clinician implicit biases in diagnoses using proximal causal inference

Kara Liu[†], Russ Altman, Vasilis Syrgkanis

Computer Science Department, Stanford University,
Stanford, CA 94305, USA
[†]*E-mail: karaliu@stanford.edu*

Clinical decisions to treat and diagnose patients are affected by implicit biases formed by racism, ableism, sexism, and other stereotypes. These biases reflect broader systemic discrimination in healthcare and risk marginalizing already disadvantaged groups. Existing methods for measuring implicit biases require controlled randomized testing and only capture individual attitudes rather than outcomes. However, the "big-data" revolution has led to the availability of large observational medical datasets, like EHRs and biobanks, that provide the opportunity to investigate discrepancies in patient health outcomes. In this work, we propose a causal inference approach to detect the effect of clinician implicit biases on patient outcomes in large-scale medical data. Specifically, our method uses proximal mediation to disentangle pathway-specific effects of a patient's sociodemographic attribute on a clinician's diagnosis decision. We test our method on real-world data from the UK Biobank. Our work can serve as a tool that initiates conversation and brings awareness to unequal health outcomes caused by implicit biases.[*]

Keywords: Implicit bias, proximal causal inference, fairness, healthcare

1. Introduction

Implicit bias refers to unconscious and automatic associations that affect how we perceive, evaluate, and interact with people from different social groups.[1] Outside of mere cognitive distortions, these biases held by healthcare professionals influence clinical decisions and alter a patient's quality of care. Implicit biases have been shown to be both harmful and pervasive in modern-day medicine, exacerbating existing inequality in the treatment and health outcomes of marginalized groups.[2,3] For instance, unconscious attitudes held by clinicians result in disparate outcomes where women are less likely then men to be diagnosed with myocardial infarction,[3] Black women in the UK and US experience higher maternal mortality than White women,[4] and low socioeconomic (SES) and non-White patients receive sub-optimal pain management treatment compared to high SES and White patients.[5,6]

The recent integration of machine learning (ML) models into clinical decision-making has highlighted the prevalence of biases in medicine. By replicating the patterns from real-world medical data, ML models perpetuate and risk amplifying existing disparities in the medical treatment of marginalized groups.[7,8] While much attention has been given to the statistical

---

objectives of fairness and the development of fair models, there has been comparatively less focus on investigating the biases present in the underlying data. A method capable of detecting implicit clinician bias in observational datasets would prevent ML models from unintentionally perpetuating biased decisions.

However, measuring implicit bias is challenging. Existing methods for quantifying implicit bias rely on the Implicit Association Test (IAT)[9] and randomized psychological experiments like affective priming.[10] While these tests are useful for initiating dialogue, they only provide a snapshot of individual clinician attitudes and do not guarantee a causal link to behavior or larger systemic discrepancies of care.[10]

In this work, we propose a computational tool to detect clinician implicit bias in observational datasets by measuring the causal effect of patient attributes, like race, SES, and other social determinants of health (SDoH), on medical diagnoses. By decomposing the causal effect into two pathways, we can separate the *biological effect* (the influence of a demographic attribute on diagnosis as mediated by valid biological traits) from the *implicit bias effect* (how the patient's attribute affects a clinician's judgement independent of their actual health state). As it is unlikely to observe a patient's true health state, we use observed medical data as proxies using proximal causal inference.[11] To estimate the effect of implicit bias, we propose a novel proximal mediation method that guarantees identifiability under several assumptions. Using real patient data from the UK Biobank, we validate our method can robustly detect several clinician implicit biases identified from prior works. We aim for the proposed method to serve as a bias-detection tool in dataset audits and initiate discussion on reducing systemic discrimination in medicine.

Disclaimer: While we use the UK Biobank data for method validation, we emphasize that this work is not a commentary on specific examples of discrimination within the UK healthcare system. Additionally, it is crucial to clarify that our method of estimating implicit bias is not intended to target clinicians but rather reflect on clinician behaviors within the context of discriminatory healthcare systems.

2. Method

2.1. *Background*

2.1.1. *Overview*

According to the Hippocratic Oath, clinicians should base their diagnostic decisions on each patient's history and current health status, unaffected by biases or stereotypes of the perceived patient identity. However, even in the ideal scenario of unbiased treatment, patient sociodemographic attributes will still influence diagnosis. Attributes including race, sex, or SES have been shown to influence a patient's true health status via mechanisms like genetics, lifestyle, and weathering from systemic oppression.[12–14] These biologically-mediated effects increase the risk of certain medical conditions. For instance, patients from lower SES backgrounds experience higher levels of stress and reduced access to healthcare, increasing their risk of cardiovascular disease.[15] In light of these known biological influences, the causal effect of a patient's sociodemographic attribute on their diagnosis by a clinician is therefore comprised of two pathway effects: the *biological effect* and the *implicit bias effect*, the latter referring to

the clinician's subjective biases of the sociodemographic attribute not mediated through the patient's actual health state.

We present the assumed causal relationships between variables as the directed acyclic graph (DAG) in Figure 1. Dashed arrows denote optional edges, and bi-directional arrows denote indirect confounding paths through latent variables. Let D be the binary sociodemographic attribute and Y the diagnosis decision we wish to measure implicit bias with respect to. M represents the latent variables encoding a patient's true underlying health state. However, as M is typically unknown, we instead observe Z and X as multivariate proxies of M. We differentiate these proxies into the variables Z that do not affect the diagnostic decision Y but could be affected by the attribute D; and the variables X which are not directly affected by the attribute D but can influence diagnosis Y. For example, X could be recent lab reports a clinician uses to make their diagnosis, and Z could be a patient's survey responses to a sleep questionnaire (assuming the survey does not influence the clinician's diagnosis). Finally, let W be sociodemographic confounders to control for.

We can now reframe *biological* and *implicit bias effects* using pathway causal effects. The *biological effect* of attribute D on diagnosis Y is the indirect effect as mediated through the true underlying health state M: $D \to M \to Y$. The *implicit bias effect* we wish to measure is the direct effect of $D \to Y$ that flows through the edge θ and is defined as the residual of the biological effect. We formally define bias in terms of controlled direct effects in Equation (1).

D: Attribute
M: Health state mediator
Z, X: Health proxies
Y: Medical outcome
W: Sociodemographic confounders
θ: Implicit bias effect
$\longrightarrow$: Causal effect
$\dashrightarrow$: Optional causal effect
$\longleftrightarrow$: Bidirectional causal effect
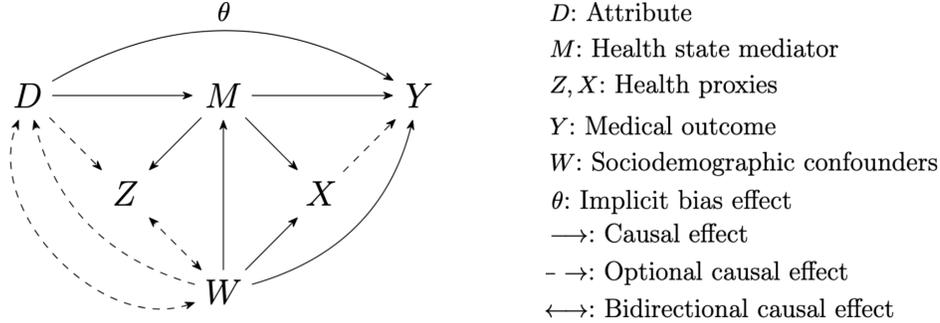
Fig. 1: Assumed causal graph.

2.1.2. *Related works*

Measuring implicit biases requires detecting the unconscious and automatic attitudes that shape behavior. The predominant method for implicit bias measurement thus far has been the Implicit Association Test (IAT),[9] a questionnaire developed in 1998 intended to measure group association through word categorization. To capture clinician biases, several works have linked clinician attitudes via their IAT score to behavioral manifestation.[10,16] Other methods for detecting implicit clinician bias include affective priming, which measures biased associations after stimulus priming; and the assumption method, which surveys clinicians' decisions after reading patient vignettes.[17] While association tests like the IAT have been integral in bringing awareness to medical biases, they are criticized for their arbitrary scoring system, inability to predict real-world patient outcomes, and context-dependency.[3,10,16,17] Furthermore, administering these controlled tests in every clinical encounter is impractical and unscalable.

Computational methods present a promising and scalable alternative for detecting implicit bias in real-world medical data. While the field of ML fairness has explored bias detection, the focus has been on identifying and mitigating bias in models rather than the data.[7] In causal inference, disentangling a causal effect into natural indirect and direct pathway effects has led to methods that control for "fair" and "unfair" causal pathway effects. [18–21] propose metrics for measuring fair pathway influence on outcomes and develop methods that mitigate the effect of unfair pathways on the predicted outcome. [22] leveraged the Fairness-Aware Causal paThs (FACTS)[23] algorithm to quantify disparate pathway influence of SDoH attributes on mortality using real-world health data. While these methods recognize an attribute's influence on an outcome contains both fair and unfair effects, prior works are limited to simple scenarios where all variables are known and observed. Our work is the first to extend pathway inference to large-scale observational data with potentially unobserved variables.

Finally, a few recent methods have explored proximal mediation analysis, where pathway effects can be measured despite unobserved mediators by using proxy variables.[24,25] However, by relying on natural direct and indirect pathway effects, these works rely on more stringent assumptions, require learning complicated bridge functions, and limit their analysis to simple datasets. In comparison, our method makes several relaxations that enable application to observational data. First, we identify controlled instead of natural effects, which presents an equally good measurement of a biased decision yet lends to a much simpler statistical problem. Additionally, we assume partially linear equations instead of requiring the identification of a complex bridge function. Finally, we do not require uniqueness of the parameters unrelated to implicit bias (i.e., the nuisance parameters for the outcome bridge function). These relaxations enable our approach to be effective at analyzing large-scale real-world medical data.

2.2. *Our method*

Our goal is to identify and estimate the following controlled direct effect:

$$\theta = \int_{m,w} \mathbb{E}[Y(1,m) - Y(0,m) \mid W = w] \; p(m,w) \; dm \; dw \tag{1}$$

where $Y(d,m)$ is the potential (or counterfactual) outcome when we intervene on the attribute D and the mediator M and set them to values (d,m); and $p(m,w)$ is the natural probability distribution in the data. If the controlled direct effect is nonzero, then there exists a direct influence of the attribute D on the outcome Y, which is evidence of implicit bias.

If we observe M, the above controlled direct effect can be identified by a simple g-formula that "controls" for M and W: $\theta = \mathbb{E}[\mathbb{E}[Y \mid D = 1, M, W] - \mathbb{E}[Y \mid D = 0, M, W]]$. Unfortunately, this equation is intractable if M is unobserved. However, we show that under a few reasonable assumptions the controlled direct effect is still identifiable.

Theorem 1 (Identification). [a] *Consider a non-parametric structural causal model (SCM) that respects the causal relationships encoded in Figure 1 (see Appendix D.1) and assume there exists a "bridge function" q that solves $\mathbb{E}[Y \mid D, M, W] = \mathbb{E}[q(D, X, W) \mid D, M, W]$. Then q also solves the Non-Parametric Instrumental Variable (NPIV) problem defined by the set of*

---

[a]We present more intuitive interpretations of each theorem and lemma in the Appendix.

conditional moment restrictions

$$\mathbb{E}[Y - q(D, X, W) \mid D, Z, W] = 0 \tag{2}$$

and the controlled direct effect can be identified as $\theta = \mathbb{E}[q(1, X, W) - q(0, X, W)]$.

Identifying parameters θ using a bridge function q (where q also solves an NPIV problem) has been extensively studied in proximal causal inference literature.[26–38] However, these approaches rely on solving saddle-point problems with adversarial training or require learning conditional density functions, both of which are statistically daunting.

We can avoid these difficult statistical tasks if we assume that the bridge function is partially linear in D and X. The following lemma shows that partial linearity of q is implied by a more primitive assumption of partial linearity of two other functions (proof in Appendix D.4).

Lemma 1 (Identification under partial linearity). *Consider a non-parametric SCM that respects the constraints encoded in Figure 1 and assume that X has dimension p_X at least as large as the dimension p_M of M. Moreover, assume that the following functions are partially linear:*

$$\mathbb{E}[Y \mid D, M, X, W] = D\,c + M^\top b + X^\top g + f_Y(W) \tag{3}$$
$$\mathbb{E}[X \mid M, W] = F\,M + f_X(W) \tag{4}$$

where F is a $p_X \times p_M$ matrix, b is a p_M-dimensional vector, g is a p_X-dimensional vector and f_Y, f_X are arbitrary non-parametric functions. If we assume the matrix F has full column rank, then there exists a partially linear outcome bridge function

$$q(D, X, W) = D\,\theta + X^\top h + f(W) \tag{5}$$

that satisfies Equation (2), where parameter $h = F^+ b + g$[b] and $\theta = c$.

Under the assumption of partial linearity, we can simplify the estimation problem by first removing the effect of W from all the remaining variables (see Appendix D.5), where for any variable V we define the residual $\tilde{V} = V - \mathbb{E}[V \mid W]$. Partial linearity of q from Equation (5), when combined with the NPIV Equation (2), implies that θ can be identified using linear instrumental variable (IV) regression where $(\tilde{Z}; \tilde{D})$[c] are the instruments and $(\tilde{X}; \tilde{D})$ are the treatments:

$$\mathbb{E}\left[(\tilde{Y} - \tilde{X}^\top h - \tilde{D}\,\theta) \begin{pmatrix} \tilde{Z} \\ \tilde{D} \end{pmatrix}\right] = 0 \qquad \text{(Primal Equation)}$$

Unique identification of θ seemingly requires unique identification of the other "nuisance" parameters like h, which might be difficult to achieve as the covariance matrix $\mathbb{E}[(\tilde{X}; \tilde{D})\,(\tilde{Z}; \tilde{D})^\top]$ is usually not full rank[d]. We invoke and simplify ideas from the recent proximal inference literature[35,39] to show that θ can be point-identified even if h is not. To achieve this, we

---

[b] F^+ is the Moore-Penrose pseudoinverse of F.

[c] We denote $(A; B)$ to be concatenation of vectors A and B.

[d] This could be the case if the number of proxies is much larger than the dimensionality of the latent mediator M.

construct a moment restriction equation that is Neyman orthogonal to the nuisance parameters h but still point-identifies θ, given sufficient quality of the proxy Z. Intuitively, we learn a new instrument $V = (\tilde{D} - \gamma^\top \tilde{Z})$ such that V is uncorrelated with $\tilde{X}$ and thus estimation of θ is not sensitive to h. Existence of such a γ is sufficient for point-identification of θ. We provide the proof for the point identification of θ in Appendix D.7 and for Neyman orthogonality in D.8.

Theorem 2. *Let h_* be the minimum norm solution to the* (Primal Equation) *and assume that the following dual equation also admits a solution γ_*:*

$$\mathbb{E}[\tilde{X}\,(\tilde{D} - \gamma^\top \tilde{Z})] = 0 \qquad\qquad \text{(Dual Equation)}$$

Furthermore, assume $\mathbb{E}[\tilde{D}\,(\tilde{D} - \gamma_*^\top Z)] \neq 0$. Then the solution θ_0 to the equation:

$$\mathbb{E}[(\tilde{Y} - \tilde{X}^\top h_* - \tilde{D}\,\theta)\,(\tilde{D} - \gamma_*^\top \tilde{Z})] = 0 \qquad\qquad (6)$$

uniquely identifies the controlled direct effect θ. Furthermore, this moment restriction is Neyman orthogonal with respect to nuisance parameters γ_*, h_*.

Theorem 2.2 allows us to invoke the general framework of [40] to construct an estimate and confidence interval for the controlled direct effect θ. The full estimation algorithm is presented in Appendix D.9.

2.3. *Testing and Removing Weak Instruments*

Our method for uniquely identifying the controlled direct effect θ relies on several assumptions, e.g., $(\tilde{Z}; \tilde{D})$ are good instruments for $(\tilde{X}; \tilde{D})$. To assess the validity of these assumptions, we developed a suite of tests that must pass for the estimate θ to be valid and can be used as validity checks by practitioners. These tests are further described in Appendix C:

(1) *Primal equation violation* - We develop a χ^2-test to check if the primal equation admits a solution, i.e., $\mathbb{E}[(\tilde{Y} - \tilde{X}^\top h_* - \tilde{D}\,\theta_0)(\tilde{Z}; \tilde{D})] \approx 0$. Intuitively, violation of the primal test implies either the variables X are insufficient proxies of the health state M or the residual proxy $\tilde{Z}$ has a direct path to $\tilde{Y}$.

(2) *Dual equation violation* - We develop a χ^2-test to check if the dual equation admits a solution, i.e., $\mathbb{E}[\tilde{X}(\tilde{D} - \gamma_*^\top \tilde{Z})] \approx 0$. Violation of the dual implies the variables Z are insufficient proxies of the health state M or that the residual proxy $\tilde{X}$ has a direct path from $\tilde{D}$.

(3) *Strength of identification* - We perform two tests to check if $V = (\tilde{D} - \gamma_*^\top \tilde{Z})$ is a good instrument for (i.e., retains enough information about) $\tilde{D}$. (a) We develop an effective F-test[41,42] to check the correlation strength of V with $\tilde{D}$. (b) We develop a z-test to check if the quantity $\mathbb{E}[\tilde{D}(\tilde{D} - \gamma_*^\top \tilde{Z})]$ is substantially bounded away from zero (see assumption in Theorem 2.2). Intuitively, these tests will fail if the hidden mediator is a very deterministic function of the attribute D.

(4) *Proxy covariance rank test* - To ensure the health proxies are sufficiently related, we check the rank of the covariance matrix of $\tilde{X}$ and $\tilde{Z}$ by identifying the number of statistically significant singular values. This rank can be viewed as an upper bound on the dimensionality of the hidden mediator M that we can control for.

2.3.1. *Proxy selection algorithm*

In practice, the initial selection of proxies X, Z may violate key assumptions, which can be detected by the failure of one or more of the aforementioned tests. In Appendix B, we provide an algorithm for identifying subsets of X and Z that satisfy the necessary assumptions and thus produce valid estimates. This proxy selection algorithm should be performed on a separate dataset from the one used to estimate θ.

3. Experiments

3.1. *Data*

To validate our approach, we use the UK Biobank, a rich and accessible repository containing genomic, imaging, and tabular health data from over 500,000 patients. Our work uses its tabular data, which includes survey questions and biometrics collected upon an individual's enrollment into the biobank. In addition, several health outcomes, including medical diagnoses via ICD10 codes, have been linked to most patients. We note and discuss the caveats of applying our method to biobank data in Section 5.2.

		Prevalence in UK Biobank (n=502411)	Prior works on implicit bias
Sociodemographic attribute D	*Race* - Asian	2.4%	43–45
	Race - Black	1.8%	3–5,46,47
	Gender - Female	54.4%	3,48
	Disability status - On disability allowance	6.2%	49,50
	Income - Household income <18,000£	20.3%	5,14,15,51
	Education - No post-secondary education	67.3%	5,51
	Weight - BMI >30	24.3%	52,53
	Insurance - Not on private insurance	31.4%	54
Medical diagnosis Y	Osteoarthritis	18.0%	47,49
	Rheumatoid arthritis	1.9%	55
	Chronic kidney disease	5.0%	56,57
	Complications during labor	2.4%	3,4
	Heart disease	10.7%	3,15,48
	Depression	6.0%	46,58
	Melanoma	1.2%	59,60

Table 1: Selected sociodemographic attributes D and diagnoses Y

Prior works have proposed sociodemographic attributes that might bias clinical decisions. For example, [48] showed that clinicians exhibited greater uncertainty when diagnosing coronary heart disease in women compared to men. We list in Table 1 most of the attributes D and diagnoses Y we test for implicit bias, and present the full list of the 102 (D, Y) pairs in Appendix E.2. To highlight the influence of clinician subjectivity, we concentrate on diagnoses that require clinician interpretation of patient-reported symptoms, e.g., chronic pain.

Selecting health proxies for Z and X relies user intuition and medical expertise to determine which variables have a direct relationship with attribute D and outcome Y, respectively. In general, proxies X could be observed by the clinician during their diagnostic decision, and proxies Z are not accessible during diagnosis but might have a direct causal relationship with attribute D. In the UK Biobank, we select X to be the biometric variables collected by the biobank at patient enrollment, which includes lab results and blood pressure readings. For Z, we use survey responses of self-reported pain levels, mental health, and sleep. We list all variables, including the sociodemographic confounders W, in Appendix E.1. Note our data contains a mix of binary, integer, and continuous variable types.

3.2. *Evaluation metrics*

3.2.1. *Semi-synthetic data validation*

We test if our method can retrieve a known implicit bias effect using semi-synthetic data. We use real data from the UK Biobank for attribute D, confounders W, and health proxies X,Z. We develop a model that computes M and a synthetic diagnosis Y with a known implicit bias effect θ_0 using linear structural equations. We test against fully continuous (Experiment 1) and both binary and continuous (Experiment 2) semi-synthetic data, the latter being more realistic in real-world medical data. Our semi-synthetic data generation method is fully described in Appendix A. As a baseline, we compare two variants of ordinary least squares (OLS): (a) given we know M, we fit an OLS model over W, D, X, and M to predict Y; (b) in the more realistic scenario where M isn't known, we learn over W, D, X, and Z. We compute the average effect estimate $\bar{\theta}$ and confidence interval based on $\pm 1.96\ \sigma$ where the average and standard deviation σ is taken over $K = 100$ iterations.

3.2.2. *Calculating the implicit bias effect in the UK Biobank*

We next run our method on the full UK Biobank data. We compute the residuals of Z, X, Y, D fitted on W using Lasso regression. For all models, the regularization term is chosen via semi-cross fitting[61,62] over 3 splits. We fit all models using the `scikit-learn` Python package. For nuisance parameters h_* and γ_* we used regularized adversarial IV estimation[35,63] with linear functions and a theoretically driven penalty choice that decays faster than the root of the number of samples.

In cases where the data may not meet the method's assumptions, we developed a proxy selection algorithm (see Section 2.3.1) that identifies an optimal subset of X, Z proxies for each (D, Y) pair using the assumption tests from Section 2.3. Although we recommend separate data splits for proxy selection and effect estimation, we use the same dataset as our intent is method demonstration rather than robust effect estimates. Details of the hyperparameters used for the selection algorithm are provided in Appendix B.

For each of the 102 pairs of attribute D and diagnosis Y, we report seven metrics: the implicit bias effect θ, the 95% confidence interval, as well as our five proposed tests from Section 2.3: (1) the primal and (2) dual violation, (3-4) the strength of identification, and (5) the $\tilde{Z}, \tilde{X}$ covariance rank test. In addition, we also run the following five analyses:

Weak identification confidence interval - If the instrument identification tests from 2.3 are violated, then effect estimation can be unstable and normality-based confidence intervals inaccurate. We thus compute an alternative confidence interval[62] developed under the assumption of weak instruments (see Appendix C.5 for the description).

Bootstrapping analyses - We perform several bootstrapping analyses to test the sensitivity of the estimate. In the first analysis, given the computational complexity of recomputing the full estimate, we compare K=10 bootstrapped iterations re-estimating the full pipeline (stage 1); K=100 iterations using the pre-computed residuals but re-estimating all other parameters (stage 2); and K=1000 iterations re-computing only the final Equation (6) (stage 3). Each

iteration samples 50% of the data without replacement. In the second analysis, we compare sampling 10%, 25%, 50% or 75% of the original data for K=10 bootstrapped iterations, re-estimating over full pipeline (stage 1). Finally, we compare different sample sizes for $K = 1000$ iterations re-estimating from stage 3 of the pipeline.

Influence points - Inspired by [64], we analyze influence scores, which measure how influential each data point is in the effect estimate. A significant change to the estimate after removing a small set of highly-influential points indicates the implicit bias calculation is highly sensitive to a few (potentially) outlier patients. We also include a preliminary interpretability analysis that explores the distinguishing phenotypes of highly influential patients, which could aid in determining if these subsets of patients correspond to some interpretable outlier group. We describe how we calculate the influence score and identify highly-influential patient sets in Appendix C.6.

Income stratification - To investigate intersectionality in implicit biases, we perform a stratified effect estimate over different income groups where $D \neq$Income.

Partial non-linearity of W - Our identification theorem allows for partial non-linearity in the effect of W. We thus re-compute the point estimate allowing for non-linear interactions with W using XGBoost[65] models instead of Lasso.

4. Results

4.1. *Synthetic data validation*

The results in Table 2 demonstrate that our method is able to retrieve the true implicit bias effect $\theta_0 = 0.5$ with high certainty for both fully continuous and mixed-type data, with comparable performance to the best-case OLS where M is known. We report our method's coverage, RMSE, bias, standard deviation, mean confidence interval, and performance on our five tests (from Section 2.3), as well as testing other values of θ, in Appendix F.1.

	θ_0	**Our method**	**OLS**(D, W, M, X)	**OLS**(D, W, Z, X)
Experiment 1: Continuous	0.5	0.54 ± 0.003	0.5 ± 0.01	1.10 ± 0.01
Experiment 2: Continuous and binary	0.5	0.53 ± 0.003	0.5 ± 0.01	1.385 ± 0.01

Table 2: Semi-synthetic data estimates $\bar{\theta} \pm 1.96\sigma$ over $K = 100$.

4.2. *Calculating the implicit bias effect in the UK Biobank*

In Appendix F.2, we show the effect estimates for the (D, Y) pairs using all proxies Z, X, adjusting the confounders W by excluding the column corresponding to the attribute D. However, as evidenced by the failure of the dual and primal tests, we found the initial sets of proxies Z, X did not meet our method's necessary assumptions. As discussed further in Appendix F.3, we believe these test failures indicate there might exist some features in X with a causal path from D that does not go through M or features within Z with a causal path to Y that doesn't flow through M. Such paths invalidate the resulting effect estimates.

We thus found applying our proxy selection algorithm (see 2.3.1) necessary for producing valid

effect estimates. After running the algorithm to select subsets of admissible X, Z proxies (the description and interpretation of the selected proxies can be found in Appendix F.3), we found 34 (D, Y) pairs that pass all tests with narrow confidence intervals. We report six in Table 3 and include the remaining estimates in Appendix F.4. Note that $\theta > 0$ implies a patient with D is more likely to be diagnosed with Y due to clinician bias, and conversely $\theta < 0$ implies a patient is less likely to be diagnosed. In Section 5.2, we offer a framework for interpreting the implications of these results.

4.2.1. *Weak instrument confidence interval*

As shown in Figure 2A, the confidence interval predicted under the weak instrument regime consistently aligns with the interval under our method, thus indicating our estimate's robustness to weak instruments.

(D, Y)	$\theta \pm 95\%$ CI	(1) Primal statistic < critical	(2) Dual statistic < critical	(3) $\mathbb{E}[\tilde{D}V] \neq 0$ statistic > critical	(4) V strength F-test statistic > critical	(5) Cov$(\tilde{X}, \tilde{Z})$ rank
Low income, Depression	0.03 ± 0.02	59.9<60.5	31.9<40.1	84.1>0.4	3332.1>23.1	3
Disability insurance, Rh. Arthritis	0.06 ± 0.0	67.3<75.6	3.4<11.1	29.2>0.4	801.1>23.1	3
Female, Heart disease	-0.19 ± 0.06	115.8<118.8	23.3<23.7	18.8>1.3	92.5>23.1	4
Black, Chronic kidney disease	0.14 ± 0.03	56.9<58.1	10.6<21.0	9.8>0.3	23.3>23.1	4
Obese, Osteoarthritis	0.09 ± 0.02	90.5<100.7	24.8<28.9	76.5>1.7	254.9>23.1	3
Asian, Osteoarthritis	-0.06 ± 0.03	94.7<101.9	33.1<33.9	13.9>0.3	74.6>23.1	5

Table 3: Six of the 34 valid UK Biobank implicit bias effect estimates after applying our X, Z proxy selection algorithm. Tests (1-5) are detailed in 2.3, where *statistic* is the given data's statistic and *critical* is the necessary critical value to be greater or less than to pass. $V = \tilde{D} - \gamma^\top \tilde{Z}$.
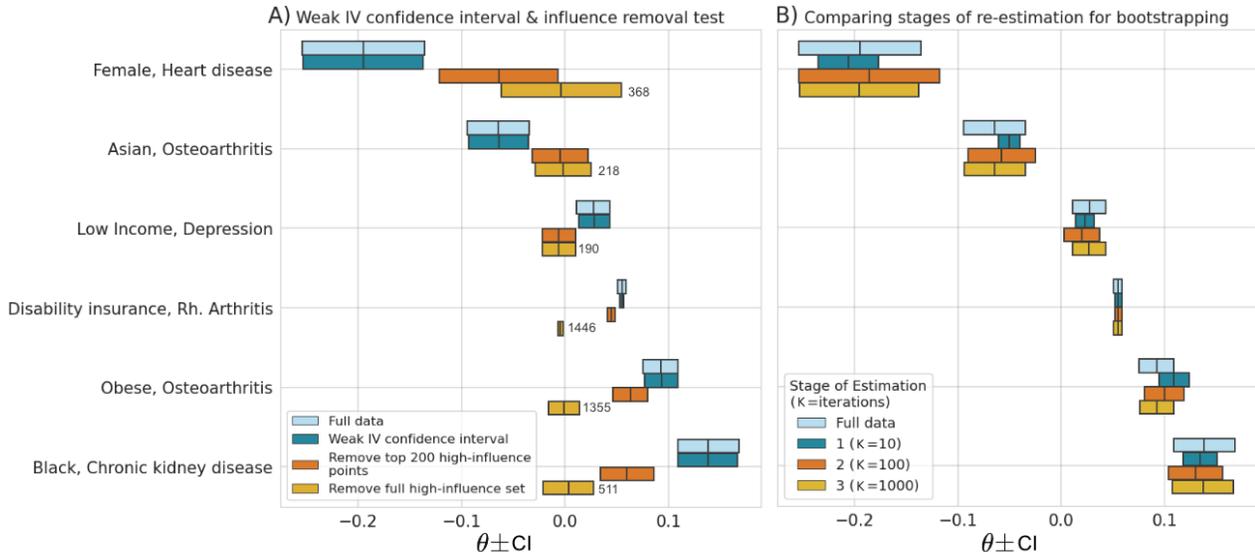
Fig. 2: Comparing effect estimates for six (D, Y) pairs using all data with: A) weak instrument and influence set removal (where the numbers next to the yellow bar reflect the set size of high-influence points); B) bootstrapped subsampling 50% of the data at different stages of re-estimation.

4.2.2. *Bootstrapping analyses*

In Figure 2B, we show the results of the first bootstrap analysis comparing different stages of re-estimation. We observe that, regardless of the estimation stage, bootstrapped estimates are consistent with the estimate from the full dataset. The consistency of the bootstrapped estimates over different sample sizes, as shown in Appendix F.6, further support the robustness of our method.

4.2.3. *Influence points*

In Figure 2A, we see that removing only a few highly-influential points leads to a significant decrease in the magnitude of the estimated effect. To investigate, we run a preliminary interpretabilty analysis where we analyze the univariate differences between patients with high influence and those with low influence. In Figure 3A patients that strongly influence the negative implicit bias estimate for (D=Female, Y=Heart disease) are more likely to be low income, unemployed due to disability, and suffer from depression. It is plausible such patients are the "outliers" driving the strong negative bias estimate.
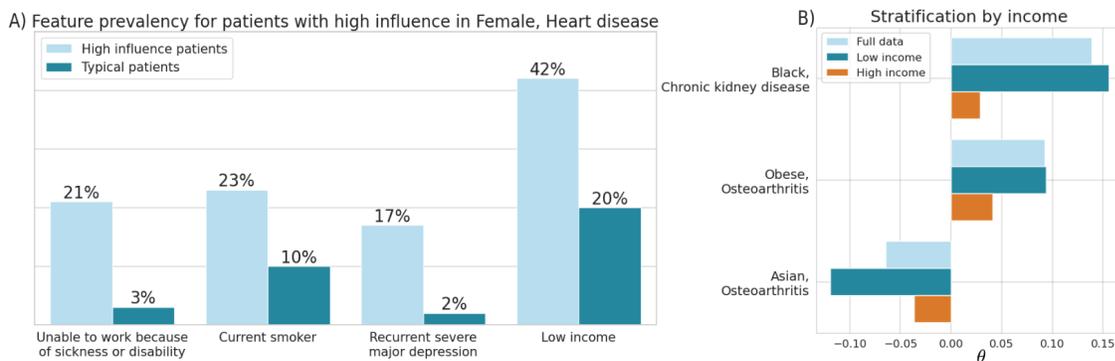
Fig. 3: A) Interpretability into high influence points. B) Income stratification

4.2.4. *Income stratification*

In Figure 3B we analyze the effect of stratification based on income. We see a general increase in bias effect estimate for the low income strata and a corresponding decrease in effect for high income strata, demonstrating potential evidence of intersectional discrimination.[5,66]

4.2.5. *Partial non-linearity of W*

In Appendix F.9, we show our implicit bias estimate with non-linear W interactions leads to a similar effect estimates of θ.

5. Discussion

5.1. **Limitations**

In this work, we propose a robust causal inference method designed to detect clinician implicit bias by estimating pathway-specific causal effects. We demonstrate the applicability of our approach to large-scale medical data by validating on both semi-synthetic and real-world datasets.

However, our work contains several limitations. First, while the UK Biobank is a rich and accessible source of medical data, most patient information is collected once upon signing up for the biobank. Although UK Biobank has synced their records to a handful of outcomes provided by EHR data (like ICD10 codes), it is unclear to what extent the available proxies for X (which were collected at patient enrollment) are used by clinicians for diagnoses. Additionally, the synced ICD10 codes are from hospital records, thus excluding primary care visits. We plan to validate our method with time-series EHR data in follow-up work.

Second, while the assumption of partially linear structural equations is crucial for enabling better identifiability of the outcome bridge function under minimal conditions, it is possible the ground truth equations are non-linear.

Finally, it is well known that intersectional identities shape complex patterns of discrimination in healthcare.[5,66] A more comprehensive analysis on the effect of implicit bias from intersectional attributes on patient treatment would be valuable for improving equity in healthcare outcomes.

5.2. *Interpretation and application of results*

While we re-iterate the intent of this work is not to diagnose specific cases of implicit bias in the UK Biobank, our method did flag several areas of clinical inequity that have been reported in literature. For instance, many works have reported gender-based inequality in cardiovascular health,[67] and we similarly detected an estimate of θ = -0.19 indicating clinicians are less likely, due to implicit biases, to diagnose D=Females with Y=heart disease. In another example, our estimate θ = -0.06 suggested clinicians are less likely to diagnose D=Asian patients with Y=osteoarthritis, and many works have highlighted both patient- and clinician-stigmas regarding pain-associated disorders, like osteoarthritis, in Asians.[68–70]

However, we did find several estimates contrary to what we expected. For example, our estimate θ = 0.14 indicated clinicians are *positively* biased towards diagnosing Black patients with chronic kidney disease. However, at the time of UK Biobank data collection, many doctors relied on a race-based equation for kidney function now known to have under-detected kidney disease in Black patients.[71]

To understand a discrepancy between a produced estimate and literature (or user intuition), we recommend (1) ensuring the data used contains sufficient health proxies and satisfy all assumptions (e.g., see biobank data limitations in 5.1); (2) investigating all mechanisms creating the medical outcome Y (e.g., hospital-specific diagnosis protocol); and (3) exploring how the discovered bias estimate fits in context, rather than opposed, to those found in literature. While our method does not offer a solution on *how* to tackle implicit biases, by bringing awareness to potential areas of discrimination within a given healthcare system, detecting biases is the first step towards creating systemic-level change through interdisciplinary collaboration and targeted anti-bias training programs.

6. Appendix

The appendix can be found at `https://github.com/syrgkanislab/hidden_mediators`.

References

1. J. Holroyd, J. Sweetman, M. Brownstein and J. Saul, The heterogeneity of implicit bias, *Implicit bias and philosophy* **1**, 80 (2016).
2. M. B. Vela, A. I. Erondu, N. A. Smith, M. E. Peek, J. N. Woodruff and M. H. Chin, Eliminating explicit and implicit biases in health care: evidence and research needs, *Annual review of public health* **43**, 477 (2022).
3. D. P. Gopal, U. Chetty, P. O'Donnell, C. Gajria and J. Blackadder-Weinstein, Implicit bias in healthcare: clinical practice, research and decision making, *Future healthcare journal* **8**, 40 (2021).
4. B. Saluja and Z. Bryant, How implicit bias contributes to racial disparities in maternal morbidity and mortality in the united states, *Journal of women's health* **30**, 270 (2021).
5. T. M. Anastas, M. M. Miller, N. A. Hollingshead, J. C. Stewart, K. L. Rand and A. T. Hirsh, The unique and interactive effects of patient race, patient socioeconomic status, and provider attitudes on chronic pain care decisions, *Annals of Behavioral Medicine* **54**, 771 (2020).
6. J. A. Sabin and A. G. Greenwald, The influence of implicit bias on treatment recommendations for 4 common pediatric conditions: pain, urinary tract infection, attention deficit hyperactivity disorder, and asthma, *American journal of public health* **102**, 988 (2012).
7. D. Ueda, T. Kakinuma, S. Fujita, K. Kamagata, Y. Fushimi, R. Ito, Y. Matsui, T. Nozaki, T. Nakaura, N. Fujima *et al.*, Fairness of artificial intelligence in healthcare: review and recommendations, *Japanese Journal of Radiology* **42**, 3 (2024).
8. S. R. Pfohl, A. Foryciarz and N. H. Shah, An empirical characterization of fair machine learning for clinical risk prediction, *Journal of biomedical informatics* **113**, p. 103621 (2021).
9. A. G. Greenwald, D. E. McGhee and J. L. Schwartz, Measuring individual differences in implicit cognition: the implicit association test., *Journal of personality and social psychology* **74**, p. 1464 (1998).
10. S. A. Arif and J. Schlotfeldt, Gaps in measuring and mitigating implicit bias in healthcare, *Frontiers in Pharmacology* **12**, p. 633565 (2021).
11. E. J. T. Tchetgen, A. Ying, Y. Cui, X. Shi and W. Miao, An introduction to proximal causal learning (2020).
12. A. T. Forde, D. M. Crookes, S. F. Suglia and R. T. Demmer, The weathering hypothesis as an explanation for racial disparities in health: a systematic review, *Annals of epidemiology* **33**, 1 (2019).
13. D. R. Williams, Stress and the mental health of populations of color: Advancing our understanding of race-related stressors, *Journal of health and social behavior* **59**, 466 (2018).
14. X. Cui and C.-T. Chang, How income influences health: decomposition based on absolute income and relative income effects, *International Journal of Environmental Research and Public Health* **18**, p. 10738 (2021).
15. A. M. K. Minhas, V. Jain, M. Li, R. W. Ariss, M. Fudim, E. D. Michos, S. S. Virani, L. Sperling and A. Mehta, Family income and cardiovascular disease risk in american adults, *Scientific reports* **13**, p. 279 (2023).
16. A. G. Greenwald, N. Dasgupta, J. F. Dovidio, J. Kang, C. A. Moss-Racusin and B. A. Teachman, Implicit-bias remedies: Treating discriminatory bias as a public-health problem, *Psychological Science in the Public Interest* **23**, 7 (2022).
17. C. FitzGerald and S. Hurst, Implicit bias in healthcare professionals: a systematic review, *BMC medical ethics* **18**, 1 (2017).
18. S. Chiappa, Path-specific counterfactual fairness, *Proceedings of the AAAI conference on artificial intelligence* **33**, 7801 (2019).
19. R. Nabi and I. Shpitser, Fair inference on outcomes, *Proceedings of the AAAI Conference on Artificial Intelligence* **32** (2018).

20. L. Zhang, Y. Wu and X. Wu, A causal framework for discovering and removing direct and indirect discrimination (2016).
21. A. I. Naimi, M. E. Schnitzer, E. E. Moodie and L. M. Bodnar, Mediation analysis for health disparities research, *American journal of epidemiology* **184**, 315 (2016).
22. I. Jun, S. E. Ser, S. A. Cohen, J. Xu, R. J. Lucero, J. Bian and M. Prosperi, Quantifying health outcome disparity in invasive methicillin-resistant staphylococcus aureus infection using fairness algorithms on real-world data (2023).
23. W. Pan, S. Cui, J. Bian, C. Zhang and F. Wang, Explaining algorithmic fairness through fairness-aware causal path decomposition (2021).
24. A. Ghassami, A. Yang, I. Shpitser and E. T. Tchetgen, Causal inference with hidden mediators (2021).
25. A. Ghassami, I. Shpitser and E. T. Tchetgen, Partial identification of causal effects using proxy variables (2023).
26. X. Chen and D. Pouzo, Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals, *Econometrica* **80**, 277 (2012).
27. C. Ai and X. Chen, Efficient estimation of models with conditional moment restrictions containing unknown functions, *Econometrica* **71**, 1795 (2003).
28. C. Ai and X. Chen, The semiparametric efficiency bound for models of sequential moment restrictions containing unknown functions, *Journal of Econometrics* **170**, 442 (2012).
29. G. Lewis and V. Syrgkanis, Adversarial generalized method of moments (2018).
30. N. Dikkala, G. Lewis, L. Mackey and V. Syrgkanis, Minimax estimation of conditional moment models, *Advances in Neural Information Processing Systems* **33**, 12248 (2020).
31. A. Bennett, N. Kallus and T. Schnabel, Deep generalized method of moments for instrumental variable analysis, *Advances in neural information processing systems* **32** (2019).
32. A. Bennett and N. Kallus, The variational method of moments (2020).
33. W. Miao and E. T. Tchetgen, A confounding bridge approach for double negative control inference on causal effects (supplement and sample codes are included) (2018).
34. Y. Cui, H. Pu, X. Shi, W. Miao and E. T. Tchetgen, Semiparametric proximal causal inference (2020).
35. A. Bennett, N. Kallus, X. Mao, W. Newey, V. Syrgkanis and M. Uehara, Inference on strongly identified functionals of weakly identified functions (2023).
36. A. Bennett, N. Kallus, X. Mao, W. Newey, V. Syrgkanis and M. Uehara, Source condition double robust inference on functionals of inverse problems (2023).
37. A. Bennett, N. Kallus, X. Mao, W. Newey, V. Syrgkanis and M. Uehara, Minimax instrumental variable regression and l_2 convergence guarantees without identification or closedness (2023).
38. J. Zhang, W. Li, W. Miao and E. T. Tchetgen, Proximal causal inference without uniqueness assumptions, *Statistics & Probability Letters* **198**, p. 109836 (2023).
39. Q. Chen, Robust and optimal estimation for partially linear instrumental variables models with partial identification, *Journal of Econometrics* **221**, 368 (2021).
40. V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen and W. Newey, Double/debiased/neyman machine learning of treatment effects, *American Economic Review* **107**, 261 (2017).
41. J. L. M. Olea and C. Pflueger, A robust test for weak instruments, *Journal of Business & Economic Statistics* **31**, 358 (2013).
42. D. W. Andrews and J. H. Stock, *Inference with Weak Instruments*, Working Paper 313, National Bureau of Economic Research (August 2005).
43. C. L. McMurtry, M. G. Findling, L. S. Casey, R. J. Blendon, J. M. Benson, J. M. Sayde and C. Miller, Discrimination in the united states: Experiences of asian americans, *Health services research* **54**, 1419 (2019).

44. O. Bougie, M. I. Yap, L. Sikora, T. Flaxman and S. Singh, Influence of race/ethnicity on prevalence and presentation of endometriosis: a systematic review and meta-analysis, *BJOG: An International Journal of Obstetrics & Gynaecology* **126**, 1104 (2019).

45. C. Wu, Y. Qian and R. Wilkes, Anti-asian discrimination and the asian-white mental health gap during covid-19 (2021).

46. H. N. Garb, Race bias and gender bias in the diagnosis of psychological disorders, *Clinical Psychology Review* **90**, p. 102087 (2021).

47. J. L. W. Taylor, C. M. Campbell, R. J. Thorpe Jr, K. E. Whitfield, M. Nkimbeng and S. L. Szanton, Pain, racial discrimination, and depressive symptoms among african american women, *Pain Management Nursing* **19**, 79 (2018).

48. N. N. Maserejian, C. L. Link, K. L. Lutfey, L. D. Marceau and J. B. McKinlay, Disparities in physicians' interpretations of heart disease symptoms by patient gender: results of a video vignette factorial experiment, *Journal of women's health* **18**, 1661 (2009).

49. J. McClendon, U. R. Essien, A. Youk, S. A. Ibrahim, E. Vina, C. K. Kwoh and L. R. Hausmann, Cumulative disadvantage and disparities in depression and pain among veterans with osteoarthritis: the role of perceived discrimination, *Arthritis Care & Research* **73**, 11 (2021).

50. L. VanPuymbrouck, C. Friedman and H. Feldner, Explicit and implicit disability attitudes of healthcare providers., *Rehabilitation psychology* **65**, p. 101 (2020).

51. I. Stepanikova and G. R. Oates, Perceived discrimination and privilege in health care: the role of socioeconomic status and race, *American journal of preventive medicine* **52**, S86 (2017).

52. S. M. Phelan, D. J. Burgess, M. W. Yeazel, W. L. Hellerstedt, J. M. Griffin and M. van Ryn, Impact of weight bias and stigma on quality of care and outcomes for patients with obesity, *Obesity Reviews* **16**, 319 (2015), Open Access, Citations: 737.

53. M. Fulton, S. Dadana and V. N. Srinivasan, Obesity, stigma, and discrimination (2023).

54. X. Han, K. T. Call, J. K. Pintor, G. Alarcon-Espinoza and A. B. Simon, Reports of insurance-based discrimination in health care and its association with access to care, *American journal of public health* **105**, S517 (2015).

55. C. A. McBurney and E. R. Vina, Racial and ethnic disparities in rheumatoid arthritis, *Current rheumatology reports* **14**, 463 (2012).

56. K. Evans, J. Coresh, L. D. Bash, T. Gary-Webb, A. Köttgen, K. Carson and L. E. Boulware, Race differences in access to health care and disparities in incident chronic kidney disease in the us, *Nephrology Dialysis Transplantation* **26**, 899 (2011).

57. K. A. Jenkins, S. Keddem, S. B. Bekele, K. E. Augustine and J. A. Long, Perspectives on racism in health care among black veterans with chronic kidney disease, *JAMA Network Open* **5**, e2211900 (2022).

58. M. Kim, Racial/ethnic disparities in depression and its theoretical perspectives, *Psychiatric Quarterly* **85**, 1 (2014).

59. Z. Rizvi, V. Kunder, H. Stewart, P. Torres, S. Moon, N. Lingappa, M. Kazaleh, V. Mallireddigari, J. Perez, N. John *et al.*, The bias of physicians and lack of education in patients of color with melanoma as causes of increased mortality: a scoping review, *Cureus* **14** (2022).

60. L. Krueger, E. Hijab, J.-A. Latkowski and N. Elbuluk, Clinical decision-making bias in darker skin types: a prospective survey study identifying diagnostic bias in decision to biopsy., *International Journal of Dermatology* **62** (2023).

61. A. Ahrens, C. B. Hansen, M. E. Schaffer and T. Wiemann, ddml: Double/debiased machine learning in stata, *The Stata Journal* **24**, 3 (2024).

62. V. Chernozhukov, C. Hansen, N. Kallus, M. Spindler and V. Syrgkanis, Applied causal inference powered by ml and ai (2024).

63. N. Dikkala, G. Lewis, L. Mackey and V. Syrgkanis, Minimax estimation of conditional moment models, *Advances in Neural Information Processing Systems* **33**, 12248 (2020).

64. T. Broderick, R. Giordano and R. Meager, An automatic finite-sample robustness metric: When can dropping a little data make a big difference? (2023).

65. T. Chen and C. Guestrin, XGBoost: A scalable tree boosting system (2016).

66. O. Ogungbe, A. K. Mitra and J. K. Roberts, A systematic review of implicit bias in health care: A call for intersectionality, *IMC Journal of Medical Science* **13**, 5 (2019).

67. J. Bosomworth and Z. Khan, Analysis of gender-based inequality in cardiovascular health: An umbrella review, *Cureus* **15** (2023).

68. S.-Y. Yang, E. Y. S. Woon, K. Griva and B. Y. Tan, A qualitative study of psychosocial factors in patients with knee osteoarthritis: insights learned from an asian population, *Clinical Orthopaedics and Related Research®* **481**, 874 (2023).

69. G. C. Gee, M. S. Spencer, J. Chen and D. Takeuchi, A nationwide study of discrimination and chronic health conditions among asian americans, *American journal of public health* **97**, 1275 (2007).

70. K. Kumar, R. J. Stack, A. Adebajo and J. Adams, Health-care professionals' perceptions of interacting with patients of south asian origin attending early inflammatory arthritis clinics, *Rheumatology Advances in Practice* **3**, p. rkz042 (2019).

71. M. A. Marzinke, D. N. Greene, P. M. Bossuyt, A. B. Chambliss, L. R. Cirrincione, C. R. Mc-Cudden, S. E. Melanson, J. H. Noguez, K. Patel, A. E. Radix *et al.*, Limited evidence for use of a black race modifier in egfr calculations: a systematic review, *Clinical chemistry* **68**, 521 (2022).

72. W. K. Newey and D. McFadden, Large sample estimation and hypothesis testing, *Handbook of econometrics* **4**, 2111 (1994).

73. P.-Å. Wedin, Perturbation bounds in connection with singular value decomposition, *BIT Numerical Mathematics* **12**, 99 (1972).

74. C. Davis and W. M. Kahan, The rotation of eigenvectors by a perturbation. iii, *SIAM Journal on Numerical Analysis* **7**, 1 (1970).

75. Y. Yu, T. Wang and R. J. Samworth, A useful variant of the davis–kahan theorem for statisticians, *Biometrika* **102**, 315 (2015).

76. H. Weyl, Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung), *Mathematische Annalen* **71**, 441 (1912).

77. G. Stewart, On the continuity of the generalized inverse, *SIAM Journal on Applied Mathematics* **17**, 33 (1969).

78. J. H. Wilkinson, *Rounding errors in algebraic processes* (SIAM, 2023).

79. J. H. Stock, Weak instruments, weak identification, and many instruments: Part 1 and part 2 (2018), NBER Summer Institute Methods Lectures.

80. G. Strang, *Introduction to linear algebra, Sixth Edition* (SIAM, 2022).

81. J.-P. Imhof, Computing the distribution of quadratic forms in normal variables, *Biometrika* **48**, 419 (1961).

82. A. Alaa and M. Van Der Schaar, Validating causal inference models via influence functions (2019).

83. D. A. Belsley, E. Kuh and R. E. Welsch, *Regression diagnostics: Identifying influential data and sources of collinearity* (John Wiley & Sons, 2005).

84. J. Fox, C. Kleiber and A. Zeileis, Diagnostics for 2sls regression `https://zeileis.github.io/ivreg/articles/Diagnostics-for-2SLS-Regression.html`, accessed: 12/18/2024.

85. V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey and J. Robins, Double/debiased machine learning for treatment and causal parameters (2016).

86. I. Montagni, T. Cariou, C. Tzourio and J.-L. González-Caballero, "i don't know","i'm not sure","i don't want to answer": a latent class analysis explaining the informative value of non-response options in an online survey on youth health, *International Journal of Social Research*

Methodology **22**, 651 (2019).

Appendix A. Semi-synthetic data generation method

We want to generate semi-synthetic data $(W_{\mathrm{gen}}, M_{\mathrm{gen}}, D_{\mathrm{gen}}, Z_{\mathrm{gen}}, X_{\mathrm{gen}}, Y_{\mathrm{gen}})$ that incorporates real data, generates all variables using partially linear structural equations, and allows for a user-specified implicit bias effect θ.

For confounders W_{gen}, we use a sample from the empirical distribution of the real confounders. In particular, W_{gen} is generated by sampling with replacement uniformly from the real confounders W. Subsequently, we want to learn from the real data a partially linear low-rank latent factor model where we are treating the mediator M as the latent factor. In particular, our goal is to learn a generative model that is of the form:

$$
\begin{aligned}
D_{\mathrm{gen}} &= \hat{f}_D(W_{\mathrm{gen}}, \hat{\epsilon}_D) \\
M_{\mathrm{gen}} &= a D_{\mathrm{gen}} + \hat{\epsilon}_M \\
Z_{\mathrm{gen}} &= \hat{G} M_{\mathrm{gen}} + \hat{f}_Z(W_{\mathrm{gen}}) + \hat{\epsilon}_Z \\
X_{\mathrm{gen}} &= \hat{F} M_{\mathrm{gen}} + \hat{f}_X(W_{\mathrm{gen}}) + \hat{\epsilon}_X
\end{aligned}
\tag{A.1}
$$

where $\hat{f}_Z, \hat{f}_X, \hat{G}, \hat{F}$ and the distributions of $\hat{\epsilon}_Z, \hat{\epsilon}_X, \hat{\epsilon}_M$, relate to real-world quantities.

Without loss of generality, the binary sensitive attribute D can be thought of as generated according to a Bernoulli distribution with success probability $p_D(W) = \mathbb{E}[D \mid W]$. Thus, to generate a sample D_{gen}, we learn a binary probabilistic classifier (or "propensity model") $\hat{p}_D(W)$ by running a logistic regression of D on W using the real data. To generate semi-synthetic D_{gen} conditioned on the already sampled W_{gen}, we sample based on a Bernoulli distribution with success probability $\hat{p}_D(W_{\mathrm{gen}})$. More formally,

$$
D_{\mathrm{gen}} = \hat{f}_D(W_{\mathrm{gen}}, \hat{\epsilon}_D) \sim \mathrm{Bernoulli}(\hat{p}_D(W_{\mathrm{gen}}))
\tag{A.2}
$$

The functions $\hat{f}_Z$ and $\hat{f}_X$ are LASSO regressions (with L1 penalty selected from cross-validation) trained to predict Z from W and X from W, respectively.

For the parameters $\hat{F}, \hat{G}$ as well as the distribution of M_{gen}, we proceed as follows. Consider the real data residuals $\tilde{Z}$ and $\tilde{X}$ after removing their predictable parts from W, based on the already trained LASSO regressions (i.e., $\tilde{Z} = Z - \hat{f}_Z(W)$). If the real world data were truly generated by the partially-linear structural equation model as in Equations (A.1), then the real data residuals $\tilde{Z}$ and $\tilde{X}$ would obey the generative equations

$$
\tilde{Z} = G\tilde{M} + \epsilon_Z \qquad\qquad\qquad \tilde{X} = F\tilde{M} + \epsilon_X
\tag{A.3}
$$

where $\tilde{M} = M - \mathbb{E}[M \mid W]$. Note that under this structural model:

$$
\mathrm{Cov}(\tilde{Z}, \tilde{X}) = \mathbb{E}[\tilde{Z}\tilde{X}^\top] = G\mathbb{E}[\tilde{M}\tilde{M}^\top]F^\top
\tag{A.4}
$$

Moreover, if we assume the latent factors of the mediator M are roughly independent, then $\mathrm{Var}(\tilde{M}) = \mathbb{E}[\tilde{M}\tilde{M}^\top] = \mathrm{diag}(\sigma_1, \ldots, \sigma_K)$, and a singular value decomposition (SVD) of the covariance $\mathrm{Cov}(\tilde{Z}, \tilde{X}) = U\Sigma V^\top$ would reveal G as the left singular vectors U, F as the right singular vectors V, and $(\sigma_1, \ldots, \sigma_K)$ as the singular values in Σ. Thus, to calculate the $\hat{G}$ and $\hat{F}$ that we will use in our model using Equations (A.1), we run SVD over the empirical covariance matrix $\mathbb{E}_n[\tilde{Z}\tilde{X}^\top] = U\Sigma V^\top$. Given $\Sigma = \mathrm{diag}(\hat{\sigma}_1, \ldots, \hat{\sigma}_n)$, we discard all singular values that fall

below a threshold calculated based on a measure of statistical significance (see appendix C.4) and consider only the remaining top K statistically significant left and right eigenvectors $\hat{G} = U_{.,1:K}$, $\hat{F} = V_{.,1:K}$ and their corresponding eigenvalues $(\hat{\sigma}_1, \ldots, \hat{\sigma}_K)$. The generated hidden mediator M_{gen} is taken to be K-dimensional and the distribution of the error variable $\hat{\epsilon}_M$ is a multivariate normal distribution with covariance matrix $\text{diag}(\hat{\sigma}_n, \ldots, \hat{\sigma}_K)$.

The noise variables $\hat{\epsilon}_Z$ and $\hat{\epsilon}_X$ are drawn from the empirical marginal distribution of the corresponding real world variables. Specifically, for a given ith random semi-synthetic sample $Z_{\text{gen}}^{(i)}$, we sample its noise $\hat{\epsilon}_Z^{(i)}$ by sampling uniformly at random a real data point $Z^{(j)}$, i.e., $\hat{\epsilon}_Z^{(i)} = Z^{(j)}$. Similarly, for the same ith random semi-synthetic sample, the noise $\hat{\epsilon}_X^{(i)}$ is generated by sampling uniformly at random another real data point $X^{(k)}$, i.e., $\hat{\epsilon}_X^{(i)} = X^{(k)}$. Finally, our algorithm generates the synthetic outcome based on the linear model:

$$Y_{\text{gen}} = \frac{b}{K} \sum_{i=1}^{K} M_{i,\text{gen}} + \theta D_{\text{gen}} + g X_{0,\text{gen}} + \hat{f}_Y(W_{\text{gen}}) + \sigma_Y \hat{\epsilon}_Y \tag{A.5}$$

where $X_{0,\text{gen}}$ denotes the first variable in X_{gen} and $M_{i,\text{gen}}$ the ith latent variable of M_{gen}. Here $\hat{f}_Y$ is a LASSO regression (with cross-validated L1 penalty) trained to predict Y from W. Note that we could generate Y_{gen} as any linear function of $M, D_{\text{gen}}, X_{\text{gen}}$; we just chose this equation for our experiments. The noise variable $\hat{\epsilon}_Y$ is drawn from the empirical distribution of Y in the real data. Specifically, for a given ith random semi-synthetic sample $\tilde{Y}_{\text{gen}}^{(i)}$, we generate a sample from $\hat{\epsilon}_Y$ by sampling uniformly at random a real data point $Y^{(m)}$, i.e., $\hat{\epsilon}_Y^{(i)} = Y^{(m)}$.

Parameters $a, b, \theta, g, \sigma_Y$ are chosen by the user as an input. a controls the degree of influence of the treatment D on the hidden mediator M. b controls the degree of direct influence of the hidden mediator M on the outcome Y. θ is the true controlled direct effect of the sensitive attribute D on Y. g controls the direct effect of the outcome proxies X on the outcome Y. σ_Y controls the level of noise of the outcome variable.

To avoid overfitting when training the parameters in our semi-synthetic generation process, we split our real data into a train-test (or train-sampling) split, and estimate $\hat{F}$, $\hat{G}$, $\text{diagonal}(\hat{\sigma}_1, \ldots, \hat{\sigma}_K)$, $\hat{p}_D$, $\hat{f}_X$, $\hat{f}_Z$ and $\hat{f}_Y$ using the training data. We use the test data to draw random samples from the empirical distributions that were used for $W_{\text{gen}}, \hat{\epsilon}_X, \hat{\epsilon}_Z, \hat{\epsilon}_Y$.

Appendix B. Proxy variable removal algorithm

For simplicity of notation, we omit the tilde's on the variables and use the short-hand notation in this section as: $Y \equiv \tilde{Y}$, $X \equiv \tilde{X}$, $Z \equiv \tilde{Z}$, $D \equiv \tilde{D}$, where recall $\tilde{V} = V - \mathbb{E}[V|W]$.

Furthermore, let $\mathcal{X} \subseteq [p_X]$ and $\mathcal{Z} \subseteq [p_Z]$ be subsets of indices of proxies[e], where p_X, p_Z is the dimension of X, Z correspondingly. Let the variables corresponding to the indices be $X_{\mathcal{X}} = (X_j)_{j \in \mathcal{X}}$ and $Z_{\mathcal{Z}} = (Z_j)_{j \in \mathcal{Z}}$.

Appendix B.1. *Overview*

---

[e]We use the notation $[n] = \{1, \ldots, n\}$ for any integer $n \geq 0$.

We develop an algorithm that selects a subset of proxy features (or variables; we use these terms interchangeably) from X, Z such that the primal and dual violation tests pass. Intuitively, this algorithm identifies pairs of indices $\mathcal{X}, \mathcal{Z}$ such that (1) there exists no direct paths from D to $X_{\mathcal{X}}$ that are not through M, and (2) there are no paths from $Z_{\mathcal{Z}}$ to Y that are not through M. In other words, there can exist no health proxy features in $X_{\mathcal{X}}$ that have an alternate mediator $M' \neq M$ that is not shared by $Z_{\mathcal{Z}}$, and vice versa. Thus, we only want to keep the proxy features $\mathcal{X}$ and $\mathcal{Z}$ that share the mediator M.

To estimate violations of $\mathcal{Z}$ and $\mathcal{X}$, we use the primal and dual equation violation tests, respectively (see Section 2.3 and Appendix C for more details on these tests, and why we believe they measure these violating paths). Let $p(\mathcal{X}, \mathcal{Z}), d(\mathcal{X}, \mathcal{Z})$ be the primal and dual test statistics when we use the subsets of proxies $X_{\mathcal{X}}, Z_{\mathcal{Z}}$. However, computing the primal and dual test statistics is expensive as it requires estimating parameters θ, h, γ. Thus we also use efficient approximations $\hat{p}(\mathcal{X}, \mathcal{Z}), \hat{d}(\mathcal{X}, \mathcal{Z})$. We discuss how we approximate $\hat{p}, \hat{d}$ in Appendix B.2.

We first compute the primal and dual approximate test scores $\hat{p}_* = \hat{p}([p_X], [p_Z]), \hat{d}_* = \hat{d}([p_X], [p_Z])$ using all X, Z proxy features. These are the baseline scores we want to beat by a substantial amount. Our proxy feature selection algorithm then starts by considering all proxy variables of Z, i.e. $\mathcal{Z} = [p_Z]$ and no proxies from X, i.e. $\mathcal{X} = \emptyset$. We increment $\mathcal{X}$ until we can no longer add more elements without incurring a large test score. At each iteration, we choose a random proxy variable $i \in [p_X]/\mathcal{X}$ and add it to $\mathcal{X}$ as long as $\hat{d}(\mathcal{X} \cup \{i\}, \mathcal{Z}) < \delta \cdot \hat{d}_*$ for some hyperparamter $\delta \in (0, 1)$; otherwise we discard the element i. Finally, when we have the largest candidate set $\mathcal{X}$ that still satisfies the aforementioned inequality, we check using the full dual test score if $d(\mathcal{X}, \mathcal{Z}) <$ critical-value, where critical-value is the rejection critical value calculated for the dual test score in Appendix C, and keep $\mathcal{X}$ if it passes. We repeat this K times over random seeds to get at most K potential candidate X proxy variable sets: $\{\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, ...\mathcal{X}^{(k)}\}$.

Next, we repeat the method to get candidate proxy variable sets for Z. For each candidate set for X, $\mathcal{X}^{(i)} \in \{\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, ...\mathcal{X}^{(K)}\}$, we start with no Z proxy variables: $\mathcal{Z} = \emptyset$. Then, we randomly add some feature $j \in [p_Z]$ into $\mathcal{Z}$ and continue adding Z variables as long as $\hat{p}(\mathcal{X}^{(i)}, \mathcal{Z} \cup \{j\}) < \delta \cdot \hat{p}_*$. Again, once we have the maximum size set $\mathcal{Z}$ that satisfies the inequality of the primal test estimate, we use the non-estimate primal score to check if $p(\mathcal{X}^{(i)}, \mathcal{Z}) <$ critical-value, where critical value is the rejection critical value calculated for the primal test in Appendix C, and only keep $\mathcal{Z}$ if it passes. For each of the $\mathcal{X}^{(i)}$ proxy variable sets $\{\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, ...\mathcal{X}^{(K)}\}$, we find K potential sets $\mathcal{Z}^{(j)}$ for a maximum of K^2 total candidate proxy variable set pairs. We repeat this one more iteration, where the only change is that instead of initializing Z to just $[p_Z]$ as we did in the first iteration when constructing the maximal X, we loop over all candidate $\mathcal{Z}^{(j)}$ sets.

The result is a list of potential candidate pairs $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$ for a specific (D, Y) estimate that minimally-violate the dual and the primal violation tests. Ideally, this list of candidates would be performed on a separate dataset from the one used to re-run the bias-detection method on each of the pairs.

Appendix B.2. *Dual and primal violation score estimates*

Both the primal and the dual tests (see Appendix C.1 and C.2, respectively) are checking whether a solution η to a linear system of equations:

$$\Sigma\eta = v$$

exists, where Σ is an $n \times m$ matrix and v is an m-dimensional vector. If a solution exists, then one such solution is the minimum norm solution $\Sigma^+ v$ and it would then satisfy:

$$\Sigma\Sigma^+ v = v \Leftrightarrow (\Sigma\Sigma^+ - I)v = 0$$

Thus, as an indication for the existence of a solution, we can use the ℓ_∞ norm of this vector[f]:

$$s = \|(\Sigma\Sigma^+ - I)v\|_\infty$$

A linear system that admits a solution will have a zero ℓ_∞ norm s. However, given access to n samples, we can only approximate Σ as Σ_n and v as v_n. Consider the error $\Sigma_n - \Sigma$, which is roughly $\mathbb{E}[\epsilon_X \epsilon_Z^\top]$ for both the dual and the primal. Because we do not expect the noise vectors ϵ_X and ϵ_Z to have a low rank structure, the error $\Sigma_n - \Sigma$ might be full rank, and thus Σ_n might not represent well the subspace spanned by the range of Σ. For this we introduce a low-rank approximation to remove the noise. For any matrix A, let $\mathrm{LRA}(A)$ denote its low-rank approximation where we only keep the singular values and corresponding eigenvectors of A are greater than a critical threshold τ (similarly to the process of calculating the threshold in the covariance rank test in Appendix C.4). In particular, if $A = UDV^\top$ is the SVD of A with $D = \mathrm{diag}(\lambda_1, \ldots, \lambda_K)$, then we let:

$$\mathrm{LRA}(A) = U\hat{D}V^\top \qquad \hat{D} = \mathrm{diag}(\lambda_1 \cdot 1\{|\lambda_1| > \tau\}, \ldots, \lambda_K \cdot 1\{|\lambda_K| > \tau\}) \qquad (\mathrm{B.1})$$

Moreover, to remove parts of v_n that span coordinates not spanned by v, we also introduce a sparsity regularization where $\mathrm{sparse}(v) = (v_1 \cdot 1\{|v_1| > \kappa_1\}, \ldots, v_K \cdot 1\{|v_K| > \kappa_K\})$ and the threshold κ_i is chosen to be a small multiple of the standard error of $\hat{v}_i$.

Let $\hat{\Sigma} = \mathrm{LRA}(\Sigma_n)$ and $\hat{v} = \mathrm{sparse}(v_n)$. Then we use as an approximate score quantity:

$$\hat{s} = \|(\hat{\Sigma}\hat{\Sigma}^+ - I)\,\hat{v}\|_\infty \qquad (\mathrm{B.2})$$

Moreover, for any pair of subsets of columns $\mathcal{S} \subseteq [n]$ and rows $\mathcal{T} \subseteq [m]$ (e.g., $\mathcal{S} = \mathcal{X}$), we use the approximate score function:

$$\hat{s}(\mathcal{S}, \mathcal{T}) = \|(\hat{\Sigma}_{\mathcal{S},\mathcal{T}}\hat{\Sigma}^+_{\mathcal{S},\mathcal{T}} - I)\,\hat{v}_{\mathcal{S}}\|_\infty \qquad (\mathrm{B.3})$$

where $\hat{\Sigma}_{\mathcal{S},\mathcal{T}}$ denotes the sub-matrix of $\hat{\Sigma}$ that contains only the rows in $\mathcal{S}$ and the columns in $\mathcal{T}$, and $\hat{v}_{\mathcal{S}}$ denotes the sub-vector of $\hat{v}$ that contains only the entries in $\mathcal{T}$.

For the case of the primal equation, we can create a proxy score $\hat{p}(\mathcal{X}, \mathcal{Z})$ by invoking the above approach with:

$$\Sigma_n = \mathbb{E}_n\left[\begin{pmatrix} Z \\ D \end{pmatrix}\begin{pmatrix} X \\ D \end{pmatrix}^\top\right] \qquad\qquad v_n = \mathbb{E}_n\left[\begin{pmatrix} Z \\ D \end{pmatrix} Y\right]$$

---

[f]Recall $\|x\|_\infty = \max_i \|x_i\|$.

and for the dual we can create a proxy score $\hat{d}(\mathcal{X}, \mathcal{Z})$ where:

$$\Sigma_n = \mathbb{E}_n[XZ^\top] \qquad\qquad v_n = \mathbb{E}_n[XD]$$

Appendix B.3. *Pseudocode*

Algorithm 1 Proxy Variable Selection Algorithm

1: **Input:** Proxy data Z and X, hyperparameter δ, number of iterations K
2: **Output:** Pairs of candidate variable sets for proxies in X and Z
3: Calculate baseline approximate primal and dual scores $\hat{p}_* = \hat{p}([p_X], [p_Z]), \hat{d}_* = \hat{d}([p_X], [p_Z])$
4: Initialize CandidateSets_Z = $\{[p_Z]\}$ ▷ List to store candidate sets of Z
5:
6: Set CandidateSets_X = $\emptyset$ ▷ List to store candidate sets of X
7: **for** each unique candidate $\mathcal{Z}^{(i)} \in$ CandidateSets_Z **do**
8:     **for** $k = 1$ to K **do**   ▷ Find candidate proxy sets of X
9:         $\mathcal{X} = \emptyset$ ▷ Start with empty set
10:         Randomly select new proxy variable $j \in [p_X]$
11:         **while** $\hat{d}(\mathcal{X} \cup \{j\}, \mathcal{Z}^{(i)}) < \delta \cdot \hat{d}_*$ **do**
12:             Set $\mathcal{X} \leftarrow \mathcal{X} \cup \{j\}$
13:             Randomly select new proxy variable $j \in [p_X]/\mathcal{X}$
14:         **end while**
15:         **if** $d(\mathcal{X}, \mathcal{Z}^{(i)}) <$ critical-value **then**   ▷ Compare true statistic to χ^2 critical value
16:             Add $\mathcal{X}$ to CandidateSets_X
17:         **end if**
18:     **end for**
19: **end for**
20: Set CandidateSetsPairs = $\emptyset$ ▷ List to store candidate variable pairs
21: Set CandidateSets_Z = $\emptyset$
22: **for** each candidate $\mathcal{X}^{(i)} \in$ CandidateSets_X **do**
23:     **for** $k = 1$ to K **do**   ▷ Find candidate proxy sets of Z
24:         $\mathcal{Z} = \emptyset$ ▷ Start with empty set
25:         Randomly select feature $j \in [p_Z]$
26:         **while** $\hat{p}(\mathcal{X}^{(i)}, \mathcal{Z} \cup \{j\}) < \delta \cdot \hat{p}_*$ **do**
27:             Set $\mathcal{Z} \leftarrow \mathcal{Z} \cup \{j\}$
28:             Randomly select new feature $j \in [p_Z]/\mathcal{Z}$
29:         **end while**
30:         **if** $p(\mathcal{X}^{(i)}, \mathcal{Z}) <$ critical value **then**   ▷ Compare true statistic to χ^2 critical value
31:             Add $(\mathcal{X}^{(i)}, \mathcal{Z})$ to CandidateSetPairs
32:             Add $\mathcal{Z}$ to CandidateSets_Z
33:         **end if**
34:     **end for**
35: **end for**
36:
37: Optionally repeat Lines 6 - 35.
38: **Output:** Candidate proxy variable pairs $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$

Appendix C. Tests

For simplicity of notation, we omit the tilde's on the variables and use the short-hand notation in this section as: $Y \equiv \tilde{Y}$, $X \equiv \tilde{X}$, $Z \equiv \tilde{Z}$, $D \equiv \tilde{D}$, where recall $\tilde{V} = V - \mathbb{E}[V|W]$.

Appendix C.1. *Primal violation test*

Appendix C.1.1. *Intuition*

Recall the (Primal Equation) is

$$\mathbb{E}\left[(Y - X^\top h - D\,\theta)\begin{pmatrix} Z \\ D \end{pmatrix}\right] = 0 \tag{C.1}$$

We want to test whether this equation admits a solution θ_0, h_*. This is a linear system of equations:

$$\mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix}\begin{pmatrix} X \\ D \end{pmatrix}^\top\right]\begin{pmatrix} h \\ \theta \end{pmatrix} = \mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix}Y\right] \tag{C.2}$$

Since all variables correspond to residuals and hence are mean-zero, the above is equivalent to:

$$\mathrm{Cov}((Z; D), (X; D))\begin{pmatrix} h \\ \theta \end{pmatrix} = \mathrm{Cov}((Z; D), Y) \tag{C.3}$$

Thus, the primal holds if $\mathrm{Cov}((Z; D), Y) \in$ column-span$(\mathrm{Cov}((Z; D), (X; D)))$. This is roughly equivalent to $\mathrm{Cov}(Z, Y) \in$ column-span$(\mathrm{Cov}(Z, X))$. Intuitively, this means that all the information flowing between Y and Z must go through the same mediator M that influences X and Z. This would not be true if either of two (equivalent) assumptions were violated: (1) there was some proxy variable Z_i which has an alternate path to Y not through X's mediator M, i.e., either $\exists Z_i$ that has a direct influence on Y or $\exists Z_i, M'$ s.t. $Z_i \to M' \to Y$ and $M' \not\to X$; (2) the proxy variables X have no direct relationship with the mediator M. The former potential violation is visualized in Figure F1.

If we view the primal violation as an issue of the first scenario (i.e., some Z_i has an alternate mediator M'), then we can mitigate the violation by identifying and removing such a Z_i, which our proxy variable removal algorithm attempts to do (see Appendix B). If we view the violation from the second scenario (i.e., Z's mediator M has no direct relationship to all proxy variables X), we could try adding stronger, more informative proxy variables to X that are likely to be mediated by M.

Appendix C.1.2. *Method*

If the primal equation admits a solution, then the minimum norm solution is one potential solution and is given by:

$$\begin{pmatrix} h_* \\ \theta_0 \end{pmatrix} = \Sigma_0^+ \, \mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix} Y \right] \qquad\qquad \Sigma_0 = \mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix} \begin{pmatrix} X \\ D \end{pmatrix}^\top \right] \qquad\qquad \text{(C.4)}$$

where Σ_0^+ is the pseudo-inverse of Σ_0.

Thus, to validate that the primal admits a solution, we need to prove the minimum norm solution θ_0, h_*, as defined by the above equations, also satisfy the primal moment restrictions:

$$\mathbb{E}\left[(Y - X^\top h_* - D\,\theta_0) \begin{pmatrix} Z \\ D \end{pmatrix} \right] = 0 \qquad\qquad \text{(C.5)}$$

Equivalently, if $m(V; \theta, h) = (Y - X^\top h - D\,\theta)(Z; D)$, with $V = (X, D, Z, Y)$, then we want to test the null hypothesis that $\mathbb{E}[m(V; \theta_0, h_*)] = 0$.

Let n be the number of samples and $p_Z + 1$ the dimension of the instruments $(Z; D)$, i.e., the dimension of the moment restrictions. For simplicity assume that n is even and that we split the data in half into training and testing samples. Let $\hat{h}, \hat{\theta}$ be estimates of the minimum norm solution using the training half of the data. Crucial to our proof is the assumption that that the minimum norm estimates are asymptotically linear estimators satisfying the following property:[g]

$$\sqrt{\frac{n}{2}} \Sigma_0 \left(\begin{pmatrix} \hat{h} \\ \hat{\theta} \end{pmatrix} - \begin{pmatrix} h_* \\ \theta_0 \end{pmatrix} \right) = \sqrt{\frac{n}{2}} \, J_0 \, \mathbb{E}_{\text{train}}\left[\begin{pmatrix} Z \\ D \end{pmatrix} (Y - X^\top h_* - D\theta_0) \right] + o_p(1) \qquad\qquad \text{(C.6)}$$

where J_0 is some matrix and $\mathbb{E}_{\text{train}}$ is the empirical average over the training samples.

If we estimate the minimum norm solution $\hat{h}, \hat{\theta}$ using the ℓ_2 regularized adversarial IV method as described in Appendix D.10.1, and furthermore assume that $\mathbb{E}[(Z; D)(Z; D)^\top]$ is invertible, the above holds with:

$$J_0 = \mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix} \begin{pmatrix} Z \\ D \end{pmatrix}^\top \right]^{1/2} \tilde{\Sigma}_0 \tilde{\Sigma}_0^+ \mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix} \begin{pmatrix} Z \\ D \end{pmatrix}^\top \right]^{-1/2} \qquad \tilde{\Sigma}_0 = \mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix} \begin{pmatrix} Z \\ D \end{pmatrix}^\top \right]^{-1/2} \Sigma_0$$

assuming we choose the ℓ_2 penalty on the parameters to decay faster than $1/\sqrt{n}$. We provide a proof for this asymptotic linearity property of the adversarial IV method in Appendix D.10.4.[h]

Consider the empirical average of the moment on the test set using the minimum norm estimates calculated on the training set, i.e.:

$$\hat{M} = \mathbb{E}_{\text{test}}[m(V; \hat{\theta}, \hat{h})]$$

---

[g]The notation $o_p(1)$ denotes terms that converge to zero in probability as the sample size n grows to infinity.

[h]We remark that if we apply a pre-processing step in this test and apply a PCA transformation to the instruments $(Z; D)$, replacing them with their PCA transformation then for the transformed instruments, we will have that $\mathbb{E}[(Z; D)(Z; D)^\top] = I$ and the formula simplifies to $J_0 = \Sigma_0 \Sigma_0^+$.

where $\mathbb{E}_{\text{test}}$ is the empirical average over the test samples. Let $M_0 = \mathbb{E}[m(V;\theta_0, h_*)]$, which is zero under the null. By influence function calculus arguments (see e.g. [72, Theorem 6.1]) and since the Jacobian of the moment m with respect to the vector $(h;\theta)$ is Σ_0, the estimate $\hat{M}$ admits the asymptotic linear representation:

$$\sqrt{\frac{n}{2}}(\hat{M} - M_0) = \sqrt{\frac{n}{2}}\,\mathbb{E}_{\text{test}}[m(V;\theta_0, h_*) - M_0] - \sqrt{\frac{n}{2}}\Sigma_0\left(\begin{pmatrix}\hat{h}\\\hat{\theta}\end{pmatrix} - \begin{pmatrix}h_*\\\theta_0\end{pmatrix}\right) + o_p(1)$$

$$= \sqrt{\frac{n}{2}}\,\mathbb{E}_{\text{test}}[m(V;\theta_0, h_*) - M_0] - \sqrt{\frac{n}{2}}J_0\,\mathbb{E}_{\text{train}}\left[\begin{pmatrix}Z\\D\end{pmatrix}(Y - X^\top h_* - D\theta_0)\right] + o_p(1)$$

where the second line derives from Equation (C.6). Thus we can derive the asymptotic linear representation of our estimate $\hat{M}$ as:

$$\sqrt{n}(\hat{M} - M_0) = \sqrt{n}\mathbb{E}_n\left[(m(V;\theta_0, h_*) - M_0)\frac{1\{\text{test}\}}{\Pr(\text{test})} - J_0\begin{pmatrix}Z\\D\end{pmatrix}(Y - X^\top h_* - D\theta_0)\frac{1\{\text{train}\}}{\Pr(\text{train})}\right] + o_p(1)$$

From this, we obtain that the estimate $\hat{M}$ is asymptotically normal:

$$\sqrt{n}(\hat{M} - M_0) \to_d N(0, A) \tag{C.7}$$

with asymptotic variance:

$$A = \mathbb{E}[(m(V;\theta_0, h_*) - M_0)^2 \mid 1\{\text{test}\}]\frac{1}{\Pr(\text{test})} + \mathbb{E}[\Phi\Phi^\top \mid 1\{\text{train}\}]\frac{1}{\Pr(\text{train})} \tag{C.8}$$

where:

$$\Phi = J_0\begin{pmatrix}Z\\D\end{pmatrix}(Y - X^\top h_* - D\theta_0) \tag{C.9}$$

$$= \mathbb{E}\left[\begin{pmatrix}Z\\D\end{pmatrix}\begin{pmatrix}Z\\D\end{pmatrix}^\top\right]^{1/2}\tilde{\Sigma}_0\tilde{\Sigma}_0^+\mathbb{E}\left[\begin{pmatrix}Z\\D\end{pmatrix}^\top\right]^{-1/2}\begin{pmatrix}Z\\D\end{pmatrix}(Y - X^\top h_* - D\theta_0) \tag{C.10}$$

We will construct an empirical estimate $\hat{A}$ of the variance A. Our first goal is to construct an estimate of the orthogonal projector matrix:

$$P \triangleq \tilde{\Sigma}_0\tilde{\Sigma}_0^+ = UDD^+U^\top \tag{C.11}$$

where the SVD of $\tilde{\Sigma}_0 = UDV^\top$, with $D = \text{diag}(\lambda_1, \ldots, \lambda_K)$ and $K = p_Z + 1$. To estimate $\hat{P}$, we consider the empirical estimate of $\tilde{\Sigma}_n$ of $\tilde{\Sigma}_0$ as

$$\tilde{\Sigma}_n = \mathbb{E}_{\text{train}}\left[\begin{pmatrix}Z\\D\end{pmatrix}\begin{pmatrix}Z\\D\end{pmatrix}^\top\right]^{-1/2}\mathbb{E}_{\text{train}}\left[\begin{pmatrix}Z\\D\end{pmatrix}\begin{pmatrix}X\\D\end{pmatrix}^\top\right]$$

and its SVD as $\tilde{\Sigma}_n = \hat{U}\hat{D}\hat{V}^\top$ with $\hat{D} = \text{diag}(\hat{\lambda}_1, \ldots, \hat{\lambda}_K)$. In Appendix C.1.3, we show we can estimate $\hat{P}$ by "soft-thresholding" the empirical singular values $\hat{\lambda}_i$ such that

$$\hat{P} = \hat{U}\hat{D}\hat{D}^+\hat{U}^\top = \hat{U}\,\text{diag}\left(\frac{\hat{\lambda}_1}{\hat{\lambda}_1 + n^{-0.2}}, \ldots, \frac{\hat{\lambda}_K}{\hat{\lambda}_K + n^{-0.2}}\right)\hat{U}^\top \tag{C.12}$$

We can then estimate the variance A using the estimates $\hat{P}, \hat{\Phi}$, and $\hat{J}$, i.e.,

$$\frac{\hat{A}}{n} = \frac{1}{n_{\text{test}}}\mathbb{E}_{\text{test}}[(m(V;\hat{\theta},\hat{h}) - \hat{M})^2] + \frac{1}{n_{\text{train}}}\mathbb{E}_{\text{train}}[\hat{\Phi}\hat{\Phi}^\top]$$

$$\hat{\Phi} = \hat{J}\begin{pmatrix} Z \\ D \end{pmatrix}(Y - X^\top\hat{h} - D\hat{\theta})$$

$$\hat{J} = \mathbb{E}_{\text{train}}\left[\begin{pmatrix} Z \\ D \end{pmatrix}\begin{pmatrix} Z \\ D \end{pmatrix}^\top\right]^{1/2} \hat{P}\ \mathbb{E}_{\text{train}}\left[\begin{pmatrix} Z \\ D \end{pmatrix}\begin{pmatrix} Z \\ D \end{pmatrix}^\top\right]^{-1/2}$$

Under regularity assumptions, the estimated variance $\hat{A}$ will be consistent and the following asymptotic normality assumption will also hold:

$$\sqrt{n}\,\hat{A}^{-1/2}\,(\hat{M} - M_0) \to_d N(0, I) \tag{C.13}$$

Since the random vector $\sqrt{n}\,\hat{A}^{-1/2}\,(\hat{M} - M_0)$ asymptotically follows the distribution of $p_Z + 1$ independent standard Gaussian random variables, then the squared ℓ_2-norm of this vector, i.e. $\left\|\sqrt{n}\,\hat{A}^{-1/2}\,(\hat{M} - M_0)\right\|_2^2$, asymptotically follows the distribution of the sum *of the squares* of $p_Z + 1$ independent standard Gaussian random variables, which corresponds to the χ^2 distribution with $p_Z + 1$ degrees of freedom. Thus we have:

$$n(\hat{M} - M_0)^\top\hat{A}^{-1}(\hat{M} - M_0) \to_d \chi^2(\text{dof} = p_Z + 1) \tag{C.14}$$

Under the null hypothesis the primal yields a solution, $M_0 = 0$ and thus we have that:

$$n\hat{M}^\top\hat{A}^{-1}\hat{M} \to_d \chi^2(\text{dof} = p_Z + 1) \tag{C.15}$$

Thus, we can use the test statistic:

$$\hat{T} = n\hat{M}^\top\hat{A}^{-1}\hat{M} \tag{Primal test statistic}$$

and compare it with the quantiles of a $\chi^2(\text{dof} = p_Z + 1)$ distribution to reject the null. In particular, if we want to erroneously reject the null with probability at most α, then we should be rejecting the null if:

$$\hat{T} > z_{1-\alpha} \tag{C.16}$$

where $z_{1-\alpha}$ is the $1 - \alpha$ percentile of a $\chi^2(\text{dof} = p_Z + 1)$ distribution.

Appendix C.1.3. *Soft-thresholding the orthogonal projection matrix P*

We expand here on estimating $P \triangleq \tilde{\Sigma}_0\tilde{\Sigma}_0^+$ using soft-thresholding of the singular values. We apply this to the primal and dual equations for their respective definitions of $\tilde{\Sigma}_0 = UDV^\top, \tilde{\Sigma}_n = \hat{U}\hat{D}\hat{V}^\top$. Explicitly, for the primal, recall that we define

$$\tilde{\Sigma}_0 = \mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix}\begin{pmatrix} Z \\ D \end{pmatrix}^\top\right]^{-1/2} \mathbb{E}\left[\begin{pmatrix} Z \\ D \end{pmatrix}\begin{pmatrix} X \\ D \end{pmatrix}^\top\right]$$

and for the dual we define

$$\tilde{\Sigma}_0 = \mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZX^\top]$$

Recall that given the SVD $\tilde{\Sigma}_0 = UDV^\top$, we have $P = UDD^+U^\top$. Note that $DD^+ = \text{diag}(1\{\lambda_1 \neq 0\}, \ldots, 1\{\lambda_K \neq 0\})$ where we expect only a small number $k < K$ of non-zero singular values λ_i. In essence, P is the orthogonal projector onto the range of $\tilde{\Sigma}_0$.

We estimate the orthogonal projector as $\hat{P}$ using the empirical estimate and corresponding SVD $\tilde{\Sigma}_n = \hat{U}\hat{D}\hat{V}^\top$ with $\hat{D} = \text{diag}(\hat{\lambda}_1, \ldots, \hat{\lambda}_K)$. Unlike the population singular values λ_i, the empirical singular values $\hat{\lambda}_i$ are typically all non-zero in practice due to sampling noise. Thus if we simply estimate P as $\hat{P} = \hat{U}\hat{D}\hat{D}^+\hat{U}^\top$, then $\hat{D}\hat{D}^+ = \vec{1}$ (the K-dimensional vector of 1's) and we will get the projection matrix $\hat{P} = \hat{U}\hat{U}^\top = I$.

To construct a projection matrix $\hat{P}$ more robust to sampling noise, note by the Wedin's theorem[73] (see also [74,75]), the space spanned by the top k (where recall $k = \text{rank}(\tilde{\Sigma}_0)$) eigenvectors of $\hat{U}$ is approximately the same as the space spanned by the top k eigenvectors of U, since $\tilde{\Sigma}_n$ is a consistent estimate of $\tilde{\Sigma}_0$. Thus if we can construct a matrix $\hat{\Lambda} = \text{diag}(v)$, where v converges to $(1, \ldots, 0)$ with a prefix of k ones, then the matrix $\hat{P} = \hat{U}\hat{\Lambda}\hat{U}^\top$ will be a consistent estimate of the matrix P. We will accomplish that by applying soft-thresholding to the empirical singular values:

$$\hat{P} = \hat{U}\,\text{diag}\left(\frac{\hat{\lambda}_1}{\hat{\lambda}_1 + n^{-0.2}}, \ldots, \frac{\hat{\lambda}_K}{\hat{\lambda}_K + n^{-0.2}}\right)\hat{U}^\top$$

To motivate this choice of soft-thresholding to estimate $\hat{P}$, first note that by Weyl's theorem,[76] the absolute difference between any of the singular values $\hat{\lambda}_i$ of $\tilde{\Sigma}_n$ and the corresponding singular values λ_i of $\tilde{\Sigma}$ are upper bounded by the operator norm of the difference $\tilde{\Sigma}_n - \tilde{\Sigma}_0$, i.e.,[i]

$$\forall j \in [K] : |\hat{\lambda}_j - \lambda_j| \leq \|\tilde{\Sigma}_n - \tilde{\Sigma}_0\|_{op}$$

Thus, we intuitively want to find an upper bound for $\|\tilde{\Sigma}_n - \tilde{\Sigma}_0\|_{op}$, which will provide insight into how to construct a new matrix $\hat{\Lambda}$ such that we can bound the error on the true singular values λ_i.

First, note that, for both the dual and the primal, $\tilde{\Sigma}_0$ takes the form $A^{-1/2}B$, with $A = \mathbb{E}[WW^\top]$ and $B = \mathbb{E}[VW^\top]$ for appropriately defined random vectors W and V and A assumed to be invertible (e.g., for the dual, $A = \mathbb{E}[ZZ^\top]$ and $B = \mathbb{E}[XZ^T]$). Accordingly, $\tilde{\Sigma}_n$ takes the form $A_n^{-1/2}B_n$, with $A_n = \mathbb{E}_n[WW^\top]$ and $B_n = \mathbb{E}_n[VW^\top]$. We can thus upper bound the operator norm:

$$\|\tilde{\Sigma}_n - \tilde{\Sigma}_0\|_{op} = \|A_n^{-1/2}B_n - A^{-1/2}B\|_{op} = \|(A_n^{-1/2} - A^{-1/2})B_n + A^{-1/2}(B_n - B)\|_{op} \qquad \text{(C.17)}$$

$$\leq \|A_n^{-1/2} - A^{-1/2}\|_{op}\|B_n\|_{op} + \|A^{-1/2}\|_{op}\|B_n - B\|_{op} \qquad \text{(C.18)}$$

Furthermore, if we use the fact that for any matrix M, $\|M\|_{op} \leq \|M\|_F$, and then apply standard concentration inequality arguments (Chernoff-Hoeffding bounds), we have that $\|A_n - A\|_{op} \leq \|A_n - A\|_F = O_p(n^{-1/2})$ and, similarly, $\|B_n - B\|_{op} \leq \|B_n - B\|_F = O_p(n^{-1/2})$. Then, by the triangle inequality we can derive that $\|B_n\|_{op} \leq \|B\|_{op} + O_p(n^{-1/2})$. Since $\|B\|_{op}$ is some

---

[i]Recall that, for a matrix A, the operator norm $\|A\|_{op}$ is intuitively the largest factor by which A "stretches" a vector. This is also the largest absolute singular value of the matrix A.

constant independent of n, then $\|B_n\|_{op} = O_p(n^{-1/2})$. Also, note that since A is invertible, we have $\|A^{-1/2}\|_{op} < \infty$ and therefore $\|A^{-1/2}\|_{op}$ is some constant independent of n. Thus $\|A^{-1/2}\|_{op} = O(1)$. Thus, applying these derivations to Equation (C.18), we get that:

$$\|\tilde{\Sigma}_n - \tilde{\Sigma}\|_{op} = O_p\left(\|A_n^{-1/2} - A^{-1/2}\|_{op}(1 + n^{-1/2}) + n^{-1/2}\right)$$

We next derive an upper bound on $\|A_n^{-1/2} - A^{-1/2}\|_{op}$. By continuity of the inverse (see e.g. [77] or [78]) we have that as long as $\|A^{-1/2}\|_{op}\|A_n^{1/2} - A^{1/2}\|_{op} < 1$ then we have that:

$$\|A_n^{-1/2} - A^{-1/2}\|_{op} \leq \frac{\|A^{-1/2}\|_{op}\|A_n^{1/2} - A^{1/2}\|_{op}}{1 - \|A^{-1/2}\|_{op}\|A_n^{1/2} - A^{1/2}\|_{op}} = O\left(\|A_n^{1/2} - A^{1/2}\|_{op}\right)$$

Moreover, we have that:[j]

$$\|A_n^{1/2} - A^{1/2}\|_{op} \leq \frac{1}{\sqrt{\lambda_{\min}(A)}}\|A_n - A\|_{op} = O(\|A_n - A\|_{op}) = O_p(n^{-1/2})$$

where $\lambda_{\min}(A)$ is the minimum eigenvalue of A. Note, then, that $\|A^{-1/2}\|_{op}\|A^{1/2} - A_n^{1/2}\|_{op} = O_p(n^{-1/2})$ will eventually be smaller than 1. Then we satisfy the aforementioned property $\|A^{-1/2}\|_{op}\|A_n^{1/2} - A^{1/2}\|_{op} < 1$ and can invoke the upper bound $\|A_n^{-1/2} - A^{-1/2}\|_{op} = O(\|A_n^{1/2} - A^{1/2}\|_{op}) = O_p(n^{-1/2})$. Finally we can conclude that:

$$\forall j \in [K] : |\lambda_j - \hat{\lambda}_j| \leq \|\tilde{\Sigma}_n - \tilde{\Sigma}\|_{op} = O_p(n^{-1/2})$$

For any $j \leq k$, the true singular value $\lambda_j \geq 0$ and thus as $n^{-0.2}$ converges to zero, $\hat{\lambda}_j$ will be bounded away from zero (since by the triangle inequality we have that $\hat{\lambda}_j \geq \lambda_j - O_p(n^{-0.5})$). Thus $\hat{\lambda}_j/(\hat{\lambda}_j + n^{-0.2}) \to 1$. For any $j > k$, we have that $\lambda_j = 0$ and therefore we expect that $|\hat{\lambda}_j| = |\hat{\lambda}_j - \lambda_j| = O_p(n^{-0.5})$. Thus eventually the $n^{-0.2}$ term will become the dominating term and we will have that $\hat{\lambda}_j/(\hat{\lambda}_j + n^{-0.2}) \to 0$. Note that any regularization $n^{-\alpha}$ for $\alpha < 0.5$ would satisfy the above properties. Experimentally we found that $\alpha = 0.2$ performs well empirically and gave nominal coverage across a wide range of dimensions of the hidden mediator.[k]

Appendix C.2. *Dual violation test*

---

[j]Consider any eigenvector x of $\sqrt{A_n} - \sqrt{A}$ with eigenvalue μ. Then:

$$x^\top(A_n - A)x = x^\top(\sqrt{A_n} - \sqrt{A})\sqrt{A_n}x + x^\top\sqrt{A}(\sqrt{A_n} - \sqrt{A})x = \mu x^\top(\sqrt{A_n} + \sqrt{A})x$$

Since the operator norm corresponds to the largest absolute eigenvalue, for symmetric matrices, and since A, A_n are symmetric, we can choose $\mu = \pm\|\sqrt{A_n} - \sqrt{A}\|_{op}$, which yields:

$$\|\sqrt{A_n} - \sqrt{A}\|_{op} \leq \frac{x^\top(A_n - A)x}{x^\top(\sqrt{A_n} + \sqrt{A})x} \leq \frac{\|A_n - A\|_{op}}{\lambda_{\min}(\sqrt{A_n} + \sqrt{A})} \leq \frac{\|A_n - A\|_{op}}{\lambda_{\min}(\sqrt{A})} = \frac{\|A_n - A\|_{op}}{\sqrt{\lambda_{\min}(A)}}$$

where we also used the fact that $\sqrt{A}$ and $\sqrt{A_n}$ are positive semi-definite.
[k]The reader can experiment with different values of α and different dimensions p_M of the hidden mediator using this notebook: `https://github.com/syrgkanislab/hidden_mediators/blob/main/MultiDimMediator.ipynb`.

Appendix C.2.1. *Intuition*

Recall the (Dual Equation) is

$$\mathbb{E}\left[(D - Z^\top\gamma)\,X\right] = 0 \tag{C.19}$$

We want to test whether this equation admits a solution γ_*. This is a linear system of equations:

$$\mathbb{E}[XZ^\top]\gamma = \mathbb{E}[XD] \tag{C.20}$$

Since all variables correspond to residuals and hence are mean-zero, the above is equivalent to:

$$\mathrm{Cov}(X, Z)\gamma = \mathrm{Cov}(X, D) \tag{C.21}$$

Thus, the dual admits a solution if $\mathrm{Cov}(X, D) \in$ column-span($\mathrm{Cov}(X, Z)$). Intuitively, this means that all the information flowing between D and X must go through the same mediator M that influences X and Z. This would not be true if either of the following two (equivalent) assumptions were violated: (1) there was some proxy variable X_i which has an alternate path from D not through Z's mediator M, i.e., either $\exists X_i$ such that D has a direct influence on X_i or $\exists X_i, M''$ s.t. $D \to M'' \to X_i$ and $M'' \not\to Z$; (2) the proxies Z are not correlated with the mediator M, i.e., Z has no direct edge from M nor an indirect correlation of M through D.

If we view the dual violation as an issue of the first scenario (i.e., some X_i has an alternate mediator M''), then we can mitigate the violation by identifying and removing such a X_i via our proxy selection algorithm. If we view the violation from the second scenario (i.e., X's mediator M has no direct relationship to all proxy variables Z), we could try adding stronger, more informative proxies to Z that are likely to be directly influenced by M.

Appendix C.2.2. *Method*

The proof is symmetric to that of the primal. If the dual equation admits a solution, then the minimum norm solution is one potential solution and is given by:

$$\gamma_* = \Sigma_0^+ \, \mathbb{E}\left[XD\right] \qquad\qquad \Sigma_0 = \mathbb{E}\left[XZ^\top\right] \tag{C.22}$$

where Σ_0^+ is the pseudo-inverse. Thus, to validate that the dual admits a solution, we need to prove the minimum norm solution γ_*, as defined by the above equations, also satisfy the dual moment restrictions:

$$\mathbb{E}\left[(D - Z^\top\gamma_*)X\right] = 0 \tag{C.23}$$

Equivalently, if $m(V; \gamma) = (D - Z^\top\gamma)X$, with $V = (X, D, Z)$, then we want to test the null hypothesis that $\mathbb{E}[m(V; \gamma_*)] = 0$.

Let n be the number of samples and p_X the dimension of X, i.e., the dimension of the dual moment restrictions. For simplicity assume that n is even and that we split the data in half into training and testing samples. Let $\hat{\gamma}$ be the estimate of the minimum norm solution using

the training half of the data. Crucial to our proof is the assumption that that the minimum norm estimate is an asymptotically linear estimator satisfying the following property:

$$\sqrt{\frac{n}{2}}\Sigma_0\left(\hat{\gamma}-\gamma_*\right)=\sqrt{\frac{n}{2}}\ J_0\ \mathbb{E}_{\text{train}}\left[X(D-Z^\top\gamma_*)\right]+o_p(1)$$

In a similar line of reasoning as with the primal, we have the above holds when solving for the minimum norm solution $\hat{\gamma}$ using adversarial IV (assuming $\mathbb{E}[XX^\top]$ is invertible) when

$$J_0=\ \mathbb{E}[XX^\top]^{1/2}\tilde{\Sigma}_0\tilde{\Sigma}_0^+\mathbb{E}[XX^\top]^{-1/2}\qquad\qquad\tilde{\Sigma}_0=\ \mathbb{E}[XX^\top]^{-1/2}\Sigma_0\qquad\text{(C.24)}$$

and the ℓ_2 penalty on the parameters decays faster than $1/\sqrt{n}$.

Consider the empirical average of the moment on the test set using the minimum norm estimates calculated on the training set, i.e.:

$$\hat{M}=\mathbb{E}_{\text{test}}[m(V;\hat{\gamma})]$$

Let $M_0=\mathbb{E}[m(V;\gamma_*)]$, which is zero under the null. By similar arguments as with the primal, we have the asymptotic linear representation

$$\sqrt{n}(\hat{M}-M_0)=\sqrt{n}\mathbb{E}_n\left[(m(V;\gamma_*)-M_0)\frac{1\{\text{test}\}}{\Pr(\text{test})}+J_0\ X(D-Z^\top\gamma_*)\frac{1\{\text{train}\}}{\Pr(\text{train})}\right]+o_p(1)$$

From this, we obtain that the estimate $\hat{M}$ is asymptotically normal:

$$\sqrt{n}(\hat{M}-M_0)\to_d N(0,A)\qquad\qquad\text{(C.25)}$$

with variance:

$$A=\mathbb{E}[(m(V;\gamma_*)-M_0)^2\mid 1\{\text{test}\}]\frac{1}{\Pr(\text{test})}+\mathbb{E}[\Phi\Phi^\top\mid 1\{\text{train}\}]\frac{1}{\Pr(\text{train})}\qquad\text{(C.26)}$$

where:

$$\Phi=\ J_0\ X(D-Z^\top\gamma_*)\qquad\qquad\qquad\text{(C.27)}$$
$$=\ \mathbb{E}[XX^\top]^{1/2}\tilde{\Sigma}_0\tilde{\Sigma}_0^+\mathbb{E}[XX^\top]^{-1/2}\ X(D-Z^\top\gamma_*)\qquad\text{(C.28)}$$

We will then construct an empirical estimate $\hat{A}$ of the variance A. If we denote the empirical variant of $\tilde{\Sigma}_0$ as

$$\tilde{\Sigma}_n=\mathbb{E}_{\text{train}}[XX^\top]^{-1/2}\mathbb{E}_{\text{train}}[XZ^\top]$$

where its SVD is $\tilde{\Sigma}_n=\hat{U}\hat{D}\hat{V}^\top$ with $\hat{D}=\text{diag}(\hat{\lambda}_1,\ldots,\hat{\lambda}_K)$ then we can estimate the orthogonal projection matrix $P\triangleq\tilde{\Sigma}_0\tilde{\Sigma}_0^+$ similarly to that of the primal by soft thresholding to the singular values (as described in Appendix C.1.3), i.e.,

$$\hat{P}=\ \hat{U}\,\text{diag}\left(\frac{\hat{\lambda}_1}{\hat{\lambda}_1+n^{-0.2}},\ldots,\frac{\hat{\lambda}_K}{\hat{\lambda}_K+n^{-0.2}}\right)\hat{U}^\top$$

We can then estimate the variance A using the estimates $\hat{P},\hat{\Phi}$, and $\hat{J}$, i.e.,

$$\frac{\hat{A}}{n}=\ \frac{1}{n_{\text{test}}}\mathbb{E}_{\text{test}}[(m(V;\hat{\gamma})-\hat{M})^2]+\frac{1}{n_{\text{train}}}\mathbb{E}_{\text{train}}[\hat{\Phi}\hat{\Phi}^\top]$$
$$\hat{\Phi}=\ \hat{J}\ X(D-Z^\top\hat{\gamma})$$
$$\hat{J}=\ \mathbb{E}_{\text{train}}[XX^\top]^{1/2}\tilde{\Sigma}_n\tilde{\Sigma}_n^+\ \mathbb{E}_{\text{train}}[XX^\top]^{-1/2}$$

Under regularity assumptions, the estimated variance $\hat{A}$ will be consistent and the following asymptotic normality assumption will also hold:

$$\sqrt{n}\,\hat{A}^{-1/2}\,(\hat{M} - M_0) \to_d N(0, I) \tag{C.29}$$

Similar to the case of primal, the squared ℓ_2 norm $\left\|\sqrt{n}\,\hat{A}^{-1/2}\,(\hat{M} - M_0)\right\|_2^2$ asymptotically follows the distribution of the sum *of the squares* of p_X independent standard Gaussian random variables, which corresponds to the χ^2 distribution with p_X degrees of freedom. Under the null hypothesis the dual yields a solution, $M_0 = 0$ and thus we have that:

$$n\hat{M}^\top \hat{A}^{-1}\hat{M} \to_d \chi^2(\text{dof} = p_X) \tag{C.30}$$

Thus, we can use the test statistic:

$$\hat{T} = n\hat{M}^\top \hat{A}^{-1}\hat{M} \tag{Dual test statistic}$$

and compare it with the quantiles of a $\chi^2(\text{dof} = p_X)$ distribution to reject the null. In particular, if we want to erroneously reject the null with probability at most α, then we should be rejecting the null if:

$$\hat{T} > z_{1-\alpha} \tag{C.31}$$

where $z_{1-\alpha}$ is the $1 - \alpha$ percentile of a $\chi^2(\text{dof} = p_X)$ distribution.

Appendix C.3. *Instrument identification strength*

Recall that our method estimates $V = (D - \gamma_*^\top Z)$ where γ_* satisfies the (Dual Equation), and V is now a new instrument in the (Primal Equation) that guarantees Neyman orthogonality to nuisance parameters h_*, γ_* when estimating θ. Let $\bar{Y} = Y - X^\top h_*$, and then the primal equation can be rewritten as the IV moment restriction equation

$$\mathbb{E}[V(\bar{Y} - D\,\theta)] = 0 \Leftrightarrow \theta_0 = \frac{\mathbb{E}[V\bar{Y}]}{\mathbb{E}[VD]} \tag{C.32}$$

where V is the instrument, $\bar{Y}$ is the outcome, and D the treatment. If the denominator $\mathbb{E}[VD]$ is small, this can cause instability to the estimate and invalidate the resulting confidence interval. This is exactly the type of condition that weak instrument tests check for. Hence, we propose two weak instrument tests with V as the instrument and D as the treatment.

Appendix C.3.1. *Intuition*

The denominator $\mathbb{E}[VD]$ is small if $D - \gamma_*^\top Z$ is uncorrelated with D. Note that γ_* is the solution to the IV moment $\mathbb{E}[X(D - \gamma^\top Z)] = 0$, where we are trying to predict D as a function of Z with instrument X. By definition, $\gamma_*^\top Z$ will capture the part of D that is predictable from Z through the mediator M. If M is very strongly correlated with D, then $\gamma_*^\top Z$ would be able to almost perfect predict D and thus the remnant variation $D - \gamma_*^\top Z$ is too small and uncorrelated with D. In essence, if the mediator is very correlated with D, then we will not be able to strongly identify the direct effect from the mediated effect of D on Y.

For instance, consider the case when D is a zero-meaned binary variable and X, Z, M are scalar random variables:

$$D \sim \text{Bernoulli}(\text{p}) - p \tag{C.33}$$
$$M = aD + \sigma_m N(0,1) \tag{C.34}$$
$$Z = eM + \sigma_z N(0,1) \tag{C.35}$$
$$X = fM + \sigma_x N(0,1) \tag{C.36}$$

Then

$$\mathbb{E}[M^2] = \sigma_m^2 + a^2 \mathbb{E}[D^2], \quad \mathbb{E}[XZ] = f \cdot e \cdot \mathbb{E}[M^2], \quad \mathbb{E}[DZ] = a \cdot e \cdot \mathbb{E}[D^2], \quad \mathbb{E}[DX] = a \cdot f \cdot \mathbb{E}[D^2]$$

Hence, the parameter γ_* and the denominator that determines the identification strength takes the form:

$$\gamma_* = \frac{\mathbb{E}[DX]}{\mathbb{E}[XZ]} = \frac{a}{e} \frac{\mathbb{E}[D^2]}{\mathbb{E}[M^2]} \qquad \mathbb{E}[VD] = \mathbb{E}[D^2] - \gamma_* \mathbb{E}[DZ] = \mathbb{E}[D^2] \frac{1}{1 + a^2 \mathbb{E}[D^2]/\sigma_m^2}$$

Thus we see that if the ratio a/σ_m (where recall a is the direct effect of $D \to M$) is very large, then we have weak identification. In other words, weak identification occurs when the mediator is primarily dictated by the treatment D and not other exogenous sources of variation. For instance, if $\sigma_m = 0$, then we cannot identify the direct effect. Moreover, we also have weak identification if $\mathbb{E}[D^2]$ is small, i.e. if the treatment is very deterministic.

Appendix C.3.2. *Method*

We consider two test statistics that target the weak identification problem. The first is based on a standard first stage F-test that is typically used in IV regression. However, here we also need to modify this statistic to account for the extra variance that stems from the estimation of the unknown parameter γ_* that enters the definition of the instrument V.

(a) V strength F-test - We develop an effective F-test[41,42] to check the correlation strength of V with D. We view our problem as a single instrument, single treatment setting with instrument V, treatment D, and outcome Y. Note, the first-stage IV equation from Equation (C.32) is $D = \pi_0 V + \epsilon$ where ϵ is the exogenous error term. The null hypothesis of a weak instrument test assumes π takes a very small value which induces a bias in the second-stage instrumental variable estimate $\hat{\theta}$ of approximate magnitude at least τ. A frequently used heuristic for τ is the Nagar bias, which posits that for a true first stage parameter π_0, an estimate $\hat{\pi}$ with variance σ_π induces a bias on the second-stage estimate of approximately $\tau = 1/\lambda^2$, where $\lambda^2 = \pi_0^2/\sigma_\pi^2$.[42] To check whether we have a weak instrument we use a heteroskedastic robust F-statistic (which is equal to the effective F-statistic when there is only one instrument) from [79] where, given a dataset with n samples, we estimate

$$F = \frac{\hat{\pi}^2}{\hat{\sigma}_\pi^2} \qquad \text{(F-test statistic)}$$

where $\hat{\pi} = \frac{\mathbb{E}_n[D\hat{V}]}{\mathbb{E}_n[D^2]}$ (with $\hat{V} = D - Z^\top \hat{\gamma}$) is the OLS estimate and $\hat{\sigma}_\pi^2$ is an estimate of the variance σ_π^2 of $\hat{\pi}$.

First let us assume that the estimate $\hat{\pi}$ satisfies asymptotic normality, i.e. $\sqrt{n}(\hat{\pi} - \pi_0) \to_d N(0, \tilde{\sigma}_\pi^2)$ where $\sigma_\pi^2 = \tilde{\sigma}_\pi^2/n$, and then we also have that

$$\frac{\hat{\pi}}{\hat{\sigma}_\pi} \approx N\left(\frac{\pi_0}{\sigma_\pi}, 1\right) \tag{C.37}$$

Then, under the null hypothesis of a weak instrument, i.e., that the Nagar bias is at most τ_* for some user-specified level τ_* (typically 0.1), we have the F-statistic $\hat{\pi}^2/\hat{\sigma}_\pi^2$ follows approximately a non-central $\chi^2_{1,c}$ distribution with 1 degree of freedom and centrality parameter $c = \lambda^2 \le 1/\tau_*$. With confidence α, we can reject the null if the test statistic F is larger than the $1-\alpha$ percentile of the $\chi^2_{1,1/\tau_*}$ distribution.

To compute the F-statistic, we need to estimate the variance $\hat{\sigma}_\pi^2$. The typical form of the heteroskedasticity robust asymptotic (not sample-size normalized) variance $\tilde{\sigma}_\pi^2$ is $\mathbb{E}[V^2(D - \pi V)^2]/\mathbb{E}[V^2]^2$. However, in our setting, the instrument $V = D - \hat{\gamma}^\top Z$ contains the estimate $\hat{\gamma}$ which introduces extra variance in the estimate $\hat{\pi}$ that we need to account for. Recall $\hat{\gamma}$ is the minimum norm solution to the regularized adversarial IV method (see Appendix D.10.1) with instrument X, treatment Z, and outcome D. For the purpose of this test, we will use a heuristic argument in the influence function of the estimate $\hat{\gamma}$:[1]

$$\sqrt{n}(\hat{\gamma} - \gamma_*) \approx \sqrt{n}\mathbb{E}_n[\Phi_\gamma] + o_p(1) \qquad \Phi_\gamma = \left(\mathbb{E}[QQ^\top] + \frac{\alpha}{n}I\right)^{-1}Q(D - Z^\top\gamma_*) \tag{C.38}$$

where $Q = \mathbb{E}[ZX^\top]\mathbb{E}[XX^\top]^+X$ is the projection of Z on X and α/n is the ℓ_2-penalty level used in the adversarial IV method. Using standard influence function calculus, we can derive the approximate asymptotic linear representation of $\hat{\pi}$ as:

$$\sqrt{n}(\hat{\pi} - \pi_0) \approx \sqrt{n}\mathbb{E}_n[\Phi_\pi] + o_p(1) \tag{C.39}$$

$$\Phi_\pi = \frac{V(D - \pi_0 V)}{\mathbb{E}[V^2]} + J_0^\top\Phi_\gamma \tag{C.40}$$

$$J_0 = \mathbb{E}[Z(2\pi_0 V - D)] \tag{C.41}$$

leading to variance of the estimate $\hat{\pi}$ that is approximately:

$$\sigma_\pi^2 = \mathbb{E}[\Phi_\pi^2]/n \tag{C.42}$$

Thus, given a dataset of n samples, we can empirically estimate this variance as

$$\hat{\sigma}_\pi^2 = \mathbb{E}_n[\hat{\Phi}_\pi^2]/n$$

where

$$\hat{\Phi}_\pi = \frac{\hat{V}(D - \hat{\pi}\hat{V})}{\mathbb{E}_n[\hat{V}^2]} + \hat{J}^\top\hat{\Phi}_\gamma \qquad\qquad \hat{J} = \mathbb{E}_n[Z(2\hat{\pi}\hat{V} - D)]$$

$$\hat{\Phi}_\gamma = (\mathbb{E}_n[\hat{Q}\hat{Q}^\top] + \frac{\alpha}{n}I)^{-1}\mathbb{E}[\hat{Q}(D - Z^\top\hat{\gamma})] \qquad\qquad \hat{Q} = \mathbb{E}_n[ZX^\top]\mathbb{E}_n[XX^\top]^+X$$

---

[1]A more formally justified correction term to the variance requires further theoretical investigation that is beyond the scope of the current work.

(b) z-test for $\mathbb{E}[DV] \geq \epsilon$ An alternative to the F-test is to simply test that the denominator $\mathbb{E}[DV]$ in Equation (C.32) is sufficiently far away from zero. Let $\pi_0 = \mathbb{E}[D(D - Z^\top \gamma_*)]$ and $\hat{\pi} = \mathbb{E}_n[D(D - Z^\top \hat{\gamma})]$ and using similar reasoning as in the test in the previous paragraph, we can derive an approximate asymptotic representation of the estimate as:

$$\sqrt{n}(\hat{\pi} - \pi_0) \approx \sqrt{n}\mathbb{E}_n[\Phi_\pi] + o_p(1) \qquad \Phi_\pi = D(D - Z^\top \gamma_*) - \pi_0 + J_0 \Phi_\gamma \qquad J_0 = -\mathbb{E}[DZ^\top]$$

Hence, we have that:

$$\sqrt{n}(\hat{\pi} - \pi_0) \to_d N(0, \tilde{\sigma}_\pi^2) \tag{C.43}$$

with $\tilde{\sigma}_\pi^2 = \mathbb{E}[\Phi_\pi^2]$.

Moreover, an estimate of the variance $\tilde{\sigma}_\pi^2$ can be calculated as:

$$\hat{\tilde{\sigma}}_\pi^2 = \mathbb{E}_n[\hat{\Phi}_\pi^2] \qquad \hat{\Phi}_\pi = D(D - Z^\top \hat{\gamma}) - \hat{\pi} + \hat{J}\hat{\Phi}_\gamma \qquad \hat{J} = -\mathbb{E}_n[DZ^\top] \tag{C.44}$$

Hence, we have that $\sqrt{n}\,\hat{\pi}$ is distributed approximately as $N(\pi_0, \hat{\tilde{\sigma}}_\pi^2)$. We can therefore use the test statistic

$$\hat{T} = \sqrt{n}\,|\hat{\pi}| = \sqrt{n}\,|\mathbb{E}_n[D(D - Z^\top \hat{\gamma})]| \qquad \text{(Z-test statistic)}$$

We can reject the null hypothesis of a weak instrument, i.e., that π_0 is smaller than some constant ϵ, if $\hat{T}$ is larger than the $1 - \alpha$ percentile of a folded normal distribution centered at ϵ and with scale $\hat{\tilde{\sigma}}_\pi$.

Appendix C.3.3. *Alternative Neyman Orthogonal Testing Method*

To avoid the heuristic argument used in the asymptotically linear representation of the regularized estimate $\hat{\gamma}$, we can instead consider a Neyman orthogonal variant of the two tests and their corresponding statistics. The empirical results presented in the paper use the heuristic argument of the previous section. Our code allows the user to specify which testing method variant they want to use, with the default being the method presented in this section.

(a) V strength F-test We can construct a Neyman orthogonal estimate of the parameter $\pi_0 = \frac{\mathbb{E}[DV]}{\mathbb{E}[D^2]}$, as follows. Consider the moment estimate that defines π_0:

$$\mathbb{E}[(D - Z^\top \gamma_*)(D - \pi_0(D - Z^\top \gamma_*))] = 0 \tag{C.45}$$

The gradient of the above equation with respect to γ_* is $\mathbb{E}[J_0 Z]$ where $J_0 = 2\pi_0 Z^\top \gamma_* - D$. Suppose that we identify a parameter ζ_* that satisfies the linear IV equation with instrument Z, treatment X and outcome J_0:

$$\mathbb{E}[(J_0 - X^\top \zeta)Z] = 0$$

Assuming that this moment equation admits a solution, the minimum norm solution ζ_* can be estimated via an ℓ_2-regularized adversarial IV method, in a manner similar to estimating γ_*. Then consider the debiased moment equation:

$$\mathbb{E}[(D - Z^\top \gamma_*)(D - \pi_0(D - Z^\top \gamma_*))] + \zeta_*^\top \mathbb{E}[X(D - Z^\top \gamma_*)] = 0 \tag{C.46}$$

This moment equation is Neyman orthogonal with respect to ζ_* since $\mathbb{E}[X(D - Z^\top \gamma_*)]$ is the gradient of the debiased moment with respect to ζ_* and by the definition of γ_*, we have that $\mathbb{E}[X(D - Z^\top \gamma_*)] = 0$. Moreover, the gradient of the debiased moment with respect to γ_* is:

$$\mathbb{E}[J_0 Z] - \mathbb{E}[ZX^\top]\zeta_* = \mathbb{E}[(J_0 - X^\top \zeta_*)Z]$$

which is also zero by the definition of ζ_*.

Thus we can construct a Neyman orthogonal moment estimate for the parameter π_0 as the solution to:

$$\mathbb{E}_n[(D - Z^\top \hat{\gamma})(D - \pi(D - Z^\top \hat{\gamma})) + \hat{\zeta}^\top X(D - Z^\top \hat{\gamma})] = 0 \qquad (C.47)$$

leading to the estimate:

$$\hat{\pi} = \frac{\mathbb{E}_n[(D - Z^\top \hat{\gamma})(D + X^\top \hat{\zeta})]}{\mathbb{E}_n[(D - Z^\top \hat{\gamma})^2]} \qquad (C.48)$$

Due to Neyman orthogonality and since the ℓ_2-regularized IV estimates converge at $o(n^{1/4})$ rates to the minimum norm solutions (as a special case of the results in [35] for the case of linear function spaces), the results of [40] imply that $\hat{\pi}$ is asymptotically linear:

$$\sqrt{n}(\hat{\pi} - \pi_0) = \sqrt{n}\mathbb{E}_n[\Phi_\pi] + o_p(1)$$

$$\Phi_\pi = \frac{(D - Z^\top \gamma_*)(D - \pi_0(D - Z^\top \gamma_*)) + {\zeta_*}^\top X(D - Z^\top \gamma_*)}{\mathbb{E}[(D - Z^\top \gamma_*)^2]}$$

Thus, we can derive the asymptotic normality statement:

$$\sqrt{n}(\hat{\pi} - \pi_0) \to_d N(0, \tilde{\sigma}_\pi^2) \qquad (C.49)$$

And we can derive an estimate of the variance $\sigma_\pi^2 = \tilde{\sigma}_\pi^2/n$ of the Neyman orthogonal estimate $\hat{\pi}$ as:

$$\hat{\sigma}_\pi^2 = \mathbb{E}_n[\hat{\Phi}_\pi^2]/n \qquad \hat{\Phi}_\pi = \frac{(D - Z^\top \hat{\gamma})(D - \hat{\pi}(D - Z^\top \hat{\gamma})) + \hat{\zeta}^\top X(D - Z^\top \hat{\gamma})}{\mathbb{E}_n[(D - Z^\top \hat{\gamma})^2]} \qquad (C.50)$$

The rest of the F-test is the same as described in Appendix C.3.2.

(b) z-test for $\mathbb{E}[DV] > \epsilon$ We can construct a Neyman orthogonal estimate of the parameter $\pi_0 = \mathbb{E}[DV]$ in a similar way to the F-test. Consider the moment estimate that defines π_0:

$$\mathbb{E}[(D - Z^\top \gamma_*)\, D] - \pi_0 = 0 \qquad (C.51)$$

The gradient with respect to γ_* is $\mathbb{E}[J_0 Z]$ where $J_0 = -D$.

Suppose that we identify a parameter ζ_* that satisfies the linear IV equation with instrument Z, treatment X and outcome J_0:

$$\mathbb{E}[(J_0 - X^\top \zeta)Z] = 0$$

Assuming that this moment equation admits a solution[m], the minimum norm solution ζ_* can be estimated via an ℓ_2-regularized adversarial IV method. Then we can consider the debiased moment equation:

$$\mathbb{E}[(D - Z^\top \gamma_*) D - \pi_0] + \zeta_*^\top \mathbb{E}[X(D - Z^\top \gamma_*)] = 0$$

This moment equation is Neyman orthogonal with respect to ζ_* since $\mathbb{E}[X(D - Z^\top \gamma_*)]$ is the gradient of the debiased moment with respect to ζ_* and by the definition of γ_*, we have that $\mathbb{E}[X(D - Z^\top \gamma_*)] = 0$. Moreover, the gradient of the debiased moment with respect to γ_* is:

$$\mathbb{E}[J_0 Z] - \mathbb{E}[ZX^\top]\zeta_* = \mathbb{E}[(J_0 - X^\top \zeta_*)Z]$$

which is also zero by the definition of ζ_*. Thus we can construct a Neyman orthogonal moment estimate for the parameter π_0 as the solution to:

$$\mathbb{E}_n[(D - Z^\top \hat{\gamma})D - \pi + \hat{\zeta}^\top X(D - Z^\top \hat{\gamma})] = 0$$

leading to the estimate:

$$\hat{\pi} = \mathbb{E}_n[(D - Z^\top \hat{\gamma})(D + X^\top \hat{\zeta})] \tag{C.52}$$

Due to Neyman orthogonality and since the ℓ_2-regularized IV estimates converge at $o(n^{1/4})$ rates to the minimum norm solutions, the results of [40] imply that $\hat{\pi}$ is asymptotically linear:

$$\sqrt{n}(\hat{\pi} - \pi_0) = \sqrt{n}\mathbb{E}_n[\Phi_\pi] + o_p(1)$$

$$\Phi_\pi = (D - Z^\top \gamma_*)D - \pi + \zeta_*^\top X(D - Z^\top \gamma_*) \tag{C.53}$$

Thus, we can derive the asymptotic normality statement:

$$\sqrt{n}(\hat{\pi} - \pi_0) \to_d N(0, \tilde{\sigma}_\pi^2) \tag{C.54}$$

where the estimate of the asymptotic variance $\tilde{\sigma}_\pi^2 = \mathbb{E}[\Phi_\pi^2]$ of the Neyman orthogonal $\hat{\pi}$ is

$$\hat{\tilde{\sigma}}_\pi^2 = \mathbb{E}_n[\hat{\Phi}_\pi^2] \qquad \hat{\Phi}_\pi = (D - Z^\top \hat{\gamma})D - \hat{\pi} + \hat{\zeta}^\top X(D - Z^\top \hat{\gamma}) \tag{C.55}$$

Thus we have that the new test statistic:

$$\hat{T} = \sqrt{n}\,|\hat{\pi}| = \sqrt{n}\,|\mathbb{E}_n[(D - Z^\top \hat{\gamma})(D + X^\top \hat{\zeta})]| \tag{C.56}$$

---

[m]Existence of a solution ζ_* is equivalent to the assumption that $\text{Cov}(Z, D)$ lies in the range of $\text{Cov}(Z, X)$. Note that our main assumption on the existence of a solution to the dual IV problem is equivalent to the assumption that $\text{Cov}(X, D)$ lies in the range of $\text{Cov}(X, Z)$. The two assumptions are slightly different but similar in nature. In Appendix D.6, we show existence of a solution γ_* to the dual IV problem is implied by the property that $\text{Cov}(M, Z)$ has full row rank, and we see a similar pattern follows for ζ_*. For instance, suppose that $\mathbb{E}[Z \mid M, D]$ is linear and independent of D, e.g., $\mathbb{E}[Z \mid M, D] = AM$ for some matrix A. (Allowing a dependence on D in this argument can also be handled, albeit the exposition becomes more complex and the sufficient assumptions harder to interpret.) Then $\text{Cov}(Z, D) = \mathbb{E}[ZD] = A\mathbb{E}[MD]$. Similarly, we have that $\text{Cov}(Z, X) = \mathbb{E}[ZX^\top] = \mathbb{E}[\mathbb{E}[Z \mid M]X^\top] = A\mathbb{E}[MX^\top]$. Thus existence of a solution is equivalent to existence of a ζ_* such that $\text{Cov}(Z, X)\zeta_* = \text{Cov}(Z, D)$, or equivalently, $A(\mathbb{E}[MD] - \mathbb{E}[MX^\top]\zeta_*) = 0$. Thus, existence of a ζ_* is implied by existence of a solution to the linear system $\mathbb{E}[MX^\top]\zeta_* = \mathbb{E}[MD]$, which is implied by the assumption that $\mathbb{E}[MX^\top] = \text{Cov}(M, X)$ has full row rank. Thus the assumptions needed for the existence of η_* are very similar to the prerequisites for the dual parameter γ_*.

approximately follows a folded normal distribution with center π_0 and scale $\hat{\hat{\sigma}}_\pi$. To test for a weak instrument, we proceed as before and reject the hypothesis that $\pi_0 \leq \epsilon$ by checking if the $z-$statistic is larger than the $1 - \alpha$ percentile of a folded normal distribution with center ϵ and scale $\hat{\hat{\sigma}}_\pi$.

Appendix C.4. *Proxy covariance rank test*

Appendix C.4.1. *Intuition*

We want to check if the health proxies X and Z are sufficiently related through the mediator M by estimating the rank of the covariance matrix $\mathrm{Cov}(X, Z)$. To do this, recall from Lemma 1 that the existence of a partially linear bridge function q relies on a partially conditional model $\mathbb{E}[X \mid M] = FM$, with F having full column rank (and recall all variables correspond to residuals after removing W).

Then the covariance of X and Z can be written as:

$$\mathrm{Cov}(X, Z) = \mathbb{E}[XZ^\top] = \mathbb{E}[\mathbb{E}[X \mid M, Z]Z^\top] = \mathbb{E}[\mathbb{E}[X \mid M]Z^\top] = F\mathbb{E}[MZ^\top] = F\,\mathrm{Cov}(M, Z)$$

If $\mathrm{Cov}(M, Z)$ has full row rank, we show in Appendix D.6 that this is an interpretable sufficient condition for the existence of a solution to the dual IV. Moreover, we also show that the full row rank assumption for $\mathrm{Cov}(M, Z)$ is necessary if we want to use the same proxies Z for any possible treatment D. Under this assumption, because $\mathrm{Cov}(X, Z)$ is the product of two matrices both of rank equal to the dimension of M, then the rank of $\mathrm{Cov}(X, Z)$ is also the dimension of M (see e.g. [80, Appendix A.1]). Thus under this sufficient (and essentially necessary) condition, the rank of $\mathrm{Cov}(X, Z)$ should reveal the dimension of M.

In practice, we do not know the true dimension of M. Therefore, by identifying the rank of $\mathrm{Cov}(X, Z)$, we are uncovering the dimension of M that such a pair of proxies can implicitly control for. If we believe that the hidden mediator has more dimensions than the identified rank of $\mathrm{Cov}(X, Z)$, then we should be concerned that this pair of proxies cannot control for all the mediators we are interested in.

Appendix C.4.2. *Method*

We calculate the rank of the covariance matrix by calculating the empirical covariance matrix $\mathbb{E}_n[XZ^\top]$ and estimating a threshold σ_* above which we determine the empirical matrix's singular values as statistically significant, and thus we report as the rank to be the number of such statistically significant singular values.

To compute the threshold singular value σ_*, we use a conservative test based on the error of the singular values between the sample (empirical) covariance matrix $\mathbb{E}_n[XZ^\top]$ and the true population covariance matrix $\mathbb{E}[XZ^\top]$. We leverage Weyl's theorem[76] and a high probability bound on the Frobenius norm of the error between these two matrices using a characterization of the asymptotic distribution of the norm as a weighted sum of χ^2 random variables.

In particular, Weyl's inequality for singular values, states that $\forall i \leq \min\{d_X, d_Z\}$

$$|\sigma_i(\mathbb{E}_n[XZ^\top]) - \sigma_i(\mathbb{E}[XZ^\top])| \leq \|\mathbb{E}_n[XZ^\top] - \mathbb{E}[XZ^\top]\|_{op} \leq \|\mathbb{E}_n[XZ^\top] - \mathbb{E}[XZ^\top]\|_F \qquad \text{(C.57)}$$

where $\|A\|_F = \sqrt{\sum_{i,j} A_{i,j}^2}$ is the Frobenius norm, $\|A\|_{op}$ is the operator norm, and $\sigma_i(A)$ denotes the i-th largest singular values of A. Thus we can control the error between the singular values of the empirical covariance matrix and the population covariance matrix, uniformly over all singular values, by constructing a high-probability bound on the Frobenius norm of the error $\mathbb{E}_n[XZ^\top] - \mathbb{E}[XZ^\top]$.

Let the Frobenius norm above be rewritten as

$$\|\mathbb{E}_n[XZ^\top] - \mathbb{E}[XZ^\top]\|_F = \sqrt{\sum_{i \leq d_X, j \leq d_Z} \vec{f}_{ij}^{\,2}}$$

where $\vec{f}$ is a $p = d_X \cdot d_Z$ dimensional vector and $\vec{f}_{ij} = \mathbb{E}_n[X_i Z_j] - \mathbb{E}[X_i Z_j]$. Note, by the multivariate CLT, the vector $\vec{f}$ approximately follows a multivariate normal distribution, i.e., $\vec{f} \to_d N(0, V)$, with covariance $V_{ij,kl} = \text{Cov}(X_i Z_j, X_k Z_l)/n$. We can approximate V with an empirical covariance $\hat{V}$ where $\hat{V}_{ij,kl} = \text{Cov}_n(X_i Z_j, X_k Z_l)/n$. Then, we have $\hat{V}^{-1/2} \vec{f} \to_d N(0, I)$. Therefore, the sum of the squares $\sum_{i,j} \vec{f}_{ij}^{\,2}$ asymptotically follows the weighted sum of p independent $\chi_2(1)$ distributions where the weights are the eigenvalues of V (see e.g. [81]), which can be approximated by the eigenvalues of $\hat{V}$.

We then calculate the $1 - \alpha$ percentile of this distribution via Monte Carlo simulation where we draw many samples of a p-dimensional vector with independent $\chi_2(1)$ entries, calculate the weighted sum for each sample using the eigenvalues of $\hat{V}$, and taking the $1 - \alpha$ percentile of this weighted sum across samples. Then the critical value is the square root of that percentile, i.e., $\sqrt{\sum_{ij} \vec{f}_{ij}^{\,2}}$. Asymptotically, the Frobenius norm of the error $\mathbb{E}_n[XZ^\top] - \mathbb{E}[XZ^\top]$ will not be larger than this critical value with probability more than α. Therefore, the error $|\sigma_i(\mathbb{E}_n[XZ^\top]) - \sigma_i(\mathbb{E}[XZ^\top])|$ for any of the singular values will not be larger than the critical value with probability more than α asymptotically as $n \to \infty$. Thus, this $1 - \alpha$ percentile we calculated will be our threshold σ_*. Any singular value which is below σ_* we declare as insignificant. The singular values above σ_* we can confidently claim are non-zero (albeit in a conservative manner, given that Weyl's inequality is only an upper bound). The number of such values is a lower bound on the rank of $\text{Cov}(X, Z)$ (and thus M) that we can be confident about.

Appendix C.5. *Weak instrument confidence interval*

Let $\hat{V} = D - \hat{\gamma}^\top Z$, $V = D - \gamma_*^\top Z$, $\hat{\bar{Y}} = Y - \hat{h}^\top X$, and $\bar{Y} = Y - h_*^\top X$, where γ_*, h_* are the true parameters and $\hat{\gamma}, \hat{h}$ are the estimates (i.e., minimum norm solutions). Consider the implicit bias effect estimate $\hat{\theta} = \mathbb{E}_n[\hat{V}\bar{Y}]/\mathbb{E}_n[\hat{V}D]$ of $\theta_0 = \mathbb{E}[V\bar{Y}]/\mathbb{E}[VD]$. Under a weak instrument regime, $\mathbb{E}[VD]$ and thus $\mathbb{E}_n[\hat{V}D]$ are very close to zero, leading to highly biased estimates and poor coverage of the confidence intervals that are based on asymptotic normality. In particular,

the intervals use the asymptotic approximation that:

$$\sqrt{n}(\hat{\theta} - \theta_0) = \sqrt{n}\left(\frac{\mathbb{E}_n[\hat{V}\hat{\bar{Y}}]}{\mathbb{E}_n[\hat{V}D]} - \theta_0\right) = \sqrt{n}\frac{\mathbb{E}_n[\hat{V}(\hat{\bar{Y}} - D\theta_0)]}{\mathbb{E}_n[\hat{V}D]} \approx \sqrt{n}\frac{\mathbb{E}_n[\hat{V}(\hat{\bar{Y}} - D\theta_0)]}{\mathbb{E}[\hat{V}D]} + o_p(1)$$

The last step is a good approximation only if the denominator $\mathbb{E}[VD]$ (and subsequently also $\mathbb{E}[\hat{V}D]$) is sufficiently separated from zero by a gap that is of magnitude larger than the $O(n^{-1/2})$ error we would expect that the empirical variant $\mathbb{E}_n[\hat{V}D]$ would have, as compared to its population analogue.

The test provided in Chapter 13.3 of [62] computes a confidence interval that does not invoke such an approximation argument for the denominator. Consider the null hypothesis that the true parameter θ_0 takes some value θ and since by construction the identifying moment equation is Neyman orthogonal with respect to all nuisance functions, we have from the results of [40] that:

$$\sqrt{\frac{n}{\mathrm{Var}((\bar{Y} - D\theta)V)}}\mathbb{E}_n[(\hat{\bar{Y}} - D\theta)\hat{V}] \to_d N(0,1) \tag{C.58}$$

and that the same also holds for the empirical estimate of the variance:

$$\sqrt{\frac{n}{\mathrm{Var}_n((\hat{\bar{Y}} - D\theta)\hat{V})}}\mathbb{E}_n[(\hat{\bar{Y}} - D\theta)\hat{V}] \to_d N(0,1) \tag{C.59}$$

Thus the square of the left-hand-side is distributed asymptotically as a $\chi^2(1)$ distribution:

$$C(\theta) \triangleq \frac{n\,\mathbb{E}_n[(\hat{\bar{Y}} - D\theta)\hat{V}]^2}{\mathrm{Var}_n((\hat{\bar{Y}} - D\theta)\hat{V})} \to_d \chi^2(1) \tag{C.60}$$

Thus, we have that, for the true parameter, $C(\theta_0) \to_d \chi^2(1)$ and therefore with probability $1-\alpha$, $C(\theta_0)$ should be less than the $(1-\alpha)$ percentile of the $\chi^2(1)$ distribution, denoted $q_{\chi^2(1)}(1-\alpha)$. Thus we can construct an $1 - \alpha$ confidence region (CR) for θ_0 as the set of all parameters θ for which $C(\theta) \le q_{\chi^2(1)}(1-\alpha)$:

$$\mathrm{CR} = \{\theta \in \mathbb{R} : C(\theta) \le q_{\chi^2(1)}(1-\alpha)\} \tag{C.61}$$

We operationalize this approach by searching over a grid candidate θ parameters in some pre-determined sufficiently large range of values and accepting a candidate parameter only if $C(\theta) \le q_{\chi^2(1)}(1-\alpha)$. In the end we return the confidence interval with endpoints that correspond to the smallest and largest accepted values. In our setting, the outcome is typically binary and hence it suffices to consider a grid of $[-1, 1]$.

Appendix C.6. *Influence score analysis*

Appendix C.6.1. *Calculating influence scores*

An influence score is a continuous value v_i assigned to each data point i that denotes how influential that point was in calculating the estimate (e.g., $\hat{\theta}$). Intuitively, the influence score v_i

approximates how much the estimate θ would change if we remove sample i from the training data, i.e. if we denote with $\hat{\theta}_{-i}$ the estimate we would have obtained if we removed the i-th sample from the data, then $v_i \approx \hat{\theta} - \hat{\theta}_{-i}$. Influence scores are often used to assess robustness of an effect estimate, where a robust estimate should not drastically change if a small number of highly-influential points are removed.[64,82]

Since our method is based on the empirical analogue of the Neyman orthogonal moment equation:

$$\mathbb{E}[(Y - X^\top h_* - D\theta)\,(D - Z^\top \gamma_*)] = 0 \qquad (C.62)$$

the estimate $\hat{\theta}$ admits an asymptotic linear representation:

$$\sqrt{n}(\hat{\theta} - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \frac{(Y_i - X_i^\top h_* - D_i\theta_0)\,(D_i - Z_i^\top \gamma_*)}{\mathbb{E}[D\,(D - Z^\top \gamma_*)]} + o_p(1) \qquad (C.63)$$

Thus the influence of each data point can be well approximated (as n goes to infinity) by:

$$\hat{\theta} - \hat{\theta}_{-i} \approx \frac{1}{n} \frac{(Y_i - X_i^\top h_* - D_i\theta_0)\,(D_i - Z_i^\top \gamma_*)}{\mathbb{E}[D\,(D - Z^\top \gamma_*)]} \qquad (C.64)$$

Thus we offer an approximation of the influence $\hat{\theta} - \hat{\theta}_{-i}$ of each data point by the score

$$v_i \triangleq \frac{1}{n} \frac{(Y_i - X_i^\top \hat{h} - D_i\hat{\theta})\,(D_i - Z_i^\top \hat{\gamma})}{\mathbb{E}_n[D\,(D - Z^\top \hat{\gamma})]} \qquad (C.65)$$

Intuitively, if $v_i \approx 0$, then point i has negligible effect on the implicit bias estimate. A large positive influence score v_i means patient i pushes the final estimate of θ to be more positive, and thus removing patient i from the data would result in a decrease of the final estimate by $\approx v_i$. Similarly, a very negative value of v_i pushes θ to be more negative, and removing patient i would increase the final estimate by $\approx |v_i|$.

In our analysis we also calculate a more accurate variant of the influence. In particular, fixing all nuisance estimates, our final estimate $\hat{\theta}$ is the result of a 2SLS algorithm with instrument $D - Z^\top \hat{\gamma}$, treatment D and outcome $Y - X^\top \hat{h}$. Exact leave-one-out influence scores have been derived for this algorithm by [83, pp. 266-68]. Thus, ignoring the influence that removing a sample has on the nuisance estimates, the leave-one-out influence can be computed exactly via the results of [83]. Our implementation uses these exact influence scores as a default. Moreover, analogues of several robust statistic diagnostics have been developed for 2SLS regression, such as Cook's distance, leverage scores, hat-values, studentized residuals and dffits (see [84]). Our implementation provides all these quantities for the final stage 2SLS regression.

Appendix C.6.2. *Estimating an influential set*

Similar to [64], we can use influence scores to create an influence set whose removal should maximally alter the estimate $\hat{\theta}$ in some pre-specified direction. For instance, we can return the minimal set of patients such that their removal from the dataset will lead to a new estimate $\hat{\theta}_{-\inf}$ and associated $(1 - \alpha)$-confidence interval $[\text{lb}_{-\inf}, \text{ub}_{-\inf}]^{\text{n}}$ that will contain zero, implying

---

[n]'lb'=lower bound; 'ub'=upper bound

that no statistically significant implicit bias can be detected for the remaining patients that are not in the minimal set.

We next describe how to calculate the minimal influence set.° Without loss of generality, suppose that $\hat{\theta} > 0$ and the confidence interval [lb, ub] using all the data (calculated either via the asymptotic normal approximation or via the weak-identification-robust confidence interval procedure in Appendix C.5) does not contain zero, i.e. lb > 0. First, we sort all n patients based on their influence scores such that $[v_0, v_1, \ldots, v_n]$ orders the patients from highest (most positive) to lowest (most negative) influence score. We only consider points whose scores $v_i \geq 0$, as removing them will push the estimate to be more negative. Suppose the first k sorted influence scores are non-negative. We then find the index m such that $\sum_{i=0}^{m} v_i >$ lb. Then our candidate minimal influence set are patients $S = \{1, 2, \ldots, m\}$, and we hypothesize that re-estimating $\hat{\theta}_{-\text{inf}}$ after removing all points in S will lead to a new confidence interval $[\text{lb}_{-\text{inf}}, \text{ub}_{-\text{inf}}]$ that will contain zero.

Appendix C.6.3. *Interpreting high-influence sets with feature interpretability*

To investigate what phenotypes characterize these identified high-influence patients, we perform a simple feature interpretability analysis. For each feature in W, Z, X, we compare the high-influence set with the rest of the patients not in the high-influence set. For categorical (including binary) features, we use a χ^2 test. For continuous features, we use a Kolmogorov–Smirnov test. We analyze all statistically significant differences between high-influence patients versus the rest of the data.

Appendix D. Proof of main theoretical claims

Appendix D.1. *Causal assumptions of Figure 1*

The following are the assumptions encoded in the structural causal model of Figure 1 and are necessary for identifying the controlled direct effect without observing the mediator M:

(1) Conditional ignorability: $Y(d, m) \perp\!\!\!\perp \{D, M(d)\} \mid W$

(2) Validity of outcome proxy X: $X \perp\!\!\!\perp D \mid (M, W)$

(3) Validity of treatment proxy Z: $Z \perp\!\!\!\perp (Y, X) \mid (D, M, W)$

(4) Consistency: $\mathbb{E}[Y(d, m) \mid M(d) = m, D = d] = \mathbb{E}[Y(d, m) \mid M = m, D = d]$

Appendix D.2. *Proof of the controlled direct effect as a g-formula*

°Given that the influence values are only approximations, after finding an influential set we remove the set and re-run estimation to confirm that the resulting confidence interval contains zero.

Recall the controlled direct effect

$$\theta = \int_{m,w} \mathbb{E}[Y(1,m) - Y(0,m) \mid W = w] \, p(m,w) \, dm \, dw$$

Here, we show that if M is known, we can identify θ using a g-formula. [P]

Lemma D.1. *The controlled direct effect of Figure 1 can be calculated by the g-formula*

$$\theta = \mathbb{E}[\mathbb{E}[Y \mid D = 1, M, W] - \mathbb{E}[Y \mid D = 0, M, W]]$$

Proof. We assume the causal constraints in D.1, and denote $M(d)$ as the potential value (also known as the counterfactual or intervened value) of the mediator M if we intervene and fix $D = d$. Given this,

$$\begin{aligned}
\mathbb{E}[Y(d,m) \mid W = w] &= \mathbb{E}[Y(d,m) \mid D = d, M(d) = m, W = w] \\
&= \mathbb{E}[Y(d,m) \mid D = d, M = m, W = w] \\
&= \mathbb{E}[Y(D,M) \mid D = d, M = m, W = w] \\
&= \mathbb{E}[Y \mid D = d, M = m, W = w]
\end{aligned}$$

Then the controlled direct effect can be identified as:

$$\begin{aligned}
\theta &= \int_{m,w} \mathbb{E}[Y(1,m) - Y(0,m) \mid W = w] \, p(m,w) \, dm \, dw \\
&= \int_{m,w} (\mathbb{E}[Y \mid D = 1, M = m, W = w] - \mathbb{E}[Y \mid D = 0, M = m, W = w]) \, p(m,w) \, dm \, dw \\
&= \mathbb{E}[\mathbb{E}[Y \mid D = 1, M, W] - \mathbb{E}[Y \mid D = 0, M, W]]
\end{aligned}$$

$\square$

Appendix D.3. *Proof of Theorem 1*

Appendix D.3.1. *Intuition*

Theorem 1 identifies how we could estimate the controlled direct θ when M isn't observed. To account for M, we leverage the information of the proxy variables X, Z via a "bridge function" q. Intuitively, if we can find a $q(D, X, W)$ that approximates the conditional expectation of Y (given D, M, W), then we show in our proof below that q also satisfies another equation:

$$\mathbb{E}[Y \mid D, M, W] = \mathbb{E}[q(D, X, W) \mid D, M, W] \tag{D.1}$$

$$\implies \mathbb{E}[Y - q(D, X, W) \mid D, Z, W] = 0 \tag{D.2}$$

Importantly, this latter equation is in the form of a Non-Parametric Instrumental Variable (NPIV) problem with confounders W, instruments $(Z; D)$, and treatment $(X; D)$. Following these properties of q, we will see in the below proof that in fact the controlled direct effect can be written as $\theta_0 = \mathbb{E}[q(1, X, W) - q(0, X, W)]$.

[P]In causal inference, a 'g-formula' estimates a causal effect by adjusting for confounding variables.

Appendix D.3.2. *Proofs*

Lemma D.2. *An outcome bridge function q which satisfies $\mathbb{E}[Y \mid D, M, W] = \mathbb{E}[q(D, X, W) \mid D, M, W]$ also satisfies $\mathbb{E}[Y - q(D, X, W) \mid D, Z, W] = 0$.*

Proof. First note $\mathbb{E}[Y \mid D, M, W] = \mathbb{E}[q(D, X, W) \mid D, M, W] \leftrightarrow \mathbb{E}[Y - q(D, X, W) \mid D, M, W] = 0$. We use the criteria of Figure 1 (see also D.1) that $Z \perp\!\!\!\perp (Y, X) \mid (D, M, W)$ for the following:

$$
\begin{aligned}
\mathbb{E}[Y - q(D, X, W) \mid D, Z, W] =&\ \mathbb{E}[\mathbb{E}[Y - q(D, X, W) \mid D, Z, M, W] \mid D, Z, W] \\
=&\ \mathbb{E}[\mathbb{E}[Y - q(D, X, W) \mid D, M, W] \mid D, Z, W] = 0
\end{aligned}
$$

Thus the outcome bridge function is also a solution to the *feasible* non-parametric Instrumental Variable (NPIV) regression problem in Equation (2). □

Lemma D.3. *If we identify an outcome bridge function q such that $\mathbb{E}[Y \mid D, M, W] = \mathbb{E}[q(D, X, W) \mid D, M, W]$, then $\theta = \mathbb{E}[q(1, X, W) - q(0, X, W)]$.*

Proof. We see that

$$
\begin{aligned}
\mathbb{E}[\mathbb{E}[Y \mid D = d, M, W]] =&\ \mathbb{E}[\mathbb{E}[q(D, X, W) \mid D = d, M, W]] \\
=&\ \mathbb{E}[\mathbb{E}[q(d, X, W) \mid D = d, M, W]] \\
=&\ \mathbb{E}[\mathbb{E}[q(d, X, W) \mid M, W]] \qquad\qquad (X \perp\!\!\!\perp D \mid M, W) \\
=&\ \mathbb{E}[q(d, X, W)]
\end{aligned}
$$

Also recall from Lemma D.1 that $\theta = \mathbb{E}[\mathbb{E}[Y \mid D = 1, M, W] - \mathbb{E}[Y \mid D = 0, M, W]]$. Then the controlled direct effect is identified as:

$$
\begin{aligned}
\theta =&\ \mathbb{E}[\mathbb{E}[Y \mid D = 1, M, W] - \mathbb{E}[Y \mid D = 0, M, W]] \\
=&\ \mathbb{E}[q(1, X, W) - q(0, X, W)]
\end{aligned}
$$

□

Appendix D.4. *Proof of Lemma 1: identification under partial linearity*

Appendix D.4.1. *Intuition*

If we assume $\mathbb{E}[Y \mid D, M, X, W]$ and $\mathbb{E}[X \mid M, W]$ are partially linear, and make a few other modest assumptions, then we can prove the bridge function q (which we saw in D.3 can be used to identify θ) is also partially linear. We also show in the proofs below that the controlled direct effect θ is simply a coefficient in the partially linear equation of $\mathbb{E}[Y \mid D, M, X, W]$.

Appendix D.4.2. *Proofs*

Lemma D.4. *(Restatement of Lemma 1) Under the assumption of partial linearity in Equation (3) and the assumptions encoded by the causal graph in Figure 1, we have $q(D, X, W) = D\,c + X^\top h + f(W)$ and $h = F^+ b + g$.*

Proof. By the assumptions of the causal graph we have that

$$Y = h_Y(M, D, X, W, u_Y)$$
$$X = h_M(M, W, u_X)$$

where h_Y, h_M are two arbitrary structural functions and u_Y, u_X are the independent exogenous error variables of the structural causal model that are independent of D, Z, M, W. Moreover, by the assumption of partially linear conditional expectation functions, we can write:

$$Y = D\,c + M^\top b + X^\top g + f_Y(W) + \epsilon_Y, \qquad \mathbb{E}[\epsilon_Y \mid D, M, X, W] = 0$$
$$X = F\,M + f_X(W) + \epsilon_X, \qquad \mathbb{E}[\epsilon_X \mid M, W] = 0$$

with $\epsilon_Y = h_Y(M, D, X, W, u_Y) - \mathbb{E}[h_Y(M, D, X, W, u_Y) \mid D, M, X, W]$ and $\epsilon_X = h_M(M, W, u_X) - \mathbb{E}[h_M(M, W, u_X) \mid M, W]$. Note that $\epsilon_X \perp\!\!\!\perp D \mid M, W$ since the only remnant randomness of ϵ_X is u_X when we condition on M, W, and $u_X \perp\!\!\!\perp D$.

We then rearrange the equations such that we no longer rely on observing M. Since F has full column rank, we have that the matrix $F^\top F$ is a $p_M \times p_M$ invertible matrix and the Moore-Penrose pseudo-inverse takes the form $F^+ = (F^\top F)^{-1} F^\top$. Thus we have

$$F^\top X = F^\top F\,M + F^\top f_X(W) + F^\top \epsilon_X \implies M = F^+(X - f_X(W) - \epsilon_X)$$

Note this expression of M only contains observable variables. We can then plug in this expression of M into the regression equation for Y to get:

$$Y = (b^\top F^+ + g^\top)X + D\,c - b^\top F^+ f_X(W) + f_Y(W) - b^\top F^+ \epsilon_X + \epsilon_Y$$

Let $h^\top = b^\top F^+ + g^\top$ and $f(W) = -b^\top F^+ f_X(W) + f_Y(W)$. Then we can write:

$$Y = h^\top X + D\,c + f(W) - b^\top F^+ \epsilon_X + \epsilon_Y$$

Moreover, the error terms satisfy that:

$$\mathbb{E}[\epsilon_Y \mid D, M, W] = \mathbb{E}[\mathbb{E}[\epsilon_Y \mid D, M, W, X] \mid D, M, W] = 0$$
$$\mathbb{E}[\epsilon_X \mid D, M, W] = \mathbb{E}[\epsilon_X \mid M, W] = 0$$

Thus, we can conclude that

$$\mathbb{E}[Y \mid D, M, W] = \mathbb{E}[h^\top X + D\,c + f_Y(W) \mid D, M, W]$$

Thus, the function $q(D, X, W) = D\,c + X^\top h + f(W)$ satisfies the property of an outcome bridge function as defined in Theorem 1 and therefore is a valid bridge function. Moreover, by Theorem 1, $q(D, X, W)$ also satisfies:

$$\mathbb{E}[Y \mid D, Z, W] = \mathbb{E}[h^\top X + D\,c + f(W) \mid D, Z, W]$$

$\square$

Lemma D.5. *Under the assumption of partial linearity in Equation* (3) *and the assumptions encoded by the causal graph in Figure 1, we have the controlled direct effect* $\theta = c$.

Proof. We know from Theorem 1 that $\theta = \mathbb{E}[q(1, X, W) - q(0, X, W)]$. Plugging in for q, we have

$$
\begin{aligned}
\theta &= \mathbb{E}[q(1, X, W) - q(0, X, W)] \\
&= (1 \cdot c + X^\top h + f(W)) - (0 \cdot c + X^\top h + f(W)) \\
&= c
\end{aligned}
$$

$\square$

Equivalence to other definitions of controlled direct effect Note that the partially linear restriction in Equation (3), together with the causal graph assumption, implies that: $\mathbb{E}[Y(d, m) \mid X, W] = d\,c + m^\top b + X^\top g + f_Y(W)$. The claim follows by the following set of equalities:

$$
\begin{aligned}
\mathbb{E}[Y \mid D = d, M = m, X, W] &= \mathbb{E}[Y(D, M) \mid D = d, M = m, X, W] \\
&= \mathbb{E}[Y(d, m) \mid D = d, M(d) = m, X, W] \\
&= \mathbb{E}[Y(d, m) \mid X, W] \qquad (Y(d, m) \perp\!\!\!\perp \{D, M(d)\} \mid X, W)
\end{aligned}
$$

Thus we conclude that:

$$
\mathbb{E}[Y(d, m) \mid X, W] = \mathbb{E}[Y \mid D = d, M = m, X, W] = d\,c + m^\top b + X^\top g + f_Y(W)
$$

Therefore we obtain that

$$
\mathbb{E}[Y(1, m) - Y(0, m) \mid X, W] = c,
$$

i.e. the direct effect is a constant, independent of X, W and independent of the value of the mediator m that we fix. Under this partially linear restriction, the controlled direct effect θ that we defined is equal to the constant c. Moreover, under this restriction, the constant c is equivalent to other definitions of a controlled direct effect, such as:

$$
\theta(m) \triangleq \mathbb{E}[Y(1, m) - Y(0, m)] = c \tag{D.3}
$$

for any fixed value m of the mediator.

Appendix D.5. *Proof of residualization of W*

For any variable V, define the residual $\tilde{V} = V - \mathbb{E}[V \mid W]$. Under the assumption of linear structural equations, we show that we can solve the NPIV moment restriction equation from Equation 2 after removing the effect of confounders W from all variables.

Proof. Note that the original conditional moment restrictions imply that (by marginalizing these constraints over X, D) yields

$$
\mathbb{E}[Y - D\,\theta - X^\top h \mid W] = \mathbb{E}[Y \mid W] - \mathbb{E}[D \mid W]\theta - \mathbb{E}[Y \mid W]^\top h = 0
$$

Subtracting these constraints from the original moment constraints, we see

$$
\mathbb{E}[\tilde{Y} - \tilde{D}\,\theta - \tilde{X}^\top h \mid D, Z, W] = 0
$$

Since $\tilde{Z}, \tilde{D}$ is measurable with respect to D, Z, W,

$$\mathbb{E}[\tilde{Y} - \tilde{D}\,\theta - \tilde{X}^\top h \mid \tilde{D}, \tilde{Z}] = 0$$

□

Appendix D.6. *Interpretable sufficient conditions for the existence of a dual solution*

We now provide a more primitive and intuitive condition for the existence of a solution γ_* to the dual IV problem. For simplicity of notation, we omit the tilde's on the variables and use the short-hand notation in this section as: $Y \equiv \tilde{Y}$, $X \equiv \tilde{X}$, $Z \equiv \tilde{Z}$, $D \equiv \tilde{D}$, where recall $\tilde{V} = V - \mathbb{E}[V|W]$.

The dual IV can be phrased as the question of the existence of a solution to the linear system:

$$\Sigma \gamma = v \qquad\qquad \Sigma \triangleq \mathbb{E}[XZ^\top] \qquad\qquad v \triangleq \mathbb{E}[XD] \qquad\qquad \text{(D.4)}$$

Under the partially linear conditional expectation function assumption, we have that:

$$v \triangleq \mathbb{E}[XD] = \mathbb{E}[\mathbb{E}[X \mid M, W, D]D] = \mathbb{E}[\mathbb{E}[X \mid M, W]D] = F\mathbb{E}[MD] = F\mathrm{Cov}(M, D)$$

Similarly:

$$\Sigma \triangleq \mathbb{E}[XZ^\top] = F\mathbb{E}[MZ^\top] = F\mathrm{Cov}(M, Z)$$

Thus existence of a solution γ to the original system is equivalent to:

$$F\left(\mathrm{Cov}(M, Z)\gamma - \mathrm{Cov}(M, D)\right) = 0 \qquad\qquad \text{(D.5)}$$

Since F has full column rank, the latter is equivalent to:

$$\mathrm{Cov}(M, Z)\gamma = \mathrm{Cov}(M, D) \qquad\qquad \text{(D.6)}$$

If $\mathrm{Cov}(M, Z)$ has full row rank, then we have an underdetermined system[q] and the above system always admits a solution, since $\mathrm{Cov}(M, Z)$ contains p_M (the dimension of M) linearly independent columns and hence its columns can span any p_M dimensional vector. Moreover, if we want to strengthen the assumption to say that we want the proxies Z to satisfy the assumption for any correlation pattern of the treatment D with the mediator, then we are asking the above system to have a solution for any vector $\mathrm{Cov}(M, D)$, which is equivalent to the assumption that $\mathrm{Cov}(M, Z)$ has full row rank. Thus if we want to use the same proxies Z to control for the mediator for many candidate treatments D (as we are in our empirical application), then such a full row-rank condition is basically both sufficient and necessary.

Corollary D.1. *If $Cov(M, Z)$ has full row rank, then there exists a solution γ_* to the dual IV, under the partially linear assumption in Lemma 1.*

---

[q]We are assuming $p_Z > p_M$.

Example 1 Consider the case where M is a scalar (e.g. binary or continuous). In this case, we have that $\text{Cov}(M, Z)$ is a row vector and $\text{Cov}(M, D)$ is a scalar. Thus, the above holds so long as $\exists Z_j$ such that $\text{Cov}(M, Z_j)$ is non-zero, i.e. some variable of the treatment proxy Z is correlated with the mediator. We see that re-interpreting solving the dual IV using the aforementioned linear system helps reveal what information the proxy Z has to carry about the hidden mediator.

Example 2 Consider the case where the mediator is a categorical variable with $K + 1$ categories and let M be the K-dimensional one-hot encoding of the mediator. Then note that row t of $\text{Cov}(M, Z)$ is a scalar multiple of the vector $\alpha_t \triangleq \mathbb{E}[Z \mid M = t]$. Thus, we need that the vectors α_t are linearly independent (i.e. none of them can be expressed as a linear combination of the rest). In other words, different values of the mediator cause linearly independent expected patterns in the treatment proxy vector.

Appendix D.7. *Proof of point identification from Theorem 2.2*

For simplicity of notation, we omit the tilde's on the variables Y, X, Z, D and use the short-hand notation in this section as: $Y \equiv \tilde{Y}$, $X \equiv \tilde{X}$, $Z \equiv \tilde{Z}$, $D \equiv \tilde{D}$.

Appendix D.7.1. *Intuition*

We prove that any h_*, γ_* that satisfy the (Primal Equation) and the (Dual Equation), respectively, will lead to the unique identification of the controlled direct effect θ_0 when solving for the moment equation $\mathbb{E}[(Y - X^\top h_* - D\theta)(D - \gamma_*^\top Z)] = 0$. This is important because it simplifies the parameter search space yet guarantees we still recover the controlled direct effect θ that we care about.

Appendix D.7.2. *Proofs*

Let h_0, θ_0 be the true parameters that satisfy the Primal Equation

$$\mathbb{E}\left[(Y - X^\top h - D\theta)\begin{pmatrix} Z \\ D \end{pmatrix}\right] = 0$$

Let h_1, θ_1 any other solution to Primal Equation. Both of these are also solutions to the modified NPIV moment equation

$$\mathbb{E}[(Y - X^\top h - D\theta)(D - \gamma^\top Z)] = 0 \tag{D.7}$$

for any γ, since this moment is a linear combination of the original moments. Let γ_* be any solution to the dual IV moment, i.e., $\mathbb{E}[X(D - \gamma_*^\top Z)] = 0$, with $\mathbb{E}[D(D - \gamma_*^\top Z)] \neq 0$. Then note that the modified NPIV equation is invariant to h, i.e.,

$$\mathbb{E}[(Y - X^\top h - D\theta)(D - \gamma_*^\top Z)] = \mathbb{E}[(Y - D\theta)(D - \gamma_*^\top Z)]$$

which implies that any solution to the modified NPIV equation takes the form (since $\mathbb{E}[D(D - \gamma_*^\top Z)] \neq 0$):

$$\mathbb{E}[(Y - D\theta)(D - \gamma_*^\top Z)] = 0 \implies \theta = \frac{\mathbb{E}[Y(D - \gamma_*^\top Z)]}{\mathbb{E}[D(D - \gamma_*^\top Z)]} \tag{D.8}$$

which is invariant to h. Thus, $\theta_1 = \theta_0$.

Appendix D.8. *Proof of Neyman orthogonality and Asymptotic Normality*

Appendix D.8.1. *Proofs*

Proof. We will show that the final moment equation:

$$\mathbb{E}[(\tilde{Y} - \tilde{X}^\top h - \tilde{D}\theta)(\tilde{D} - \gamma^\top \tilde{Z})] = 0 \tag{D.9}$$

is Neyman orthogonal (as defined in [40]) with respect to nuisance parameters h, γ and the nuisance functions that estimate the residual variables $\tilde{Y}, \tilde{X}, \tilde{D}, \tilde{Z}$.

- *Neyman orthogonality with respect to h:* Consider any solution to the dual equation γ_*. The orthogonality with respect to h is verified since the derivative with respect to h is zero:

$$\partial_h \mathbb{E}[(\tilde{Y} - \tilde{X}^\top h - \tilde{D}\theta)(\tilde{D} - \gamma_*^\top \tilde{Z})] = \mathbb{E}[\tilde{X}(\tilde{D} - \gamma_*^\top \tilde{Z})] = 0 \tag{D.10}$$

where we use the fact that γ_*, by definition, must satisfy $\mathbb{E}[\tilde{X}(\tilde{D} - \gamma^\top \tilde{Z})] = 0$.

- *Neyman orthogonality with respect to γ:* Consider any solution h_*, θ_0 to the primal equation and modified NPIV problem. We can prove Neyman orthogonality with respect to γ using the fact that h_*, θ_0 also solve the primal equation moment restriction:

$$\partial_\gamma \mathbb{E}[(\tilde{Y} - \tilde{X}^\top h_* - \tilde{D}\theta_0)(\tilde{D} - \gamma^\top Z)] = -\mathbb{E}[(\tilde{Y} - \tilde{X}^\top h_* - \tilde{D}\theta_0)\tilde{Z}^\top] = 0 \tag{D.11}$$

- *Neyman orthogonality with respect to the residual functions:* Finally, we show Neyman orthogonality with respect to the conditional expectation functions that appear in the residuals $\tilde{Y}, \tilde{X}, \tilde{D}, \tilde{Z}$. First note, that for the correct residuals,

$$\mathbb{E}[\tilde{V} \mid W] = \mathbb{E}[V - \mathbb{E}[V \mid W] \mid W] = 0 \qquad \forall V \in [Y, Z, X, D] \tag{D.12}$$

$$\mathbb{E}[\tilde{Y} - \tilde{X}^\top h - \tilde{D}\theta \mid W] = 0 \tag{D.13}$$

The directional derivative with respect to $q_Y(\cdot) = \mathbb{E}[Y \mid W = \cdot]$ in any direction ν_Y (which is any function that lies in the space of differences of q_Y) is:

$$\partial_{q_Y}[\nu_Y] \triangleq \partial_\tau \mathbb{E}[(Y - q_Y(W) - \tau\nu_Y(W) - \tilde{X}^\top h_* - \tilde{D}\theta_0)(\tilde{D} - \gamma_*^\top Z)] \mid_{\tau=0}$$

$$= -\mathbb{E}\left[\nu_Y(W)(\tilde{D} - \gamma_*^\top \tilde{Z})\right] = 0$$

Similarly, the directional derivative with respect to $q_D(\cdot) = \mathbb{E}[D \mid W = \cdot]$ in direction ν_D is:

$$\partial_{q_D}[\nu_D] \triangleq \partial_\tau \mathbb{E}[(\tilde{Y} - \tilde{X}^\top h_* - D\theta_0 + q_D(W) + \tau\theta_0\nu_D(W))(\tilde{D} - q_D(W) - \tau\nu_D(W) - \gamma_*^\top Z)] \mid_{\tau=0}$$

$$= \mathbb{E}\left[-\nu_D(W)\left(\tilde{Y} - \tilde{X}^\top h - \tilde{D}\theta_0\right) + \nu_D(W)\theta_0(\tilde{D} - \gamma_*^\top \tilde{Z})\right] = 0$$

Similarly, the directional derivatives with respect to $q_X = \mathbb{E}[X \mid W = \cdot]$ and $q_Z = \mathbb{E}[Z \mid W = \cdot]$ can be verified to be zero.

$\square$

Appendix D.8.2. *Asymptotic Normality*

We first define some notation. Let U be the vector of random variables (W, X, Z, D, Y) and one instance as $u = (w, x, z, d, y)$. Denote the nuisance parameter vector as $\eta = (h, \gamma, q_X, q_Z, q_D, q_Y)$, with $q_V(\cdot) = \mathbb{E}[V \mid W = \cdot]$ and $q_{0,V}(W) = E[V \mid W]$ for $V \in [Y, Z, X, D]$. Denote the function $\nu_V = \hat{q}_V - q_{0,V}$. For any function ν over an input space W, denote with $\|\nu\| = \sqrt{\mathbb{E}[\nu(W)^2]}$, the L^2 norm, and for any parameter vector $x \in \mathbb{R}^p$, we denote the squared ℓ_2−norm with $\|x\| = \sqrt{\sum_{i=1}^p x_i^2}$.

Our estimator is based on a moment restriction that is Neyman orthogonal and linear in the target parameter θ:

$$\psi(u; \theta, \eta) = \psi^a(u; \eta)\theta + \psi^b(u; \eta) \tag{D.14}$$

where

$$\psi^a(u; \eta) = \ -(D - q_D(W))\left(D - q_D(W) - \gamma^\top(Z - q_Z(W))\right) \tag{D.15}$$

$$\psi^b(u; \eta) = \ (Y - q_Y(W) - (X - q_X(W))^\top h)\left(D - q_D(W) - \gamma^\top(Z - q_Z(W))\right) \tag{D.16}$$

We want to prove such an estimator for θ is both asymptotically linear and asymptotically normal. To do so, we apply [85, Theorem 3.1] and show all the necessary regularity conditions for asymptotic behavior is satisfied.

Theorem D.1. *Suppose that all random variables W, X, Z, D, Y are almost surely and absolutely bounded by a constant, and that the nuisance estimates are consistent:*

$$\|\hat{h} - h_*\| = o_p(1) \qquad \|\hat{\gamma} - \gamma_*\| = o_p(1) \qquad \|\hat{q}_V - q_{0,V}\| = o_p(1) \tag{D.17}$$

Furthermore, suppose that the nuisance estimates satisfy the rate conditions:

$$\sqrt{n}\left(\|\nu_D\|^2 + \|\nu_D\|\,\|\nu_Z\| + \|\nu_Y\|\,\|\nu_D\| + \|\nu_Y\|\,\|\nu_Z\| + \|\nu_D\|\,\|\nu_X\| + \|\nu_Z\|\,\|\nu_X\|\right) = o_p(1) \tag{D.18}$$

and:

$$\sqrt{n}(\hat{\gamma} - \gamma_*)\mathbb{E}[\tilde{Z}\tilde{X}^\top](\hat{h} - h_*) = o_p(1) \tag{D.19}$$

Then the parameter estimate $\hat{\theta}$ satisfies:

$$\sqrt{n}(\hat{\theta} - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{(\tilde{Y}_i - \tilde{X}_i^\top h_* - \tilde{D}_i\theta_0)\,(\tilde{D}_i - \gamma_*^\top \tilde{Z}_i)}{\mathbb{E}[\tilde{D}\,(\tilde{D} - \gamma_*^\top \tilde{Z})]} + o_p(1) \to_d N\left(0, \sigma^2\right) \tag{D.20}$$

with

$$\sigma^2 = \frac{\mathbb{E}[(\tilde{Y} - \tilde{X}^\top h_* - \tilde{D}\theta_0)^2\,(\tilde{D} - \gamma_*^\top \tilde{Z})^2]}{\mathbb{E}[\tilde{D}\,(\tilde{D} - \gamma_*^\top \tilde{Z})]^2}. \tag{D.21}$$

Moreover, $(1-\alpha)$-confidence intervals can be constructed as:

$$\Pr\left(\theta_0 \in [\hat{\theta} \pm \Phi^{-1}(1-\alpha/2)\sqrt{\hat{\sigma}^2/n}]\right) \to 1-\alpha \tag{D.22}$$

where Φ is the CDF of the standard normal distribution and [r]:

$$\hat{\sigma}^2 = \frac{\mathbb{E}_n[(\hat{\tilde{Y}} - \hat{\tilde{X}}^\top\hat{h} - \hat{\tilde{D}}\hat{\theta})^2\,(\hat{\tilde{D}} - \hat{\gamma}^\top\hat{\tilde{Z}})^2]}{\mathbb{E}_n[\hat{\tilde{D}}\,(\hat{\tilde{D}} - \hat{\gamma}^\top\hat{\tilde{Z}})]^2} \tag{D.23}$$

Proof. We need to verify the regularity conditions of [85, Theorem 3.1]. We already have satisfied three conditions: (1) that the moment must satisfy Neyman orthogonality; (2) the nuisance functions must be consistent, i.e.:

$$\|\hat{\eta} - \eta_0\| = o_p(1) \tag{D.24}$$

and (3) the true Jacobian of the moment with respect to the parameter θ is invertible, which is satisfied as the true Jacobian $\mathbb{E}[\phi^a(U;\eta_0)] = \mathbb{E}[\tilde{D}(\tilde{D} - \gamma_*^\top\tilde{Z})]$ is assumed to be non-zero. Next, note that since all the variables are uniformly and absolutely bounded, the function ψ^α is Lipschitz in the nuisance functions, i.e.:

$$|\mathbb{E}[\psi^a(U;\eta)] - \mathbb{E}[\psi^a(U;\eta_0)]| \leq C\,\|\eta - \eta_0\| \tag{D.25}$$

for some large enough constant C, where $\|\eta - \eta_0\|^2 = \sum_{V\in[X,Z,D,Y]}\|q_V - q_{0,V}\|^2 + \|h - h_0\|^2 + \|\gamma - \gamma_0\|^2$. Moreover, the moments ψ and ψ^a both satisfy the mean-squared continuity condition:

$$\mathbb{E}[\|\psi(U;\theta_0,\eta) - \phi(U;\theta_0,\eta_0)\|^2] \leq C\|\eta - \eta_0\|^2 \tag{D.26}$$

$$\mathbb{E}\left[\left(\psi^a(U;\eta) - \psi^a(U;\eta_0)\right)^2\right] \leq C\,\|\eta - \eta_0\|^2 \tag{D.27}$$

Next, we need to satisfy that the Hessian of the moment with respect to the nuisance functions in the directions of their errors converges to zero faster than $n^{-1/2}$. That is,

$$\sqrt{n}\left(\partial_\tau^2\mathbb{E}[\psi(U;\theta_0,\eta_0 + \tau(\hat{\eta} - \eta_0))]\,|_{\tau=\bar{\tau}}\right) = o_p(1) \tag{D.28}$$

Let $\nu_V(W) = \hat{q}_V(W) - q_{0,V}(W)$. The Hessian can be decomposed as a sum of separate terms. First, we have the terms related to the second order derivatives with respect to the nuisance functions q_V, for $V \in [Z,X,D,Y]$:

$$A_1 = 2\mathbb{E}[\nu_D(W)^2] - 2\mathbb{E}[\nu_D(W)\,\gamma_*^\top\nu_Z(W)] = O\left(\|\nu_D\|^2 + \|\nu_D\|\,\|\nu_Z\|\right)$$

$$A_2 = 2\mathbb{E}[\nu_Y(W)\nu_D(W)] - 2\mathbb{E}[\nu_Y(W)\gamma_*^\top\nu_Z(W)] - 2\mathbb{E}[\nu_D(W)\,h_*^\top\nu_X(W)] + 2\mathbb{E}[\gamma_*^\top\nu_Z(W)\,h_*^\top\nu_X(W)]$$

$$= O\left(\|\nu_Y\|\,\|\nu_D\| + \|\nu_Y\|\,\|\nu_Z\| + \|\nu_D\|\,\|\nu_X\| + \|\nu_Z\|\,\|\nu_X\|\right)$$

Finally, we have the term related to the second order derivative with respect to the vector (h,γ):

$$A_3 = 2(\hat{\gamma} - \gamma_*)^\top\mathbb{E}[\tilde{X}\tilde{Z}^\top](\hat{h} - h_*)$$

Thus we need that:

$$\sqrt{n}(A_1 + A_2 + A_3) = o_p(1) \tag{D.29}$$

---

[r]We denote $\hat{\tilde{V}} = V - \hat{q}_V(W)$, where $\hat{q}_V$ is the estimate of the regression function $q_V(\cdot) = \mathbb{E}[V \mid W = \cdot]$.

This condition is satisfied under our assumptions. Finally, we note that our estimation algorithm does not invoke sample splitting when estimating the nuisance parameters h, γ. However, these parameters are low dimensional and hence the corresponding spaces satisfy the Donsker condition. Therefore, sample splitting can be avoided for these particular nuisance parameters. Given that the regularity conditions are all met, we can invoke the conclusion of [85, Theorem 3.1]] which states that the estimate $\hat{\theta}$ satisfies the asymptotic properties stated in Equations (D.20)-(D.23). □

Appendix D.9. *Estimation and confidence interval construction algorithm*

Algorithm 2 Estimation Algorithm

1: **Input:** samples of W, D, Z, X, Y
2: **Output:** controlled direct effect estimate $\hat{\theta}$ and standard error $\hat{\sigma}$
3: Residualize W from (D, Z, X, Y) using LASSO regularization and cross-fitting to select λ:

$$\tilde{D} = D - \beta_D^\top W \qquad \beta_D = \arg\min_\beta \mathbb{E}_n[(D - \beta_D^\top W)^2] + \lambda\|\beta_D\|_1 \qquad \text{(D.30)}$$

$$\forall i \in [p_z] : \tilde{Z}_i = Z_i - \beta_{Z_i}^\top W \qquad \beta_{Z_i} = \arg\min_\beta \mathbb{E}_n[(Z_i - \beta_{Z_i}^\top W)^2] + \lambda\|\beta_{Z_i}\|_1 \qquad \text{(D.31)}$$

$$\forall i \in [p_x] : \tilde{X}_i = X_i - \beta_{X_i}^\top W \qquad \beta_{X_i} = \arg\min_\beta \mathbb{E}_n[(X_i - \beta_{X_i}^\top W)^2] + \lambda\|\beta_{X_i}\|_1 \qquad \text{(D.32)}$$

$$\tilde{Y} = Y - \beta_Y^\top W \qquad \beta_Y = \arg\min_\beta \mathbb{E}_n[(Y - \beta_Y^\top W)^2] + \lambda\|\beta_Y\|_1 \qquad \text{(D.33)}$$

4: Calculate estimates $\hat{h}, \hat{\theta}_{pre}$ of the minimum norm solution (h_*, θ_0) to the primal moment equation

$$\mathbb{E}\left[(\tilde{Y} - \tilde{X}^\top h - \tilde{D}\theta)\begin{pmatrix}\tilde{Z} \\ \tilde{D}\end{pmatrix}\right] = 0 \qquad \text{(D.34)}$$

using an linear regularized adversarial IV solver (see Appendix D.10.1) with instruments $(\tilde{Z}; \tilde{D})$, treatments $(\tilde{X}; \tilde{D})$, outcome $\tilde{Y}$ and penalty parameter $\alpha \sim n^{0.3}$ (equiv. penalty level $\lambda \sim n^{-0.7}$).

5: Calculate estimate $\hat{\gamma}$ of the minimum norm solution γ_* to the dual moment equation

$$E[(\tilde{D} - \gamma^\top \tilde{Z})\tilde{X}] = 0$$

using an linear regularized adversarial IV solver (see Appendix D.10.1) with instruments $\tilde{X}$, treatments $\tilde{Z}$ and outcome $\tilde{D}$ and $\alpha \sim n^{0.3}$ (equiv. penalty level $\lambda \sim n^{-0.7}$).

6: Calculate final estimate $\hat{\theta}$ of the solution θ_0 to the Neyman orthogonal moment equation

$$\mathbb{E}[(\tilde{Y} - \tilde{X}^\top \hat{h} - \tilde{D}\theta)\,(\tilde{D} - \tilde{Z}^\top \hat{\gamma})] = 0,$$

and corresponding standard error of estimate as:

$$\hat{\theta} = \frac{\mathbb{E}_n[(\tilde{Y} - \hat{h}^\top \tilde{X})\,(\tilde{D} - \hat{\gamma}^\top \tilde{Z})]}{\mathbb{E}_n[\tilde{D}(\tilde{D} - \hat{\gamma}^\top \tilde{Z})]} \qquad \hat{\sigma} = \frac{1}{\sqrt{n}}\sqrt{\frac{\mathbb{E}_n[(\tilde{Y} - \hat{h}^\top \tilde{X} - \tilde{D}\hat{\theta})^2(\tilde{D} - \hat{\gamma}^\top \tilde{Z})^2]}{\mathbb{E}_n[\tilde{D}(\tilde{D} - \hat{\gamma}^\top \tilde{Z})]^2}} \qquad \text{(D.35)}$$

Appendix D.10. *Linear Regularized Adversarial IV*

Appendix D.10.1. *Procedure*

We define in detail a generic procedure for estimating the minimum norm solution θ_0 to a linear instrumental variable (IV) regression problem of the form:

$$\mathbb{E}[(Y - \theta^\top X)Z] = 0$$

for some outcome Y, instrument vector Z, and treatment vector X, where the covariance matrix $\mathbb{E}[ZX^\top]$ is ill-posed, i.e., has a zero singular value. (Note that the optimal solution $\hat{\theta}$ to this IV problem is a general parameter and thus separate from the implicit bias effect estimate.) Our procedure is based on the ℓ_2-regularized adversarial IV method that optimizes the empirical criterion defined in [63] that solves:

$$\min_\theta \max_\beta \mathbb{E}_n[2(Y - \theta^\top X)Z^\top \beta - \beta^\top ZZ^\top \beta] + \lambda \|\theta\|^2 \tag{D.36}$$

The optimal solution to the inner optimization problem takes the form:

$$\beta_* = \mathbb{E}_n[ZZ^\top]^+ \mathbb{E}_n[Z(Y - \theta^\top X)] \tag{D.37}$$

and thus, using the optimal solution, we get:

$$\max_\beta \mathbb{E}_n[2(Y - \theta^\top X)Z^\top \beta - \beta^\top ZZ^\top \beta] = \mathbb{E}_n[(Y - \theta^\top X)Z^\top]\mathbb{E}_n[ZZ^\top]^+ \mathbb{E}_n[Z(Y - \theta^\top X)Z] \tag{D.38}$$

Thus $\hat{\theta}$ is defined as the minimizer of:

$$\min_\theta \mathbb{E}_n[(Y - \theta^\top X)Z^\top]\mathbb{E}_n[ZZ^\top]^+ \mathbb{E}_n[Z(Y - \theta^\top X)] + \lambda \|\theta\|^2 \tag{D.39}$$

We can define the random variable:

$$\hat{Q} = \mathbb{E}_n[XZ^\top]\mathbb{E}_n[ZZ^\top]^+ Z \tag{D.40}$$

Note that $\hat{B} = \mathbb{E}_n[ZZ^\top]^+ \mathbb{E}_n[ZX^\top]$ is the minimum norm solution to the OLS problem of predicting X linearly as a function of Z. Hence, $\hat{Q} = \hat{B}^\top Z$ is a best linear prediction of X as a function of Z. Notably, $\hat{B}$ satisfies the first order conditions to the OLS problem, i.e.

$$\mathbb{E}_n[Z(X - \hat{B}^\top Z)^\top] = 0 \implies \hat{B}^\top \mathbb{E}_n[Z(X - \hat{B}^\top Z)^\top] = 0 \implies \mathbb{E}_n[\hat{Q}X^\top] = \mathbb{E}_n[\hat{Q}\hat{Q}^\top] \tag{D.41}$$

Then the solution $\hat{\theta}$ to the regularized adversarial IV problem is the solution to the equation:

$$\mathbb{E}_n[\hat{Q}(Y - X^\top \theta)] + \lambda \theta = 0 \Leftrightarrow \mathbb{E}_n[\hat{Q}Y] - (\mathbb{E}_n[\hat{Q}X^\top] + \lambda I)\theta = 0 \tag{D.42}$$

which takes the closed form:

$$\hat{\theta} = (\mathbb{E}_n[\hat{Q}X^\top] + \lambda I)^+ \mathbb{E}_n[\hat{Q}Y] = (\mathbb{E}_n[\hat{Q}\hat{Q}^\top] + \lambda I)^+ \mathbb{E}_n[\hat{Q}Y] \tag{D.43}$$

Thus in this setting, the regularized adversarial IV solution is equivalent to a regularized two-stage-least-squares procedure where we use the empirical minimum norm solution for the first stage and we use a ℓ_2-regularized second stage with penalty level λ.

Algorithm 3 Regularized Adversarial IV Procedure: $\text{AdvIV}(Z, X, Y; \alpha)$

1: **Input:** n samples of the variables (Z, X, Y), with instrument vector Z, treatment vector X, outcome Y; ℓ_2 penalty hyperparameter α
2: **Output:** Estimate of the minimum norm solution $\hat{\theta}$ to the linear IV problem defined via the moment restriction equations: $\mathbb{E}[(Y - \theta^\top X)Z] = 0$
3: Calculate the OLS solution of the first stage, giving the minimum mean-squared-projection $\hat{Q}$ of X on Z:

$$\hat{Q} = \hat{B}^\top Z \qquad\qquad \hat{B} = \mathbb{E}_n[ZZ^\top]^+ \mathbb{E}_n[ZX^\top] \qquad\qquad \text{(D.44)}$$

4: Calculate the regularized second stage solution:

$$\hat{\theta} = \left(\mathbb{E}_n[\hat{Q}X^\top] + \frac{\alpha}{n}I \right)^+ \mathbb{E}_n[\hat{Q}Y] = \left(\mathbb{E}_n[\hat{Q}\hat{Q}^\top]^{-1} + \frac{\alpha}{n}I \right)^+ \mathbb{E}_n[\hat{Q}Y] \qquad \text{(D.45)}$$

Following very similar arguments as the ones presented in [36, Theorem 4] (see Appendix D.10.2), we can argue that $\hat{\theta}$ converges to the minimum norm solution θ_0, as long as $\lambda = \omega(1/n)$ and the covariance matrix $\mathbb{E}[ZZ^\top]$ is full rank, as well as several other regularity assumptions. Moreover, the convergence of $\|\hat{\theta} - \theta_0\|$ follows a rate of $\sqrt{n^{-1}/\lambda + \lambda}$, while the convergence in terms of the weak norm

$$\|\hat{\theta} - \theta_0\|_w \triangleq \|\mathbb{E}[ZX^\top](\hat{\theta} - \theta_0)\|$$

should follow a rate of $\sqrt{n^{-1} + \lambda^2}$. In particular, in Appendix D.10.2 we prove the following theorem:

Theorem D.2 (Rates). *Assume that $\mathbb{E}[ZZ^\top]$ is full rank. Let $\Sigma_0 = \mathbb{E}[ZX^\top]$ and $\tilde{\Sigma}_0 = \mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZX^\top]$ and θ_0 be the minimum norm solution $\theta_0 = \tilde{\Sigma}_0^+ \mathbb{E}[ZY]$ to the IV problem $\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[Z(Y - \theta^\top X)] = 0$. Assume that X, Z, Y are uniformly and absolutely bounded. Given n samples of (Z, X, Y), let $\hat{\theta}$ be the solution to the adversarial IV method described in Algorithm 3, with $\lambda = \alpha/n$ chosen such that $\lambda = o(1)$ and $n\lambda \to \infty$. Then $\hat{\theta}$ satisfies:*

$$\|\hat{\theta} - \theta_0\|_w^2 \leq O_p\left(\lambda^2 + \frac{1}{n}(1 + \|\theta_0\|^2) \right) \qquad\qquad \text{(D.46)}$$

$$\|\hat{\theta} - \theta_0\|^2 \leq O_p\left(\lambda + \frac{1}{\lambda n}(1 + \|\theta_0\|^2) \right) \qquad\qquad \text{(D.47)}$$

In essence, this theorem states that $\hat{\theta}$ will eventually converge, in probability, to the true θ_0, i.e., that the estimate is consistent.

We will use the regularized IV algorithm to estimate both of our nuisance parameters h_*, γ_*, with the corresponding definitions of an instrument, treatment and outcome and the corresponding definitions of convergence metrics (note, then, that the parameter θ in this section then refers to these nuisance parameters). Recall under Theorem D.1, to prove asymptotic normality of our target parameter estimate $\hat{\theta}$ (i.e., the implicit bias effect) in the context of our main algorithm, we need the product of the errors of the primal estimate $\hat{h}$ and the dual

estimate $\hat{\gamma}$, both obtained by the adversarial IV method, to satisfy the convergence property:

$$\sqrt{n}(\hat{h} - h_0)^\top \mathbb{E}\left[XZ^\top\right](\hat{\gamma} - \gamma_*) = o_p(1) \tag{D.48}$$

Note that a sufficient condition is that:

$$\sqrt{n}\|\hat{h} - h_*\| \, \|\gamma_* - \hat{\gamma}\|_w = o_p(1)$$

It suffices to show that $n\|\hat{h} - h_*\|^2 \, \|\gamma_* - \hat{\gamma}\|_w^2 = o_p(1)$. Applying the results from Theorem D.2, we have that $n\|\hat{h} - h_*\|^2 \, \|\gamma_* - \hat{\gamma}\|_w^2$ is of order:

$$n\left(\frac{1}{n\lambda} + \lambda\right)\left(\frac{1}{n} + \lambda^2\right) = \left(\frac{1}{n\lambda} + \lambda\right)\left(1 + n\lambda^2\right) = \frac{1}{n\lambda} + 2\lambda + n\lambda^3$$

Thus if we choose $\lambda = o(n^{-1/3})$ and $\lambda = \omega(n^{-1})$, then $\frac{1}{n\lambda} + 2\lambda + n\lambda^3$ will be $o(1)$, and we have from Theorem D.1 that the primal and dual IV estimates are asymptotically normal. In our experiments we chose $\alpha \sim n^{0.3} \implies \lambda \sim 1/n^{0.7}$, which satisfies both properties.

Appendix D.10.2. *Convergence Rate of Estimate*

In this section we prove Theorem D.10.2 of the convergence of the regularized adversarial estimate $\hat{\theta}$ to the true estimate θ_0, with respect to an upper bound as a function of the penalty parameter λ.

First, let θ_0 be the minimum norm solution to the un-regularized minimization problem:

$$\min_\theta \|\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[(Y - \theta^\top X)Z]\|^2 \tag{D.49}$$

This minimization problem is assumed to have an optimal value of zero, or, equivalently, θ_0 is the minimum norm solution to the linear system $\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZX^\top]\theta = \mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZY]$, or $\tilde{\Sigma}_0\theta = \mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZY]$. Since $\mathbb{E}[ZZ^\top]$ is invertible, the minimum norm solution θ_0 also satisfies $\mathbb{E}[ZY] = \mathbb{E}[ZX^\top]\theta_0$.

Then, we define a parameter θ_λ as a solution to the *population* regularized adversarial criteria for a general parameter λ. We can then find the error of θ_λ from the true θ_0, both in terms of strong and weak norm (i.e., where the error is multiplied by Σ_0).

To define θ_λ, first consider the *population* adversarial criterion:

$$\min_\theta \max_\beta \mathbb{E}[2(Y - \theta^\top X)Z^\top\beta - 2\beta^\top ZZ^\top\beta] + \lambda\|\theta\|^2 \tag{D.50}$$

By similar argument to the empirical adversarial case, we can solve for the optimal β,

$$\beta_* = \frac{1}{2}\mathbb{E}[ZZ^\top]^{-1}\mathbb{E}[Z(Y - \theta^\top X)] \tag{D.51}$$

and plugging in, the inner optimal solution takes the form

$$\max_\beta \mathbb{E}[2(Y - \theta^\top X)Z^\top\beta - 2\beta^\top ZZ^\top\beta] + \lambda\|\theta\|^2 = \frac{1}{2}\mathbb{E}[(Y - \theta^\top X)Z^\top]\,\mathbb{E}[ZZ^\top]^{-1}\mathbb{E}[Z(Y - \theta^\top X)Z]$$

$$= \frac{1}{2}\|\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[(Y - \theta^\top X)Z]\|^2$$

We can furthermore rewrite the population criterion in terms of θ_0 where we see that:

$$\theta_\lambda = \arg\min_\theta \frac{1}{2}\|\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[(ZY - ZX^\top\theta)]\|^2 + \lambda\|\theta\|^2 \tag{D.52}$$

$$= \arg\min_\theta \frac{1}{2}\|\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZX^\top](\theta_0 - \theta)\|^2 + \lambda\|\theta\|^2 \tag{D.53}$$

$$= \arg\min_\theta \frac{1}{2}\|\tilde{\Sigma}_0(\theta - \theta_0)\|^2 + \lambda\|\theta\|^2 \triangleq L_\lambda(\theta) \tag{D.54}$$

We will first bound the error of our empirical regularized estimate $\hat{\theta}$ with respect to this population regularized estimate θ_λ. Then we will combine it with a bound on the error between θ_λ and the minimum norm solution θ_0 to conclude the proof.

First, note that θ_λ minimizes the loss $L_\lambda(\theta)$. Moreover, we show that the loss is strongly convex, which implies a bound on the distance between the true minimizer θ_λ and any other θ (including $\hat{\theta}$). In particular, we can invoke strong convexity of this loss function to upper bound the error of our estimate $\hat{\theta}$ with respect to θ_λ, as measured by a combination of a strong and a weak norm.

Lemma D.6 (Strong Convexity). *The loss function $L_\lambda(\theta) = \|\tilde{\Sigma}_0(\theta - \theta_0)\|^2 + 2\lambda\|\theta\|^2$ is strongly convex with:*

$$\partial_\theta^2 L(\theta) = 2(\tilde{\Sigma}_0^\top\tilde{\Sigma}_0 + 2\lambda I) \tag{D.55}$$

Since θ_λ optimizes the objective, we have:

$$L_\lambda(\theta) - L_\lambda(\theta_\lambda) \geq (\theta - \theta_\lambda)^\top(\tilde{\Sigma}_0^\top\tilde{\Sigma}_0 + 2\lambda I)(\theta - \theta_\lambda) = \|\tilde{\Sigma}_0(\theta - \theta_\lambda)\|^2 + 2\lambda\|\theta - \theta_\lambda\|^2 \tag{D.56}$$

Applying Lemma D.6 with $\theta = \hat{\theta}$, we have that:

$$, \|\tilde{\Sigma}_0(\hat{\theta} - \theta_\lambda)\|^2 + 2\lambda\|\hat{\theta} - \theta_\lambda\|^2 \leq \|\tilde{\Sigma}_0(\hat{\theta} - \theta_0)\|^2 - \|\tilde{\Sigma}_0(\theta_\lambda - \theta_0)\|^2 + 2\lambda(\|\hat{\theta}\|^2 - \|\theta_\lambda\|^2) \tag{D.57}$$

Recall the weak norm error is defined as $L(\theta) = \|\tilde{\Sigma}_0(\theta - \theta_0)\|$. Thus we need to upper bound the difference on the right hand side between the weak norm error $L(\theta_\lambda)$ of the population adversarial estimate θ_λ and the error $L(\hat{\lambda})$, evaluated at the empirical estimate $\hat{\theta}$. If we can prove such a bound exists, then we have the empirical adversarial estimate $\hat{\theta}$ satisfies an oracle inequality with respect to the weak norm error. The following lemma states that a bound exists as long as n is sufficiently large. We defer the proof to Appendix D.10.3.

Lemma D.7 (Oracle inequality). *Let θ_0 be the minimum norm solution to the linear system $\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[Z(Y - \theta^\top X)] = 0$, θ_λ the population ℓ_2-regularized adversarial estimate, and $\hat{\theta}$ the empirical ℓ_2-regularized adversarial estimate with appropriate choice of λ. Then the adversarial estimate satisfies the oracle inequality as long as $n \geq C\log(p_X \cdot p_Z/\delta)$ for some large enough universal constant C with probability $1 - \delta$:*

$$\|\tilde{\Sigma}_0(\hat{\theta} - \theta_0)\|^2 - \|\tilde{\Sigma}_0(\theta_\lambda - \theta_0)\|^2 \leq 7\|\tilde{\Sigma}_0(\theta_\lambda - \theta_0)\|^2 + 2\lambda(\|\theta_\lambda\|^2 - \|\hat{\theta}\|^2)$$

$$+ O\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}\|\tilde{\Sigma}_0(\hat{\theta} - \theta_\lambda)\|(1 + \|\hat{\theta} - \theta_\lambda\| + \|\theta_0\|) + \frac{\log(p_X \cdot p_Z/\delta)}{n}(1 + \|\theta_0\|)^2\right)$$

If we apply Lemma D.7 to the right hand side of Equation (D.57), we have:

$$\|\tilde{\Sigma}_0(\hat{\theta}-\theta_\lambda)\|^2+2\lambda\|\hat{\theta}-\theta_\lambda\|^2 \leq 7\|\tilde{\Sigma}_0(\theta_\lambda-\theta_0)\|^2+O\left(\sqrt{\frac{\log(p_X\cdot p_Z/\delta)}{n}}\|\tilde{\Sigma}_0(\hat{\theta}-\theta_\lambda)\|(1+\|\hat{\theta}-\theta_\lambda\|+\|\theta_0\|)\right)$$

$$+O\left(\frac{\log(p_X\cdot p_Z/\delta)}{n}(1+\|\theta_0\|)^2\right)$$

Applying the AM-GM inequality to the second term of the right hand side, we get:[s]

$$\frac{1}{2}\|\tilde{\Sigma}_0(\hat{\theta}-\theta_\lambda)\|^2+2\lambda\|\hat{\theta}-\theta_\lambda\|^2 \leq 7\|\tilde{\Sigma}_0(\theta_\lambda-\theta_0)\|^2+C\left(\frac{\log(p_X\cdot p_Z/\delta)}{n}(1+\|\hat{\theta}-\theta_\lambda\|^2+\|\theta_0\|^2)\right)$$

for some constant C. For $n\geq\frac{2}{3\lambda}C\log(p_X\cdot p_Z/\delta)$ (which will eventually hold if λ is chosen such that $n\lambda\to\infty$), we have that $C\frac{\log(p_X\cdot p_Z/\delta)}{n}\|\hat{\theta}-\theta_\lambda\|^2\leq\frac{3}{2}\lambda\|\hat{\theta}-\theta_\lambda\|^2$, hence we can bring this part from the right hand side to the left hand side and simplify the above as:

$$\frac{1}{2}\|\tilde{\Sigma}_0(\hat{\theta}-\theta_\lambda)\|^2+\frac{\lambda}{2}\|\hat{\theta}-\theta_\lambda\|^2 \leq 7\|\tilde{\Sigma}_0(\theta_\lambda-\theta_0)\|^2+C\left(\frac{\log(p_X\cdot p_Z/\delta)}{n}(1+\|\theta_0\|^2)\right) \qquad (D.58)$$

Thus we have derived a bound on the error between $\hat{\theta}$ and θ_λ as a function of the erorr between θ_λ and θ_0. We next bound the error of the regularized population adversarial solution θ_λ with respect to θ_0.

Lemma D.8 (Bias of θ_λ). *Let θ_0 be the minimum norm solution to the linear system $\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[Z(Y-\theta^\top X)]=0$. Let $\tilde{\Sigma}_0=UDV^\top$ by the SVD of $\tilde{\Sigma}_0$, with $D=diag(\sigma_1,\ldots,\sigma_K,0,\ldots,0)$ and K the rank of $\tilde{\Sigma}_0$. Let $a_0=V^\top\theta_0$. Then the population solution θ_λ using penalty parameter λ satisfies that:*

$$\|\theta_\lambda-\theta_0\|^2 \leq 4\lambda^2\sum_{i=1}^K\frac{a_{0,i}^2}{\sigma_i^4}=O(\lambda^2) \qquad \|\tilde{\Sigma}_0(\theta_\lambda-\theta_0)\|^2 \leq 4\lambda^2\sum_{i=1}^K\frac{a_{0,i}^2}{\sigma_i^2}=O(\lambda^2) \qquad (D.59)$$

Proof. Let $a=V^\top\theta$ and $a_0=V^\top\theta_0$, and note $\tilde{\Sigma}_0^\top\tilde{\Sigma}_0=VD^2V^\top$ and $\|\theta\|=\|V^\top\theta\|=\|a\|$. Then we rewrite the regularized problem defining θ_λ from Equation (D.52) as:

$$\min_a(a-a_0)^\top D^2(a-a_0)+2\lambda\|a\|^2=\sum_{i=1}^{p_X}\sigma_i^2(a_i-a_{0,i})^2+2\lambda a_i^2 \qquad (D.60)$$

Considering the first order condition optimi a_i, we can write the optimal solution as:

$$a_{\lambda,i}=\frac{\sigma_i^2}{\sigma_i^2+2\lambda}a_{i,0} \qquad (D.61)$$

---

[s]We use that $O(a\cdot b)\leq\frac{1}{2}a^2+O(b^2)$ for any $a,b>0$, and $O(a\cdot b)\leq C\cdot a\cdot b$ for some constant C. Applying the AM-GM inequality $C\cdot a\cdot b=a\cdot(C\cdot b)\leq\frac{1}{2}a^2+\frac{1}{2}C^2b^2=\frac{1}{2}a^2+O(b^2)$. Then the result follows be invoking this property for $a=\|\tilde{\Sigma}_0(\hat{\theta}-\theta_\lambda)\|$ and $b=\sqrt{\frac{\log(p_X\cdot p_Z/\delta)}{n}}(1+\|\hat{\theta}-\theta_\lambda\|+\|\theta_0\|)$. Moreover, $(a+b+c)^2=O(a^2+b^2+c^2)$.

Moreover, note that for $i > K$, $\sigma_i = 0$ and the minimum norm solution θ_0 has zero inner product with the eigenvectors $V_{\cdot,i}$. Thus, we have that:

$$\|\theta_\lambda - \theta_0\|^2 = \|a_\lambda - a_0\|^2 = \sum_{i=1}^{K} a_{i,0}^2 \left(1 - \frac{\sigma_i^2}{\sigma_i^2 + 2\lambda}\right)^2 = \sum_{i=1}^{K} a_{i,0}^2 \frac{4\lambda^2}{(\sigma_i^2 + 2\lambda)^2} \leq 4\lambda^2 \sum_{i=1}^{K} \frac{a_{i,0}^2}{\sigma_i^4} \quad \text{(D.62)}$$

Then the weak norm error $\|\tilde{\Sigma}_0(\theta_\lambda - \theta_0)\|$ can be written as:

$$\|\tilde{\Sigma}_0(\theta_\lambda - \theta_0)\|^2 = \sum_{i=1}^{K} \sigma_i^2 (a_{\lambda,i} - a_{0,i})^2 \leq 4\lambda^2 \sum_{i=1}^{K} \frac{a_{i,0}^2}{\sigma_i^2} \quad \text{(D.63)}$$

$\square$

We can now combine all these derived bounds to get the desired error rates.

Weak norm rate For the weak norm convergence rate (i.e., $\|\mathbb{E}[ZX^\top](\hat{\theta} - \theta_0)\|$), we can apply the triangle inequality and Equation (D.58) to get:

$$\|\tilde{\Sigma}_0(\hat{\theta} - \theta_0)\|^2 \leq \|\tilde{\Sigma}_0(\hat{\theta} - \theta_\lambda)\|^2 + \|\tilde{\Sigma}_0(\theta_\lambda - \theta_0)\|^2 \leq 2\|\tilde{\Sigma}_0(\hat{\theta} - \theta_\lambda)\|^2 + 2\|\tilde{\Sigma}_0(\theta_\lambda - \theta_0)\|^2$$

$$\leq O\left(\|\tilde{\Sigma}_0(\theta_0 - \theta_\lambda)\|^2 + \frac{\log(p_X \cdot p_Z/\delta)}{n}(1 + \|\theta_0\|^2)\right)$$

Subsequently, by Lemma D.8:

$$\|\tilde{\Sigma}_0(\hat{\theta} - \theta_0)\|^2 \leq O\left(\lambda^2 + \frac{\log(p_X \cdot p_Z/\delta)}{n}(1 + \|\theta_0\|^2)\right)$$

Recall $\tilde{\Sigma}_0 = \mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZX^\top]$. Let $K\Lambda K^\top$ be the eigendecomposition of $\mathbb{E}[ZZ^\top]$ and for any $x \in R^{p_Z}$, let $a_x = K^\top x$. Then we see that

$$\|\mathbb{E}[ZZ^\top]^{-1/2}x\|^2 = x^\top K\Lambda^{-1}K^\top x = a_x^\top \Lambda^{-1} a_x = \sum_{i=1}^{p_z} \frac{1}{\lambda_i} a_{x,i}^2 \geq \frac{1}{\lambda_{\max}} \sum_{i=1}^{p_z} a_{x,i}^2$$

$$= \frac{1}{\lambda_{\max}} a_x^\top a_x = \frac{1}{\lambda_{\max}} x^\top KK^\top x = \frac{1}{\lambda_{\max}} x^\top x = \frac{1}{\lambda_{\max}}\|x\|^2$$

where $\lambda_{\max}$ is the maximum eigenvalue of $\mathbb{E}[ZZ^\top]$ (which is non-zero). Thus we get that for any vector ν:

$$\|\tilde{\Sigma}_0\nu\|^2 = \|\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZX^\top]\nu\|^2 \geq \frac{1}{\lambda_{\max}}\|\mathbb{E}[ZX^\top]\nu\|^2 = \frac{1}{\lambda_{\max}}\|\nu\|_w^2$$

from which we deduce that:

$$\|\hat{\theta} - \theta_0\|_w^2 = \|\mathbb{E}[ZX^\top](\hat{\theta} - \theta_0)\|^2 \leq \lambda_{\max}\|\tilde{\Sigma}_0(\hat{\theta} - \theta_0)\|^2$$

$$= O\left(\lambda^2 + \frac{\log(p_X \cdot p_Z/\delta)}{n}(1 + \|\theta_0\|^2)\right)$$

Strong norm rate For the ℓ_2-norm convergence rate, we can invoke the triangle inequality and Equation (D.58)

$$\|\hat{\theta} - \theta_0\|^2 \leq \|\hat{\theta} - \theta_\lambda\|^2 + \|\theta_\lambda - \theta_0\|^2 \leq 2\|\hat{\theta} - \theta_\lambda\|^2 + 2\|\theta_\lambda - \theta_0\|^2$$

$$\leq O\left(\frac{1}{\lambda}\|\tilde{\Sigma}_0(\theta_\lambda - \theta_0)^2\|^2 + \|\theta_\lambda - \theta_0\|^2 + \frac{\log(p_X \cdot p_Z/\delta)}{\lambda n}(1 + \|\theta_0\|^2)\right)$$

Subsequently, by Lemma D.8:

$$\|\hat{\theta} - \theta_0\|^2 \leq O\left(\lambda + \lambda^2 + \frac{\log(p_X \cdot p_Z/\delta)}{\lambda n}(1 + \|\theta_0\|^2)\right)$$

$$= O\left(\lambda + \frac{\log(p_X \cdot p_Z/\delta)}{\lambda n}(1 + \|\theta_0\|^2)\right)$$

Thus, overall we can conclude the rates of the estimator:

$$\|\hat{\theta} - \theta_0\|_w^2 \leq O_p\left(\lambda^2 + \frac{1}{n}(1 + \|\theta_0\|^2)\right)$$

$$\|\hat{\theta} - \theta_0\|^2 \leq O_p\left(\lambda + \frac{1}{\lambda n}(1 + \|\theta_0\|^2)\right)$$

Appendix D.10.3. *Proof of Lemma D.7*

Proof. Let the optimal solution to the inner population optimization in the adversarial criterion when $\theta = \hat{\theta}$ be $\hat{\beta} = \frac{1}{2}\mathbb{E}[ZZ^\top]^{-1}\mathbb{E}[ZX^\top](\theta_0 - \hat{\theta}) = \frac{1}{2}\mathbb{E}[ZZ^\top]^{-1/2}\tilde{\Sigma}_0(\theta_0 - \hat{\theta})$. Since θ_0 also satisfies $\mathbb{E}[ZY] = \mathbb{E}[ZX^\top]\theta_0$, we see that:

$$A \triangleq \frac{1}{2}\|\tilde{\Sigma}_0(\theta_0 - \hat{\theta})\|^2 = 2\hat{\beta}^\top\mathbb{E}[ZX^\top](\theta_0 - \hat{\theta}) - 2\hat{\beta}^\top\mathbb{E}[ZZ^\top]\hat{\beta}$$

$$= 2\hat{\beta}^\top(\mathbb{E}[ZY] - \mathbb{E}[ZX^\top]\hat{\theta}) - 2\hat{\beta}^\top\mathbb{E}[ZZ^\top]\hat{\beta}$$

Since X, Z, Y are bounded, by a Chernoff bound and a union bound, we have that with probability $1 - \delta$:

$$\|\mathbb{E}[ZX^\top] - \mathbb{E}_n[ZX^\top]\|_\infty, \|\mathbb{E}[ZY] - \mathbb{E}_n[ZY]\|_\infty, \|\mathbb{E}[ZZ^\top] - \mathbb{E}_n[ZZ^\top]\|_\infty \leq O\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}\right)$$

Thus, we can replace the population covariances with the empirical covariances (leaving one population covariance $\hat{\beta}^\top\mathbb{E}[ZZ^\top]\hat{\beta}$ intact, which will serve its purpose later on in the proof):

$$A \leq 2\hat{\beta}^\top(\mathbb{E}_n[ZY] - \mathbb{E}_n[ZX^\top]\hat{\theta}) - \hat{\beta}^\top\mathbb{E}_n[ZZ^\top]\hat{\beta} - \hat{\beta}^\top\mathbb{E}[ZZ^\top]\hat{\beta} + \epsilon(n)$$

$$\leq \max_\beta 2\beta^\top\mathbb{E}_n[Z(Y - X^\top\hat{\theta})] - \beta^\top\mathbb{E}_n[ZZ^\top]\beta - \hat{\beta}^\top\mathbb{E}[ZZ^\top]\hat{\beta} + \epsilon(n)$$

where:

$$\epsilon(n) = O\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}(\|\hat{\beta}\| + \|\hat{\beta}\|^2 + \|\hat{\beta}\|\,\|\hat{\theta}\|)\right)$$

$$\leq C\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}(\|\hat{\beta}\|(1 + \|\hat{\theta}\|) + \hat{\beta}^\top\mathbb{E}[ZZ^\top]\hat{\beta})\right)$$

for some constant C.

For $n \geq C^2\log(p_X \cdot p_Z/\delta)$, we have that $C\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}\hat{\beta}^\top\mathbb{E}[ZZ^\top]\hat{\beta}$ (which will be a positive scalar) will be smaller than the absolute value of the negative term in the optimization criterion $-\hat{\beta}^\top\mathbb{E}[ZZ^\top]\hat{\beta}$. Thus these two terms can be ignored as they sum up to something negative. Thus, we can derive:

$$A \leq \max_\beta 2\beta^\top\mathbb{E}_n[Z(Y - X^\top\hat{\theta})] - \beta^\top\mathbb{E}_n[ZZ^\top]\beta + \zeta(n)$$

with:
$$\zeta(n) = O\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}\|\hat{\beta}\|(1 + \|\theta_0\| + \|\hat{\theta} - \theta_\lambda\|)\right)$$

where we also invoked the triangle inequality $\|\hat{\theta}\| \leq \|\theta_\lambda\| + \|\hat{\theta} - \theta_\lambda\|$ and the fact that by the bias Lemma D.8 $\|\theta_\lambda\| = O(\|\theta_0\|)$. Moreover, note that

$$\|\hat{\beta}\| = \left\|\frac{1}{2}\mathbb{E}[ZZ^\top]^{-1/2}\tilde{\Sigma}_0(\theta_0 - \hat{\theta})\right\| \tag{D.64}$$

$$= O\left(\|\tilde{\Sigma}_0(\theta_0 - \hat{\theta})\|\right) \tag{D.65}$$

$$= O(\sqrt{A}) \tag{D.66}$$

so we can rewrite

$$\zeta(n) = O\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}\sqrt{A}(1 + \|\theta_0\| + \|\hat{\theta} - \theta_\lambda\|)\right)$$

Since $\hat{\theta}$ optimizes the regularized empirical criterion $L_{n,\lambda}(\theta) \triangleq \max_\beta 2\beta^\top \mathbb{E}_n[Z(Y - X^\top\theta)] - \beta^\top \mathbb{E}_n[ZZ^\top] + \lambda\|\theta\|^2$ over all parameters θ, it means that $L_{n,\lambda}(\hat{\theta}) \leq L_{n,\lambda}(\theta_\lambda)$. Thus, we can derive:

$$A \leq \max_\beta 2\beta^\top \mathbb{E}_n[Z(Y - X^\top\theta_\lambda)] - \beta^\top \mathbb{E}_n[ZZ^\top]\beta + \zeta(n) + \lambda(\|\theta_\lambda\|^2 - \|\hat{\theta}\|^2)$$

$$\leq \max_\beta 2\beta^\top \mathbb{E}[Z(Y - X^\top\theta_\lambda)] - \beta^\top \mathbb{E}[ZZ^\top]\beta + \delta(n,\beta) + \zeta(n) + r(n)$$

with:

$$r(n) = \lambda(\|\theta_\lambda\|^2 - \|\hat{\theta}\|^2)$$

$$\delta(n,\beta) = O\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}(\|\beta\|(1 + \|\theta_\lambda\|) + \|\beta\|^2)\right)$$

$$\leq C\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}(\sqrt{\beta^\top \mathbb{E}[ZZ^\top]\beta}(1 + \|\theta_\lambda\|) + \beta^\top \mathbb{E}[ZZ^\top]\beta)\right)$$

for some constant C. For $n \geq 4C^2 \log(p_X \cdot p_Z/\delta)$ and since by the Lemma D.8 $\|\theta_\lambda\| = O(\|\theta_0\|)$:

$$\delta(n,\beta) \leq \frac{1}{2}\beta^\top \mathbb{E}[ZZ^\top]\beta + O\left(\sqrt{\frac{\log(p_X \cdot p_Z/\delta)}{n}}(\sqrt{\beta^\top \mathbb{E}[ZZ^\top]\beta}(1 + \|\theta_0\|)\right)$$

Applying the AM-GM inequality to the second term on the right hand side, we can further bound:

$$\delta(n,\beta) \leq \frac{3}{4}\beta^\top \mathbb{E}[ZZ^\top]\beta + O\left(\frac{\log(p_X \cdot p_Z/\delta)}{n}(1 + \|\theta_0\|)^2\right)$$

Letting $\kappa(n) = O\left(\frac{\log(p_X \cdot p_Z/\delta)}{n}(1 + \|\theta_0\|)^2\right)$, we can bound A as:

$$A \leq \max_\beta 2\beta^\top \mathbb{E}[Z(Y - X^\top\theta_\lambda)] - \frac{1}{4}\beta^\top \mathbb{E}[ZZ^\top]\beta + \zeta(n) + \kappa(n) + r(n)$$

Solving for the optimization problem, we derive an optimal solution of:

$$\max_{\beta} 2\beta^\top \mathbb{E}[Z(Y - X^\top \theta_\lambda)] - \frac{1}{4}\beta^\top \mathbb{E}[ZZ^\top]\beta = 4\|\tilde{\Sigma}_0(\theta_0 - \theta_\lambda)\|^2$$

where we used that θ_0 satisfies $\mathbb{E}[ZY] = \mathbb{E}[ZX^\top]\theta_0$.

Overall, we have derived that:

$$\|\tilde{\Sigma}_0(\theta_0 - \hat{\theta})\|^2 \le 8\|\tilde{\Sigma}_0(\theta_0 - \theta_\lambda)\|^2 + 2\zeta(n) + 2\kappa(n) + 2r(n)$$

$\square$

Appendix D.10.4. *Asymptotic Linearity of Estimate*

We prove that the error of the estimate $\hat{\theta}$, when projected onto the range of the covariance Σ_0, i.e. the quantity $\Sigma_0(\hat{\theta} - \theta_0)$, satisfies an asymptotic linearity property. Note that this does not necessarily imply that $\hat{\theta} - \theta_0$ is asymptotically linear, since Σ_0 is not necessarily invertible (in our application, it is most definitely rank-deficient).

Theorem D.3. *Assume that $\mathbb{E}[ZZ^\top]$ is full rank. Let $\Sigma_0 = \mathbb{E}[ZX^\top]$, $\tilde{\Sigma}_0 = \mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[ZX^\top]$, and θ_0 the minimum norm solution (i.e., $\theta_0 = \tilde{\Sigma}_0^+ \mathbb{E}[ZY]$) to the IV problem $\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}[Z(Y - \theta^\top X)] = 0$. Assume that X, Z, Y are uniformly and absolutely bounded. Given n samples of (Z, X, Y), let $\hat{\theta}$ be the solution to the adversarial IV method described in Algorithm 3, with $\lambda = \alpha/n = o(n^{-1/2})$ and $\lambda = \omega(n^{-1})$. Then $\hat{\theta}$ satisfies the following asymptotic linearity property:*

$$\sqrt{n}\Sigma_0(\hat{\theta} - \theta_0) = \sqrt{n}\mathbb{E}[ZZ^\top]^{1/2} \tilde{P} \,\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}_n\left[Z(Y - X^\top \theta_0)\right] + o_p(1) \qquad (\text{D}.67)$$

where $\tilde{P} = \tilde{\Sigma}_0\tilde{\Sigma}_0^+$ is the orthogonal projector onto the range of $\tilde{\Sigma}_0$.

Proof. Recall from Appendix D.10.1 that $\hat{\theta}$ solves the empirical moment equation

$$\mathbb{E}_n[\hat{Q}\,(Y - X^\top \theta)] + \lambda\theta = 0 \qquad (\text{D}.68)$$

where

$$\hat{Q} = \hat{B}^\top Z = (\mathbb{E}_n[ZZ^\top]^+\mathbb{E}_n[ZX^\top])^\top Z = \mathbb{E}_n[XZ^\top]\mathbb{E}_n[ZZ^\top]^+ Z \qquad (\text{D}.69)$$

Then, using the fact that $\Sigma_0\theta_0 = E[ZY]$, we have that

$$\Sigma_0(\hat{\theta} - \theta_0) = \Sigma_0\hat{\theta} - \Sigma_0\theta_0 = \Sigma_0\hat{\theta} - \mathbb{E}[ZY] = \mathbb{E}[Z(X^\top \hat{\theta} - Y)] \qquad (\text{D}.70)$$

Thus for any matrix B and hence for the projection matrix $\hat{B}$ (of Z onto X):

$$\begin{aligned}
\hat{B}^\top \Sigma_0(\hat{\theta} - \theta_0) =\ & \mathbb{E}[\hat{B}^\top Z(X^\top \hat{\theta} - Y)] & (\text{D}.71)\\
=\ & \mathbb{E}[\hat{Q}(X^\top \hat{\theta} - Y)] & (\text{D}.72)\\
=\ & \mathbb{E}[\hat{Q}(X^\top \hat{\theta} - Y)] - \mathbb{E}_n[\hat{Q}(X^\top \hat{\theta} - Y)] - \lambda\hat{\theta} & (\text{D}.73)
\end{aligned}$$

since $\mathbb{E}_n[\hat{Q}(X^\top \hat{\theta} - Y)] + \lambda\hat{\theta} = 0$. If we choose $\lambda = \alpha/n = o_p(n^{-1/2})$, then we have $\sqrt{n}\lambda = o_p(1)$. By Theorem D.2, and since by our assumption $\lambda = o(1)$ and $n\lambda = \omega(1) \to \infty$, we know that

$\|\hat\theta - \theta_0\| = o_p(1)$ and therefore $\|\hat\theta\| \le \|\theta_0\| + \|\hat\theta - \theta_0\| = O_p(\|\theta_0\|) = O_p(1)$. Thus, we have that $\sqrt{n}\lambda\|\hat\theta\| = o_p(1)$. Thus:

$$
\begin{aligned}
\sqrt{n}\hat{B}^\top\Sigma_0(\hat\theta - \theta_0) =\ & \sqrt{n}\mathbb{E}[\hat{Q}(X^\top\hat\theta - Y)] - \sqrt{n}\mathbb{E}_n[\hat{Q}(X^\top\hat\theta - Y)] + \sqrt{n}\lambda\hat\theta && \text{(D.74)} \\
=\ & \sqrt{n}\mathbb{E}[\hat{Q}(X^\top\hat\theta - Y)] - \sqrt{n}\mathbb{E}_n[\hat{Q}(X^\top\hat\theta - Y)] + o_p(1) && \text{(D.75)} \\
=\ & \sqrt{n}(\mathbb{E} - \mathbb{E}_n)\left[\hat{Q}(X^\top\hat\theta - Y)\right] + o_p(1) && \text{(D.76)} \\
=\ & \sqrt{n}(\mathbb{E} - \mathbb{E}_n)\left[\hat{Q}(X^\top\hat\theta - Y)\right] + && \text{(D.77)} \\
& (\sqrt{n}(\mathbb{E} - \mathbb{E}_n)\left[\hat{Q}X^\top\theta_0\right] - \sqrt{n}(\mathbb{E} - \mathbb{E}_n)\left[\hat{Q}X^\top\theta_0\right]) + o_p(1) && \text{(D.78)} \\
=\ & \sqrt{n}(\mathbb{E} - \mathbb{E}_n)\left[\hat{Q}(X^\top\theta_0 - Y)\right] + \sqrt{n}(\mathbb{E} - \mathbb{E}_n)\left[\hat{Q}X^\top\right](\hat\theta - \theta_0) + o_p(1) && \text{(D.79)}
\end{aligned}
$$

Note that the term $\sqrt{n}(\mathbb{E} - \mathbb{E}_n)[\hat{Q}X^\top] = \hat{B}^\top\sqrt{n}(\mathbb{E} - \mathbb{E}_n)[ZX^\top] = O_p(1)$. Moreover, since the regularized adversarial IV estimate converges to the minimum norm solution, we have $\|\hat\theta - \theta_0\| = o_p(1)$. Thus we have $\sqrt{n}(\mathbb{E} - \mathbb{E}_n)[\hat{Q}X^\top](\hat\theta - \theta_0) = o_p(1)$. This leads to:

$$
\begin{aligned}
\sqrt{n}\hat{B}^\top\Sigma_0(\hat\theta - \theta_0) =\ & \sqrt{n}(\mathbb{E} - \mathbb{E}_n)\left[\hat{Q}(X^\top\theta_0 - Y)\right] + o_p(1) \\
=\ & \sqrt{n}\hat{B}^\top(\mathbb{E} - \mathbb{E}_n)\left[Z(X^\top\theta_0 - Y)\right] + o_p(1)
\end{aligned}
$$

Moreover, since $\mathbb{E}[ZZ^\top]$ is full rank, then by the continuity of the inverse and the law of large numbers, $\hat{B} = \mathbb{E}_n[XZ^\top]\mathbb{E}_n[ZZ^\top]^{-1}$ converges in probability to the true projection matrix $B = \mathbb{E}[XZ^\top]\mathbb{E}[ZZ^\top]^{-1}$. Using the fact that $\mathbb{E}[Z(Y - \theta_0^\top X)] = 0$, we can also conclude that

$$
\begin{aligned}
\sqrt{n}B^\top\Sigma_0(\hat\theta - \theta_0) =\ & \sqrt{n}B^\top(\mathbb{E} - \mathbb{E}_n)\left[Z(X^\top\theta_0 - Y)\right] + o_p(1) \\
=\ & \sqrt{n}B^\top(-\mathbb{E}_n)\left[Z(X^\top\theta_0 - Y)\right] + o_p(1) \\
=\ & \sqrt{n}B^\top\mathbb{E}_n\left[Z(Y - X^\top\theta_0)\right] + o_p(1)
\end{aligned}
$$

Since $\Sigma_0^\top\mathbb{E}[ZZ^\top]^{-1} = B^\top$, we derive that:

$$
\sqrt{n}\Sigma_0^\top\mathbb{E}[ZZ^\top]^{-1}\Sigma_0(\hat\theta - \theta_0) = \sqrt{n}\Sigma_0^\top\mathbb{E}[ZZ^\top]^{-1}\mathbb{E}_n[Z(Y - X^\top\theta_0)] + o_p(1)
$$

Let $\tilde\Sigma_0 = \mathbb{E}[ZZ^\top]^{-1/2}\Sigma_0$. Then we can re-write the above as:

$$
\sqrt{n}\tilde\Sigma_0^\top\tilde\Sigma_0(\hat\theta - \theta_0) = \sqrt{n}\tilde\Sigma_0^\top\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}_n[Z(Y - X^\top\theta_0)] + o_p(1)
$$

Let the SVD of $\tilde\Sigma_0 = UDV^\top$, and then $\tilde\Sigma_0^\top\tilde\Sigma_0 = VD^2V^\top$ and $(\tilde\Sigma_0^\top)^+ = UD^+V^\top$. Thus $(\tilde\Sigma_0^\top)^+\tilde\Sigma_0^\top\tilde\Sigma_0 = UD^+V^\top VD^2V^\top = UD^+D^2V^\top = UDV^\top = \tilde\Sigma_0$. Thus left-multiplying the above equation with $(\tilde\Sigma_0^\top)^+$, we get:

$$
\sqrt{n}\tilde\Sigma_0(\hat\theta - \theta_0) = \sqrt{n}(\tilde\Sigma_0^\top)^+\tilde\Sigma_0^\top\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}_n[Z(Y - X^\top\theta_0)] + o_p(1)
$$

Moreover, $(\tilde\Sigma_0^\top)^+\tilde\Sigma_0^\top = UD^+V^\top VDU^\top = UD^+DU^\top = \tilde\Sigma_0\tilde\Sigma_0^+$. Thus we can equivalently write:

$$
\sqrt{n}\tilde\Sigma_0(\hat\theta - \theta_0) = \sqrt{n}\tilde\Sigma_0\tilde\Sigma_0^+\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}_n[Z(Y - X^\top\theta_0)] + o_p(1)
$$

Since $\tilde\Sigma_0 = \mathbb{E}[ZZ^\top]^{-1/2}\Sigma_0$, we can left-multiply both sides by $\mathbb{E}[ZZ^\top]^{1/2}$ to get that:

$$
\sqrt{n}\Sigma_0(\hat\theta - \theta_0) = \sqrt{n}\mathbb{E}[ZZ^\top]^{1/2}\tilde\Sigma_0\tilde\Sigma_0^+\mathbb{E}[ZZ^\top]^{-1/2}\mathbb{E}_n\left[Z(Y - X^\top\theta_0)\right] + o_p(1)
$$

$\square$

Appendix E. Data

All UK Biobank data can be explored in their accessible online search engine: `https://biobank.ndph.ox.ac.uk/showcase/search.cgi`.

For preprocessing our data, we compress all variables over time to a single feature. (Although the UK Biobank data is primarily static, there are a few follow-up analyses on a small subset of the original population, and thus some variables are collected multiple times.) We treat all categorical variables as multiple-categorical variables, i.e., a patient could check off multiple categories for a single feature. For W, Z, and X, categorical features are multi-hot encoded and zero-meaned (i.e., scaled to -0.5 or 0.5), and continuous features are standardized.

We additionally ran into moderate levels of missingness for a few features, a common problem in structured healthcare data. We left categorical data unchanged (as a value of zero in categorical data indicates both missing and unobserved) as we did not find high levels of missingness in these variables. For continuous data, we mean-imputed missing values.

Appendix E.1. *List of variables for W, Z, X*

Appendix E.1.1. *W variables*

For the sociodemographic confounders W, we used 34 features, which, after multi-hot encoding, resulted in 171 total binary and continuous variables that are listed in Table E1. Note that we include all possible patient attributes D in W except the current attribute we are analyzing for.

Table E1: W features

Feature	Variable Type	Cardinality
Age when attended assessment centre	Continuous	1
Alcohol intake frequency.	Categorical	7
Attendance/disability/mobility allowance	Categorical	6
Average total household income before tax	Categorical	7
Current employment status	Categorical	9
Distance between home and job workplace	Continuous	1
Duration of moderate activity	Continuous	1
Frequency of travelling from home to job workplace	Continuous	1
Gas or solid-fuel cooking/heating	Categorical	6
Hand grip strength (left)	Continuous	1
Hand grip strength (right)	Continuous	1
Heating type(s) in home	Categorical	9
How are people in household related to participant	Categorical	9
Job involves heavy manual or physical work	Categorical	6
Job involves mainly walking or standing	Categorical	6
Job involves night shift work	Categorical	6
Job involves shift work	Categorical	6
Length of time at current address	Continuous	1
Length of working week for main job	Continuous	1
Number in household	Continuous	1
Number of vehicles in household	Categorical	7
Own or rent accommodation lived in	Categorical	8
Private healthcare	Categorical	6
Education level	Categorical	8
Race	Categorical	2
Sex	Categorical	1
Smoking status	Categorical	4
Standing height	Continuous	1
Time employed in main current job	Continuous	1
Transport type for commuting to job workplace	Categorical	6
Type of accommodation lived in	Categorical	7
Types of physical activity in last 4 weeks	Categorical	7
UK Biobank assessment centre	Categorical	26
Weight	Continuous	1

Appendix E.1.2. *X variables*

For the outcome proxies X, we used 65 continuous variables representing several different biomarkers. These listed in Table E2.

Table E2: X features

Feature	Variable Type	Cardinality
Alanine aminotransferase	Continuous	1
Albumin	Continuous	1
Alkaline phosphatase	Continuous	1
Apolipoprotein A	Continuous	1
Apolipoprotein B	Continuous	1
Aspartate aminotransferase	Continuous	1
Basophill count	Continuous	1
Basophill percentage	Continuous	1
C-reactive protein	Continuous	1
Calcium	Continuous	1
Cholesterol	Continuous	1
Creatinine	Continuous	1
Cystatin C	Continuous	1
Diastolic blood pressure, automated reading	Continuous	1
Diastolic blood pressure, manual reading	Continuous	1
Direct bilirubin	Continuous	1
Eosinophill count	Continuous	1
Eosinophill percentage	Continuous	1
Gamma glutamyltransferase	Continuous	1
Glucose	Continuous	1
Glycated haemoglobin (HbA1c)	Continuous	1
HDL cholesterol	Continuous	1
Haematocrit percentage	Continuous	1
Haemoglobin concentration	Continuous	1
High light scatter reticulocyte count	Continuous	1
High light scatter reticulocyte percentage	Continuous	1
IGF-1	Continuous	1
Immature reticulocyte fraction	Continuous	1
LDL direct	Continuous	1

Table E3: *X* features, continued

Feature	Variable Type	Cardinality
Lipoprotein A	Continuous	1
Lymphocyte count	Continuous	1
Lymphocyte percentage	Continuous	1
Mean corpuscular haemoglobin	Continuous	1
Mean corpuscular haemoglobin concentration	Continuous	1
Mean corpuscular volume	Continuous	1
Mean platelet (thrombocyte) volume	Continuous	1
Mean reticulocyte volume	Continuous	1
Mean sphered cell volume	Continuous	1
Monocyte count	Continuous	1
Monocyte percentage	Continuous	1
Neutrophill count	Continuous	1
Neutrophill percentage	Continuous	1
Nucleated red blood cell count	Continuous	1
Nucleated red blood cell percentage	Continuous	1
Phosphate	Continuous	1
Platelet count	Continuous	1
Platelet crit	Continuous	1
Platelet distribution width	Continuous	1
Pulse rate (during blood-pressure measurement)	Continuous	1
Pulse rate, automated reading	Continuous	1
Red blood cell (erythrocyte) count	Continuous	1
Red blood cell (erythrocyte) distribution width	Continuous	1
Reticulocyte count	Continuous	1
Reticulocyte percentage	Continuous	1
Rheumatoid factor	Continuous	1
SHBG	Continuous	1
Systolic blood pressure, automated reading	Continuous	1
Systolic blood pressure, manual reading	Continuous	1
Total bilirubin	Continuous	1
Total protein	Continuous	1
Triglycerides	Continuous	1
Urate	Continuous	1
Urea	Continuous	1
Vitamin D	Continuous	1
White blood cell (leukocyte) count	Continuous	1

Appendix E.1.3. *Z variables*

For the proxies Z, we used 67 features, which, after multi-hot encoding, equals 196 total binary and continuous variables, and are listed in Table E4.

Table E4: Z features

Feature	Variable Type	Cardinality
Back pain for 3+ months	Categorical	1
Bipolar and major depression status	Categorical	6
Bipolar disorder status	Categorical	2
Chest pain due to walking ceases when standing still	Categorical	1
Chest pain or discomfort	Categorical	1
Chest pain or discomfort walking normally	Categorical	2
Chest pain or discomfort when walking uphill or hurrying	Categorical	2
Daytime dozing / sleeping (narcolepsy)	Categorical	4
Ever depressed for a whole week	Categorical	1
Ever highly irritable/argumentative for 2 days	Categorical	1
Ever manic/hyper for 2 days	Categorical	1
Ever unenthusiastic/disinterested for a whole week	Categorical	1
Facial pains for 3+ months	Categorical	1
Fed-up feelings	Categorical	1
Frequency of depressed mood in last 2 weeks	Categorical	4
Frequency of tenseness / restlessness in last 2 weeks	Categorical	4
Frequency of tiredness / lethargy in last 2 weeks	Categorical	4
Frequency of unenthusiasm / disinterest in last 2 weeks	Categorical	4
General pain for 3+ months	Categorical	1
Getting up in morning	Categorical	4
Guilty feelings	Categorical	1
Headaches for 3+ months	Categorical	1
Hip pain for 3+ months	Categorical	1
Illness, injury, bereavement, stress in last 2 years	Categorical	7
Irritability	Categorical	1
Knee pain for 3+ months	Categorical	1
Length of longest manic/irritable episode	Categorical	3
Loneliness, isolation	Categorical	1
Longest period of depression	Continuous	1
Longest period of unenthusiasm / disinterest	Continuous	1
Manic/hyper symptoms	Categorical	6
Miserableness	Categorical	1
Mood swings	Categorical	1
Morning/evening person (chronotype)	Categorical	4
Nap during day	Categorical	3
Neck/shoulder pain for 3+ months	Categorical	1
Nervous feelings	Categorical	1
Neuroticism score	Continuous	1
Number of depression episodes	Continuous	1
Number of unenthusiastic/disinterested episodes	Continuous	1
Pain type(s) experienced in last month	Categorical	9
Probable recurrent major depression (moderate)	Categorical	1
Probable recurrent major depression (severe)	Categorical	1
Risk taking	Categorical	1
Seen a psychiatrist for nerves, anxiety, tension or depression	Categorical	1
Seen doctor (GP) for nerves, anxiety, tension or depression	Categorical	1
Sensitivity / hurt feelings	Categorical	1
Severity of manic/irritable episodes	Categorical	2
Single episode of probable major depression	Categorical	1
Sleep duration	Continuous	1
Sleeplessness / insomnia	Categorical	3
Snoring	Categorical	1
Stomach/abdominal pain for 3+ months	Categorical	1
Suffer from 'nerves'	Categorical	1
Tense / 'highly strung'	Categorical	1
Worrier / anxious feelings	Categorical	1

Appendix E.2. *List of full D, Y pairs*

We use eight D attributes total, which were detailed in Table 1: Race=Asian, Race=Black, Gender=Female, On disability allowance, Low income, No post-secondary education, Clinically considered Obese (BMI > 30), and Not on private health insurance.[t]

In addition, we use thirteen Y medical diagnoses, where we use ICD10 codes as proxies and where the binary label of a diagnosis is aggregated across all of a patient's UK synced hospital records to represent if diagnosis Y was every given at any point.[u] In addition to the seven diagnoses Y mentioned in Table 1 (osteoarthritis, rheumatoid arthritis, chronic kidney disease, complications during labor, heart disease, depression, and melanoma), we also look at the diagnoses listed in Table E5: back pain, fibromyalgia, migraines, endometriosis, and inflammatory bowel disease.

In total, we look at 102 pairs (all D attributes × all Y attributes, except D =Female and $Y \in$ [endometriosis, compilations during labor]).

		Prevalence in UK Biobank (n=502411)
Medical diagnosis Y	Migraine	1.7%
	Inflammatory bowel disease	1.5%
	Back pain	6.4%
	Endometriosis	2.0%
	Fibromyalgia	0.7%

Table E5

Appendix F.   Additional results

Appendix F.1. *Semi-synthetic experiments*

Appendix F.1.1. *Hyperparameters*

Recall our data generation process uses the linear structural equations:

$$M_{\text{gen}} = aD_{\text{gen}} + \hat{\epsilon}_M \tag{F.1}$$

$$Y_{\text{gen}} = \frac{b}{K} \sum_{i=1}^{K} M_{i,\text{gen}} + \theta D_{\text{gen}} + gX_{0,\text{gen}} + \hat{f}_Y(W_{\text{gen}}) + \sigma_Y \hat{\epsilon}_Y \tag{F.2}$$

---

[t]Note that the UK offers universal healthcare, and thus public vs. private insurance is less of a proxy for income as in the US.

[u]Note that, unlike the US, ICD10 codes in the UK are not the primary source for billing information, and thus is likely a stronger proxy for diagnoses than using ICD10 codes in the US.

For each experiment, we select $a = 1$, $b = 1$, $g = 0$, and $\sigma_Y = 1$. For the implicit bias effect θ_0 we try $\theta_0 \in [-0.5, -0.1, 0.1, 0.5]$. We fit the propensity model $\mathbb{E}[D|W]$ with a logistic regression model using semi-cross fitting where the penalty parameter λ is first chosen from cross-validation to minimize the out-of-sample error. To estimate G, F, and $\text{Cov}(M)$ $= \Sigma = \text{diagonal}(\sigma_1, ..., \sigma_K)$, we compute the SVD of $\text{Cov}(\tilde{X}_{\text{train}}, \tilde{Z}_{\text{train}})$ where $\tilde{V}_{\text{train}}$ denotes 50% of the data from a train-test split. In addition, we keep only the non-zero singular values of $\text{Cov}(\tilde{X}_{\text{train}}, \tilde{Z}_{\text{train}})$, using our proposed covariance rank test from Appendix C.4, to generate $\text{Cov}(M)$. During inference, we use the data from the test split to estimate $\epsilon_X, \epsilon_Z, \epsilon_Y \sim \sigma_Y F_n(\tilde{Y}_{\text{test}})$ and generate new samples given the linear equations.

All V_{gen} for $V \in [X, Z, Y]$ are sampled as continuous. To test on semi-synthetic data that is more realistic of mixed-type medical data, we convert the sampled data to binary by thresholding all features which are binary in the real data. Specifically, since only Z contains binary data, if b_Z are the indices of all real binary Z features, we set $Z_{\text{gen,binary}}[:, b_Z] = Z_{\text{gen}}[:, b_Z] > 0$.

We fit our structural equations once and then run our method over $K = 100$ randomly generated semi-synthetic datasets of size $n = 50{,}000$.

For baselines, in addition to the two OLS models mentioned in Section 3.2.1, we also consider adversarial IV estimation using the (Primal Equation) where $(\tilde{D}_{\text{gen}}; \tilde{Z}_{\text{gen}})$ are the instruments and $(\tilde{D}_{\text{gen}}; \tilde{X}_{\text{gen}})$ are the treatments (note this is naturally run in our method when we estimate h_*; our method just "throws" out the computed θ). We also ran LASSO regression using cross validation to choose the penalty parameter λ but found very similar results to OLS and thus excluded for brevity.

Appendix F.1.2. *Evaluation metrics*

We report the following evaluation metrics:
- Average point estimate $\bar{\theta} = \frac{1}{K} \sum_i \theta_i$ over the $K = 100$ iterations
- Confidence interval (CI) $= \bar{\theta} \pm 1.96\sigma$, where σ is the standard deviation over all estimates θ_i
- Average CI $= \frac{1}{K} \sum_i 1.96\sigma_i$
- Coverage, or the percentage of the K iterations $\theta_i \pm 1.96\sigma_i$ that contains the true θ_0
- RMSE (root mean squared error) $= \sqrt{\frac{1}{K} \sum_i (\theta_i - \theta_0)^2}$
- Bias $= \left| \frac{1}{K} \sum_i (\theta_i - \theta_0) \right|$
- Average success rate of our five tests proposed in Section 2.3

Appendix F.1.3. *Results*

We report results for the mixed-type synthetic data, i.e., continuous and binary. In Table F1, we test $\theta_0 \in [-0.5, -0.1, 0.1, 0.5]$ and compare our method to the baselines. In Table F2 we report the aforementioned evaluation metrics of our method over different θ_0.

θ_0	Our method	OLS(D,W,M,X)	OLS(D,W,Z,X)	AdvIV$((\tilde{D};\tilde{Z}),(\tilde{D};\tilde{X}))$
-0.5	-0.46 ± 0.003	-0.5 ± 0.01	0.38 ± 0.01	-0.39 ± 0.002
-0.1	-0.07 ± 0.003	-0.1 ± 0.01	0.78 ± 0.01	0.00 ± 0.002
0.1	0.13 ± 0.003	0.1 ± 0.01	0.98 ± 0.01	0.19 ± 0.002
0.5	0.53 ± 0.003	0.5 ± 0.01	1.38 ± 0.01	0.58 ± 0.002

Table F1: Semi-synthetic data: our method versus baselines, reporting $\bar{\theta} \pm 1.96\sigma$ over $K = 100$ iterations for mixed-type data, i.e., continuous and binary.

θ_0	$\bar{\theta} \pm 1.96\sigma$	$\pm 1.96\sigma_i$	Coverage	RMSE	Bias	(1) % Passing Primal	(2) % Passing Dual	(3) % Passing $\mathbb{E}[\tilde{D}V] \neq 0$	(4) % Passing V strength F-test	(5) % Passing Cov$(\bar{X},\bar{Z})$ rank test
-0.5	-0.46 ± 0.0	0.07	0.78	0.05	0.04	91%	94%	100%	100%	100%
-0.1	-0.07 ± 0.0	0.07	0.80	0.05	0.03	93%	94%	100%	100%	100%
0.1	0.13 ± 0.0	0.07	0.82	0.05	0.03	95%	94%	100%	100%	100%
0.5	0.53 ± 0.0	0.07	0.85	0.05	0.03	95%	94%	100%	100%	100%

Table F2: Semi-synthetic data: evaluation metrics on our method over $K = 100$ iterations for mixed-type data, i.e., continuous and binary.

Appendix F.2. *UK Biobank effect estimates using all X, Z proxy data*

In Table F3 we show the results of our method on 20 (D,Y) pairs using all of the X, Z as the proxies for M. We show the effect estimate $\hat{\theta}$, 95% confidence interval, and the results from the first four tests we proposed. Note we exclude the covariance rank test; since we are using the same X, Z for all (D,Y) pairs, we only need to compute the test once and note that it passes with four statistically-significant non-zero singular values.

We see that, while the identification tests pass, both the dual and the primal violation tests fail for all (D,Y) pairs. This invalidates any estimates $\hat{\theta}$.

(D, Y)	$\theta \pm 95\%$ **CI**	**(1) Primal** statistic $<$ critical	**(2) Dual** statistic $<$ critical	**(3)** $\mathbb{E}[\check{D}V] \neq 0$ statistic $>$ critical	**(4)** V **strength F-test** statistic $>$ critical
Black, Osteoarthritis	0.18±0.24	422.8≮230.7	144.1≮84.8	7.0>0.3	37.8>23.1
Black, Heart disease	-0.03±0.13	243.4≮230.7	144.1≮84.8	7.0>0.3	37.8>23.1
Black, Depression	0.48±0.2	417.9≮230.7	144.1≮84.8	7.0>0.3	37.8>23.1
Black, Back pain	-0.43±0.13	344.5≮230.7	144.1≮84.8	7.0>0.3	37.8>23.1
Black, Rh. Arthritis	0.01±0.05	246.1≮230.7	144.1≮84.8	7.0>0.3	37.8>23.1
Female, Osteoarthritis	-0.14±0.31	440.0≮230.7	136.2≮84.8	9.5>1.1	33.4>23.1
Female, Heart disease	0.4±0.18	248.9≮230.7	136.2≮84.8	9.5>1.1	33.4>23.1
Female, Depression	0.27±0.27	411.1≮230.7	136.2≮84.8	9.5>1.1	33.4>23.1
Female, Back pain	0.66±0.17	361.5≮230.7	136.2≮84.8	9.5>1.1	33.4>23.1
Female, Rh. Arthritis	-0.04±0.06	248.1≮230.7	136.2≮84.8	9.5>1.1	33.4>23.1
Obese, Osteoarthritis	-0.31±0.05	407.9≮230.7	180.4≮84.8	64.9>1.3	131.2>23.1
Obese, Heart disease	-0.07±0.03	241.0≮230.7	180.4≮84.8	64.9>1.3	131.2>23.1
Obese, Depression	-0.06±0.04	403.7≮230.7	180.4≮84.8	64.9>1.3	131.2>23.1
Obese, Back pain	-0.09±0.02	354.9≮230.7	180.4≮84.8	64.9>1.3	131.2>23.1
Obese, Rh. Arthritis	-0.04±0.01	249.4≮230.7	180.4≮84.8	64.9>1.3	131.2>23.1
Asian, Osteoarthritis	0.04±0.1	438.9≮230.7	157.1≮84.8	10.5>0.3	40.4>23.1
Asian, Heart disease	-0.03±0.06	243.5≮230.7	157.1≮84.8	10.5>0.3	40.4>23.1
Asian, Depression	0.11±0.09	424.3≮230.7	157.1≮84.8	10.5>0.3	40.4>23.1
Asian, Back pain	0.0±0.05	324.3≮230.7	157.1≮84.8	10.5>0.3	40.4>23.1
Asian, Rh. Arthritis	-0.01±0.02	254.1≮230.7	157.1≮84.8	10.5>0.3	40.4>23.1

Table F3: Results using all X, Z proxies (no proxy selection algorithm). Note that the primal and dual tests fail, as the statistic is never $<$ the critical value, as is needed to pass.

Appendix F.3. *Discovery of invalid proxies and selection of non-violating proxy subsets*

As see in the results in Table F3, using all X, Z proxies leads to the violation of the primal and dual tests, regardless of the (D, Y) pair. Given the covariance rank test is still passing, we hypothesized this is due to the existence of some subset of Z proxies (say, Z') that have a causal path to Y that isn't through M, and a subset of X proxies (say, X') that have a causal path to Y that isn't M. We provide an example of the hypothesized causal graph of this scenario in Figure F1, where we remove confounders W for clarity.
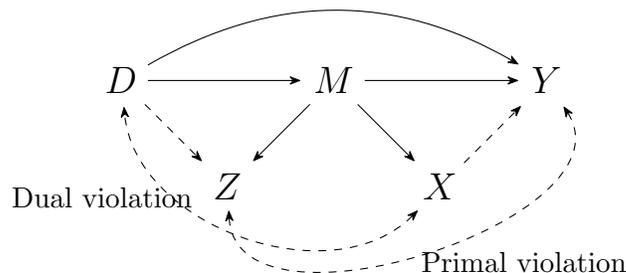
Fig. F1: Potential causal graph causing the violations of the primal and the dual. The violation of the primal is caused by some $Z' \to M_Z \to Y$, and the violation of the dual is caused by some X' such that $D \to M_X \to X'$

Given this hypothesis, we altered our data in two ways that proved integral to retrieving statistically significant estimates. First, we removed all Z proxy binary features indicating a survey response as 'Do not know' or 'Prefer not to answer' (e.g., Z_j representing the binary variable 'Back pain in last 3+ months = Do not know' would be removed). While similar ambiguous or non-responses can contain very real phenotypes in certain cases,[86] we observed they had little association with other X proxies, D, or Y. Second, we developed a proxy selection algorithm (see the full description of this algorithm in Appendix B) that finds an admissible set of proxy features $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$ as subsets of the original proxies such that we exclude any proxy features that the algorithm finds to be in violation of the assumptions.

For the proxy variable selection algorithm hyperparameters, we set $K = 150$, $\delta = 0.1$, and we run the algorithm for two iterations. 68 of the 102 (D, Y) pairs find no admissible proxy sets.[v] We found 34 (D, Y) pairs that find one or more candidate proxy pairs $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$.

Appendix F.3.1. *Interpretability into the proxy selection algorithm*

We next investigate which proxies were selected as admissible and if these proxies share any meaningful biological significance. We run several interpretability analysis:

- For a given (D, Y) pair, are the proxy candidates $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$ similar for all candidates i?
- Are the proxies biologically relevant to D or Y?
- Is there a pattern which proxies are selected with respect to D or Y?

First, we look for any consistency in the selected candidates for a given (D, Y) pair that have multiple candidates selected. We find most, such as D =Female, Y =Heart disease, to have almost perfect consistency in X and Z proxies selected. Other pairs, however, such as D =Black, Y =Chronic kidney disease, had high variance in the proxy subsets chosen and warrants further investigation.

Next we examined if the selected proxies are biologically reasonable. We can sort X (or Z) features based on how often the feature is selected into $\mathcal{X}^{(i)} (or \mathcal{Z}^{(i)})$ across all candidates i for a given (D, Y) pair. Although D =Black, Y =Chronic kidney disease had high variance in candidate features selected, a qualitative analysis showed the features commonly selected were biologically reasonable. Two of the most select biomarker features of X, aspartate aminotransferase (AST) and alanine aminotransferase (ALT), are enzymes found commonly in healthy livers and are often biomarkers used to check for kidney disease. We also probed Z proxies for biological significance but find most (D, Y) pairs are consistent in which proxies are selected versus removed. For example, Z proxies related to a patient's self-reported bipolar disorder status or daily lethargy consistently were selected into the Z proxy subset, while reports of neck, shoulder, or chest pain were rarely selected. We give three examples of the most common proxy features selected in Table F4.

---

[v]Increasing K or varying δ might produce an admissible set in this case, but we did not explore these further.

(D, Y)	Top 10 Z proxies selected	Top 10 X proxies selected
Disability Insurance, Rh. arthritis	Bipolar disorder status=Bipolar Type I (Mania), Bipolar and major depression status=Single Probable major depression episode, Manic/hyper symptoms=I needed less sleep than usual, Manic/hyper symptoms=I was more creative or had more ideas than usual, Manic/hyper symptoms=All of the above, Frequency of depressed mood in last 2 weeks=Nearly every day, Length of longest manic/irritable episode=A week or more, Length of longest manic/irritable episode=Less than a week, Illness, injury, bereavement, stress in last 2 years=Death of a spouse or partner, Bipolar and major depression status=Probable Recurrent major depression (severe)	White blood cell (leukocyte) count, Albumin, Aspartate aminotransferase, Pulse rate, Basophill percentage, Immature reticulocyte fraction, High light scatter reticulocyte percentage, Monocyte percentage, Nucleated red blood cell percentage, Vitamin D
Female, Heart disease	Frequency of tiredness / lethargy in last 2 weeks=Not at all, Frequency of unenthusiasm / disinterest in last 2 weeks=Not at all, Bipolar and major depression status=Bipolar II Disorder, Bipolar and major depression status=Probable Recurrent major depression (severe), Bipolar and major depression status=Probable Recurrent major depression (moderate), Bipolar and major depression status=Single Probable major depression episode, Loneliness, isolation=Yes, Guilty feelings=Yes, Frequency of depressed mood in last 2 weeks=Not at all, Frequency of depressed mood in last 2 weeks=Several days	Red blood cell (erythrocyte) count, LDL direct, Direct bilirubin, Mean platelet (thrombocyte) volume, Monocyte count, C-reactive protein, Gamma glutamyltransferase, Mean sphered cell volume, Aspartate aminotransferase, Lipoprotein A
Black, Chronic kidney disease	Bipolar and major depression status=Bipolar II Disorder, Bipolar disorder status=Bipolar Type II (Hypomania), Daytime dozing / sleeping (narcolepsy)=All of the time, Bipolar and major depression status=Bipolar I Disorder, Bipolar disorder status=Bipolar Type I (Mania), Illness, injury, bereavement, stress in last 2 years=Death of a spouse or partner, Manic/hyper symptoms=I needed less sleep than usual, Chest pain or discomfort walking normally=Unable to walk on the level, Single episode of probable major depression=Yes,	Total protein, Pulse rate, Aspartate aminotransferase, Alanine aminotransferase, SHBG, Rheumatoid factor, Urate, Monocyte count, IGF-1

Table F4: Top 10 proxies selected for three (D, Y) pairs that generate multiple proxy set candidates. X and Z features are sorted based on how often they were selected by each $\mathcal{X}^{(i)}$ or $\mathcal{Z}^{(i)}$, respectively, where $\mathcal{X}^{(i)}$ is a candidate of admissible X proxy features for a given (D, Y) pair.

To select a single proxy feature set given a list of candidates for a (D, Y) pair, we compute the final point estimate for all candidates i in $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$ and we choose the median estimate and its corresponding proxy features. We then compare the proxies selected across (D, Y) pairs by taking the correlation of the binary matrix indicating if a feature was selected as admissible for each (D, Y) pair or not. To only consider cases where the proxies contributed to a non-zero bias effect being detected, we filter the (D, Y) pairs if $|\theta| > 0.05$.

In Tables F2 and F3 we show the correlation clustermap for X and Z features, respectively. We observe a strong similarity of X proxies selected based on the attribute D. For example, the admissible proxy subsets S_X for $D =$Black X are almost identical. We believe this finding supports our hypothesis as depicted in Figure F1 that there exists a violating path $D \to X'$, as our algorithm appears to remove the same violating X' proxies for a specific attribute D. For selected proxies Z, there are less evident patterns of features selected across D or Y, although the clustermap slighly indicates Z proxies are selected according to Y diagnoses. Instead, it appears that most (D, Y) pairs have a non-trivial positive correlation with each other. This could indicate there is a subset Z' of Z features that have little association with any of the X proxies available in the UK Biobank. This violation could be remedied by having a richer set of biomarker proxies X.

In Table F6 we list the number of X, Z proxies chosen for each (D, Y) pair.
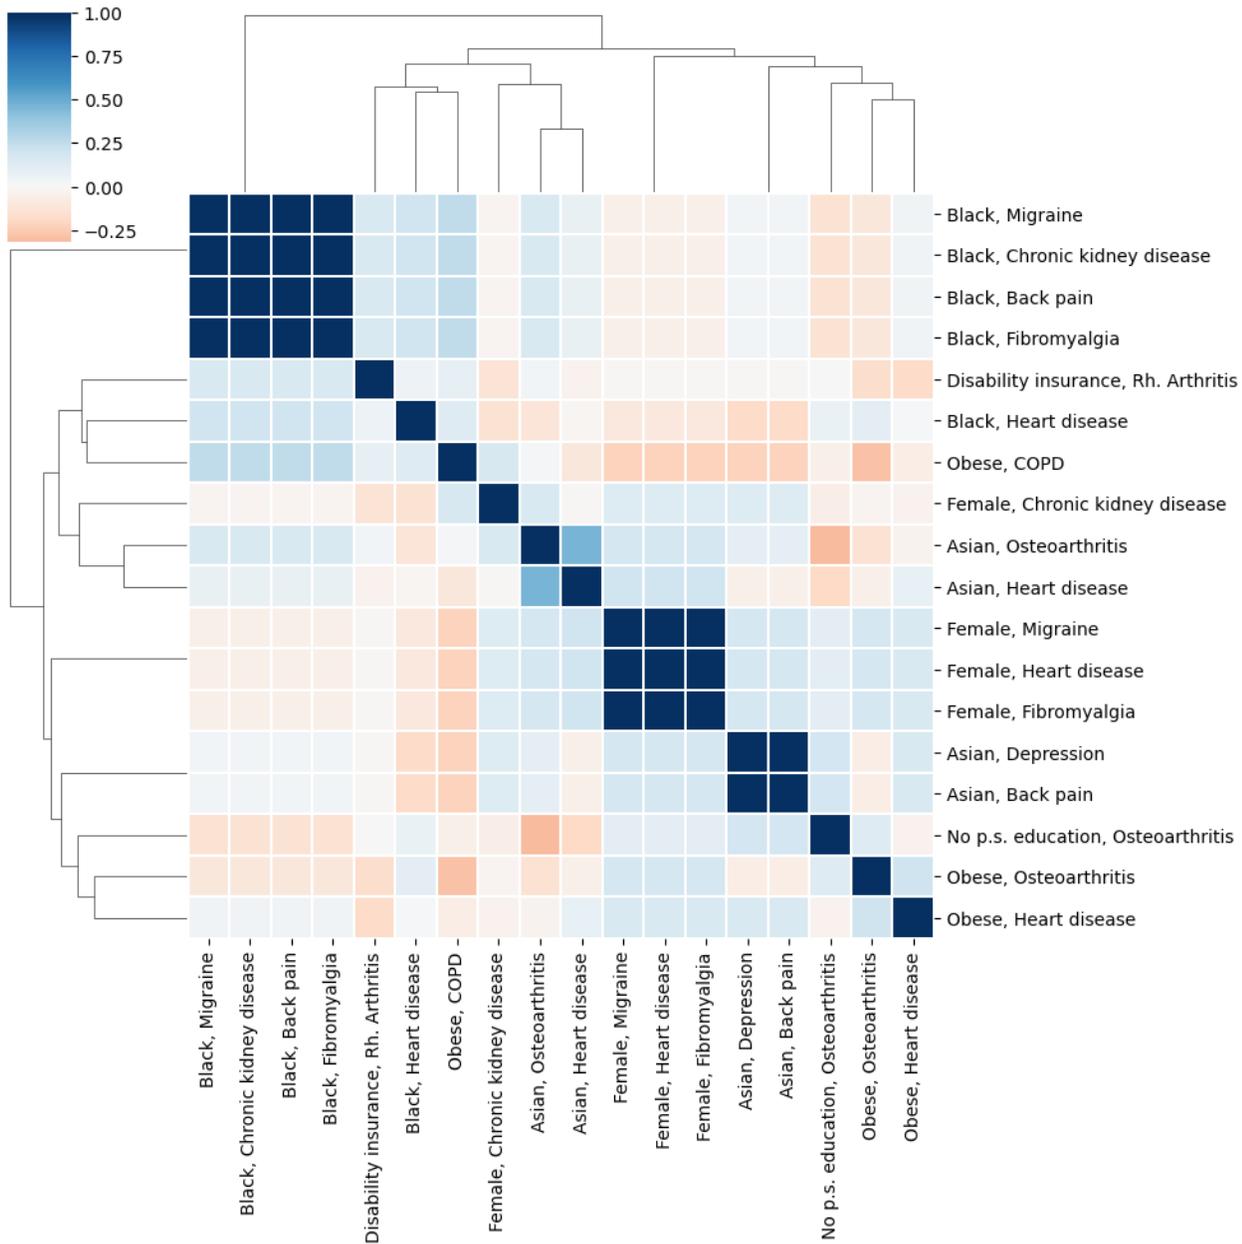
Fig. F2: Correlation clustermap comparing proxy X features selected by the proxy removal algorithm across (D, Y) pairs that (1) yielded admissible $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$ candidates and (2) whose median point estimate $|\theta| > 0.05$. COPD = Chronic obstructive pulmonary disease. p.s. = post-secondary.

Appendix F.4. *All valid estimates after running our proxy selection algorithm*

After running the proxy selection algorithm, we found 34 (D, Y) pairs that produced admissible proxy subsets $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$ which pass all five of our proposed tests and thus produce valid effect estimates for θ. In Table F5, we show the results for the remaining 28 (the other six were presented in Table 3). We highlight strong bias esti-
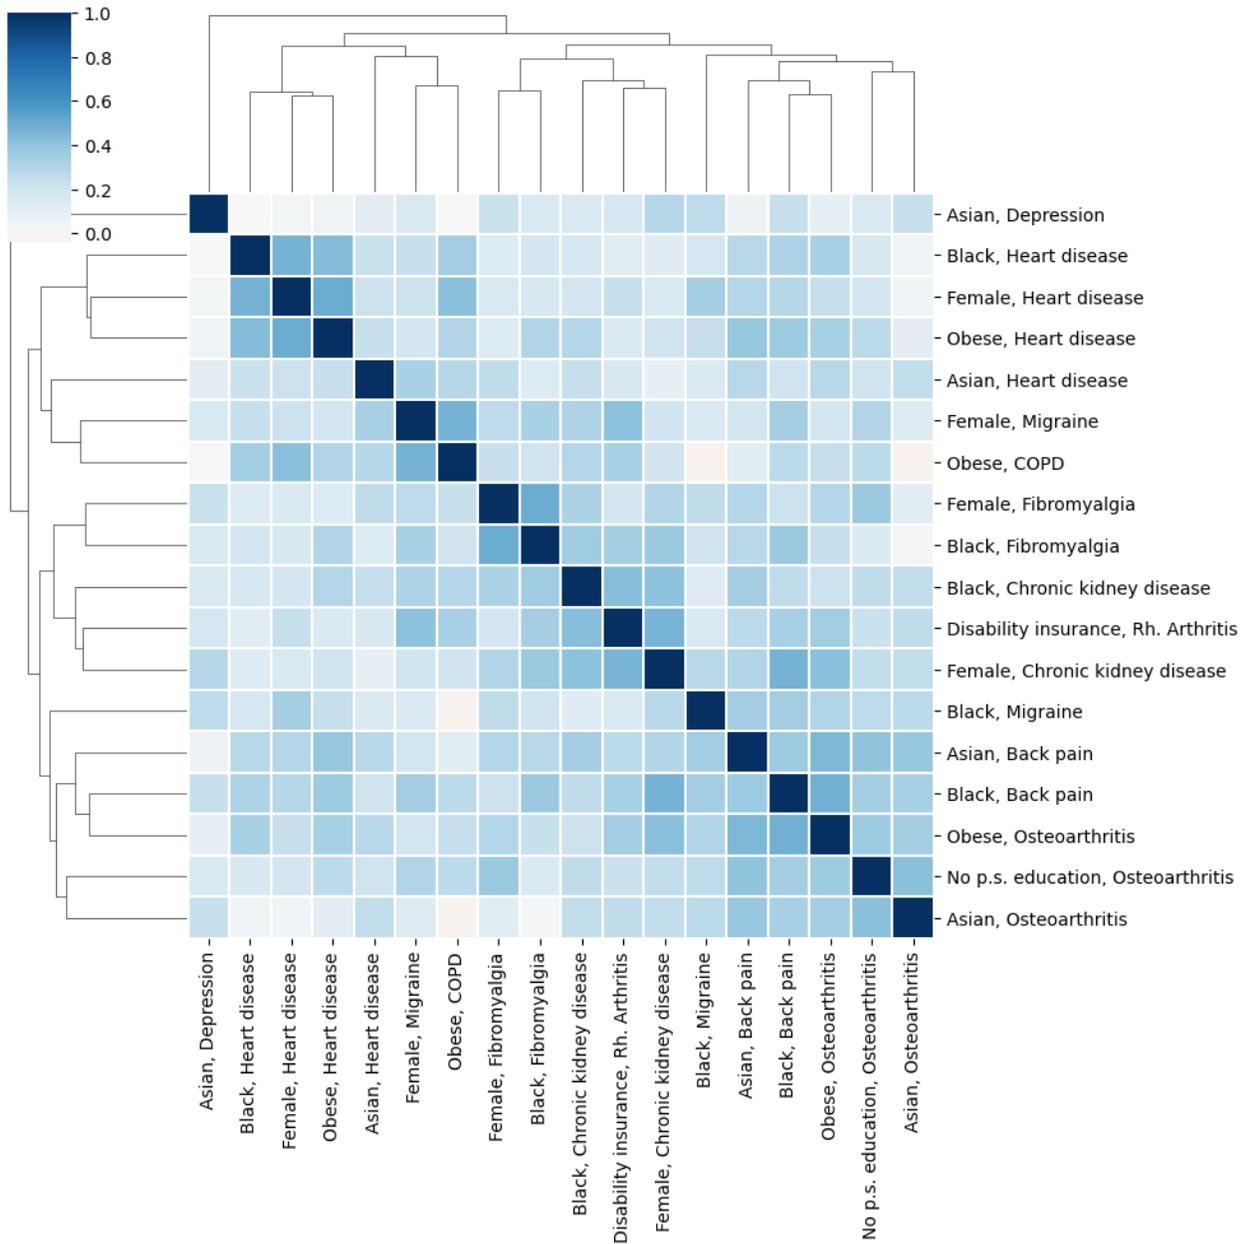
Fig. F3: Correlation clustermap comparing proxy Z features selected by the proxy removal algorithm across (D, Y) pairs that (1) yielded admissible $\{(\mathcal{X}^{(i)}, \mathcal{Z}^{(i)})\}_i$ candidates and (2) whose median point estimate $|\theta| > 0.05$. COPD = Chronic obstructive pulmonary disease. p.s. = post-secondary.

mates for (D=Female, Y=Migraine) with θ=-0.14; (D=Female, Y=Chronic kidney disease) with θ=-0.09; (D=Female, Y=Fibromyalgia) with θ=-0.08; (D=No post-secondary education, Y=Osteoarthritis) with θ=0.06; (D=Black, Y=Back pain) with θ=0.16; (D=Black, Y=Fibromyalgia) with θ=0.06; (D=Black, Y=Migraine) with θ=0.08; (D=Obese, Y=Heart disease) with θ=-0.06; (D=Obese, Y=COPD) with θ=-0.08; (D=Asian, Y=Heart disease)

with θ=0.09; (D=Asian, Y=Depression) with θ=0.32; and (D=Asian, Y=Back pain) with θ=0.09.

(D, Y)	$\theta \pm 95\%$ **CI**	**(1) Primal** statistic < critical	**(2) Dual** statistic < critical	**(3)** $\mathbb{E}[\tilde{D}V] \neq 0$ statistic > critical	**(4)** V **strength F-test** statistic > critical	**(5) Cov**$(\tilde{X}, \tilde{Z})$ rank
Low income, Endometriosis	0.0±0.0	59.2<76.8	28.1<31.4	64.6>0.5	1947.8>23.1	2
Disability insr., Back pain	0.03±0.01	91.1<97.4	31.6<40.1	27.6>0.3	402.6>23.1	4
Disability insr., Complications during labor	0.0±0.0	63.3<66.3	15.9<25.0	20.7>0.4	739.5>23.1	2
Not on private insr., Chronic kidney disease	0.0±0.0	19.3<33.9	16.4<22.4	59.1>0.3	82373.3>23.1	2
Not on private insr., Endometriosis	0.0±0.0	51.3<58.1	7.4<11.1	43.4>0.3	3546.8>23.1	2
No p.s. education, Osteoarthritis	0.06±0.01	92.7<93.9	37.8<38.9	123.6>1.3	150.5>23.1	5
No p.s. education, IBD	0.0±0.0	80.1<97.4	39.6<42.6	119.7>1.3	91.4>23.1	4
No p.s. education, COPD	0.0±0.0	63.5<71.0	41.8<42.6	125.6>0.6	676.6>23.1	4
Female, Fibromyalgia	-0.08±0.03	75.6<84.8	20.5<23.7	21.2>1.6	58.3>23.1	4
Female, Chronic kidney disease	-0.09±0.04	61.9<71.0	19.6<21.0	22.0>2.1	42.1>23.1	4
Female, Migraine	-0.14±0.06	76.3<79.1	22.7<23.7	20.8>1.7	43.5>23.1	4
Female, Melanoma	0.02±0.02	84.8<88.3	23.1<23.7	19.7>1.5	81.9>23.1	4
Black, Heart disease	0.1±0.03	97.7<115.4	25.5<31.4	9.7>0.2	216.8>23.1	4
Black, Back pain	0.16±0.04	84.0<93.9	15.4<21.0	10.1>0.3	68.7>23.1	5
Black, Fibromyalgia	0.06±0.02	68.6<72.2	16.1<21.0	9.8>0.3	44.4>23.1	4
Black, IBD	0.01±0.01	56.6<84.8	20.8<21.0	9.4>0.3	28.0>23.1	4
Black, Migraine	0.08±0.01	104.3<105.3	18.1<21.0	10.9>0.2	118.3>23.1	5
Black, Melanoma	-0.02±0.01	84.6<84.8	13.3<21.0	10.2>0.3	73.1>23.1	4
Obese, Heart disease	-0.06±0.01	112.9<113.1	24.8<30.1	67.4>1.6	184.1>23.1	4
Obese, Fibromyalgia	0.01±0.01	56.3<72.2	28.5<28.9	63.9>1.6	156.0>23.1	4
Obese, COPD	-0.08±0.01	95.5<97.4	24.1<27.6	69.0>1.4	293.7>23.1	4
Asian, Heart disease	0.09±0.02	108.1<109.8	23.6<27.6	14.2>0.3	163.1>23.1	5
Asian, Depression	0.32±0.05	56.2<73.3	14.6<23.7	14.4>0.2	54.4>23.1	4
Asian, Back pain	0.09±0.02	96.0<100.7	23.7<23.7	14.4>0.2	211.7>23.1	4
Asian, Rh. Arthritis	-0.0±0.0	65.9<66.3	11.4<15.5	14.9>0.2	689.8>23.1	4
Asian, Fibromyalgia	0.04±0.01	57.3<61.7	15.6<16.9	14.0>0.3	102.3>23.1	2
Asian, Chronic kidney disease	0.01±0.02	45.5<48.6	14.4<16.9	14.3>0.3	38.1>23.1	3
Asian, Endometriosis	0.02±0.01	86.1<88.3	31.3<32.7	9.9>0.2	208.4>23.1	3

Table F5: The remaining 28 valid UK Biobank implicit bias effect estimates after applying our X, Z proxy selection algorithm. Tests (1-5) are detailed in Section 2.3, where *statistic* is the given data's statistic and *critical* is the necessary critical value to be greater or less than to pass. $V = \tilde{D} - \gamma^\top \tilde{Z}$. COPD = Chronic obstructive pulmonary disease. IBD = Inflammatory bowel disease. p.s. = post-secondary. insr. = insurance.

Appendix F.5. *Weak instrument confidence interval for all valid estimates*

We ran the weak instrument confidence interval test on the remaining 28 (D, Y) pairs after proxy removal. We see in Table F4 that the test yields intervals consistent with those from our method, demonstrating our estimate is robust to weak instruments.

Appendix F.6. *Additional bootstrap analysis*

Here we show the two other bootstrapping analyses on the original six (D, Y) pairs and mentioned in Section 3.2.2. In Figure F5A we show the influence of sampling $10\%, 25\%, 50\%$ or 75% of the original data size for $K = 10$ bootstrapped iterations, re-estimating over the full pipeline (stage 1). Similarly, in Figure F5B we show the influence of sampling $10\%, 25\%, 50\%$ or 75% of the original data size for $K = 1000$ bootstrapped iterations, re-estimating over the final estimate only (stage 3). We see relatively consistent estimates regardless of sample size and
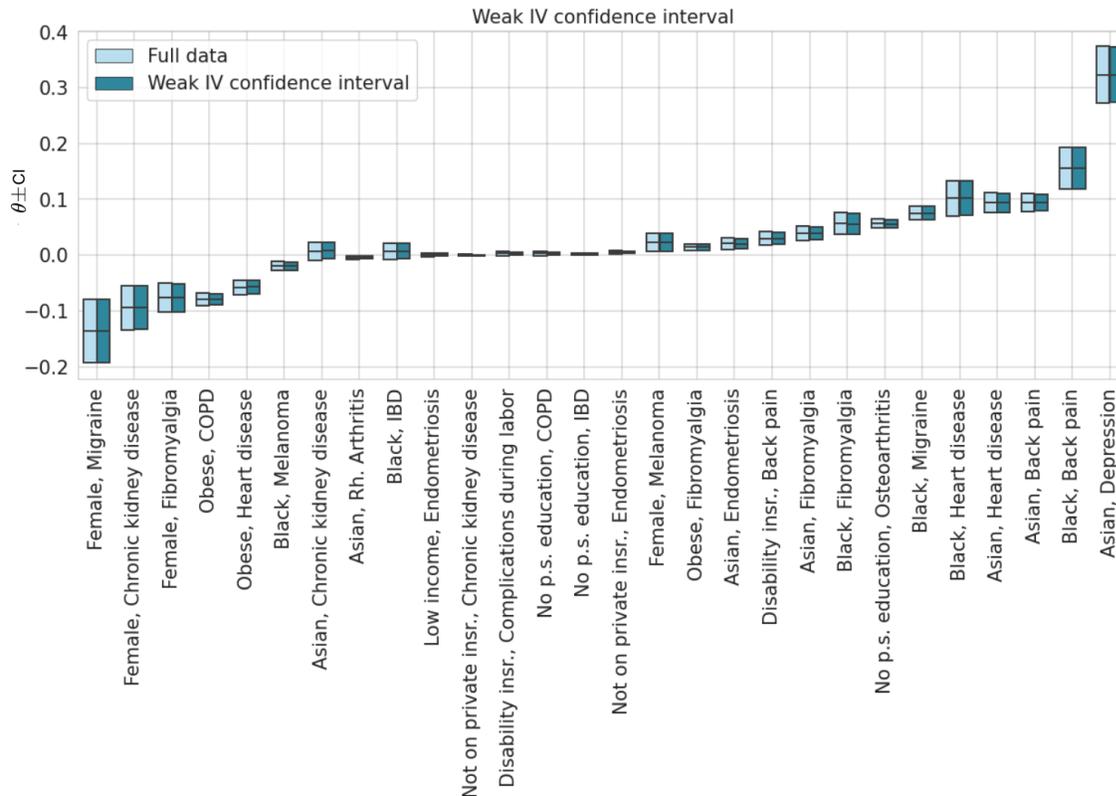
Fig. F4: Weak instrumental variable (IV) confidence interval test. COPD = Chronic obstructive pulmonary disease. IBD = Inflammatory bowel disease. p.s. = post-secondary. insr. = insurance.

re-estimation stage. As expected, the confidence interval increases as sample size decreases.

Appendix F.7. *Additional influence points results*

In Table F6, we list the size of the high-influence set for each (D, Y) pair such that the estimated estimate after the set's removal is $\theta_{\text{inf}} = 0$. Most of the sizes of the minimally influential set are decently small (recall $n = 502411$).

We perform two interpretability analyses to better understand these results. First, we compare high-influence sets to see if similar points are being scored as highly influential across (D, Y) pairs. We take the Szymkiewicz-Simpson coefficient $(\frac{|A \cap B|}{\min(|A|, |B|)})$ for the high-influence set across all (D, Y) pairs where $| \theta | > .01$ The resulting heatmap in Figure F6 shows a potential pattern for overlap among diagnoses Y, particularly for $Y \in$ [Fibromyalgia, Chronic kidney disease, Melanoma]. This could potentially indicate that, for these Y diagnoses, the same "outlier" patients are driving the detected implicit biases.

To investigate what phenotypes characterize these identified high-influence patients, we perform a simple feature interpretability analysis, as described in Appendix C.6.3. Similar to Figure 3A, we report significant features between high-influence set versus the rest of the data for several other (D, Y) pairs in Figure F7. These features discrepancies could be poten-
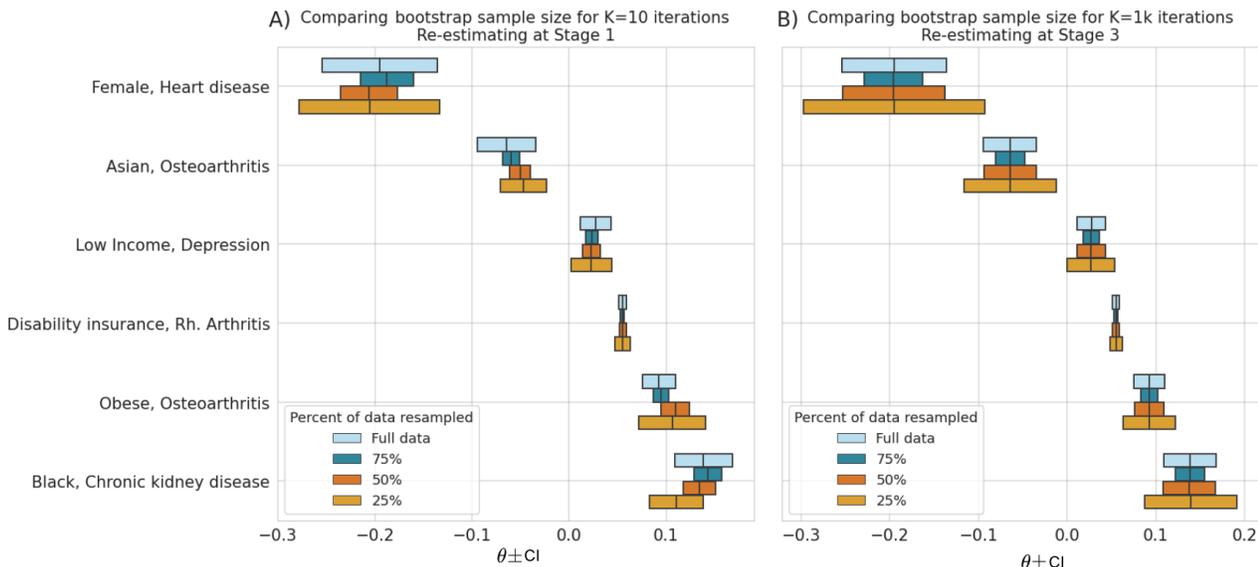
Fig. F5: Bootstrap analyses for six (D, Y) pairs. A) Effect of sample size on bootstrapped estimate for $K=10$ iterations re-estimating from stage 1. B) Effect of sample size on bootstrapped estimate for $K=1000$ iterations re-estimating from stage 3.

tial phenotypes that characterize "outlier" patient groups that drive the implicit bias effect estimate.

Appendix F.8. *Additional income stratification results*

In Figure 3B we show the intersectional effect of income on implicit bias for three (D, Y) pairs where $D \neq$Income. We re-estimate the new estimate θ from stage 2 of the pipeline (see Section 3.2.2) on the patients corresponding to one of two income strata: Low income=Average total household income before tax is less than 18,000£and High income=Average total household income before tax is greater than 100,000£. We report the remaining (D, Y) pairs in Figure F8 and show there are variations in implicit bias estimates based on income strata.

Appendix F.9. *Partial non-linearity of W*

We model the residuals $\mathbb{E}[V|W]$ for $V \in [Y, D, X, Z]$ using XGBoost regression[65] where the learning rate is chosen via semi-cross fitting. For the six (D, Y) pairs we use the same admissible proxy sets for X and Z as mentioned previously and recompute the estimate θ and the 95% confidence interval using the new XGBoost residuals. In Table F7, we see the XGBoost residuals perform equivalently to using Lasso residuals for our data.

(D, Y)	$\theta \pm 95\%$ **CI**	**Inf. set size**	**#Z proxies**	**#X proxies**
Low income, Depression	0.03±0.02	190	43	27
Disability insr., Rh. Arthritis	0.06±0.0	1466	56	5
Female, Heart disease	-0.19±0.06	368	94	14
Black, Chronic kidney disease	0.14±0.03	511	41	12
Obese, Osteoarthritis	0.09±0.02	1355	78	18
Asian, Osteoarthritis	-0.06±0.03	218	79	22
Low income, Endometriosis	0.0±0.0	9	57	20
Disability insr., Back pain	0.03±0.01	320	75	27
Disability insr., Complications during labor	0.0±0.0	30	48	15
Not on private insr., Chronic kidney disease	0.0±0.0	8	21	13
Not on private insr., Endometriosis	0.0±0.0	148	41	5
No p.s. education, Osteoarthritis	0.06±0.01	1650	72	26
No p.s. education, IBD	0.0±0.0	23	75	29
No p.s. education, COPD	0.0±0.0	78	52	29
Female, Fibromyalgia	-0.08±0.03	301	64	14
Female, Chronic kidney disease	-0.09±0.04	113	52	12
Female, Migraine	-0.14±0.06	382	59	14
Female, Melanoma	0.02±0.02	90	67	14
Black, Heart disease	0.1±0.03	477	91	20
Black, Back pain	0.16±0.04	499	72	12
Black, Fibromyalgia	0.06±0.02	155	53	12
Black, IBD	0.01±0.01	18	64	12
Black, Migraine	0.08±0.01	724	82	12
Black, Melanoma	-0.02±0.01	203	64	12
Obese, Heart disease	-0.06±0.01	852	89	19
Obese, Fibromyalgia	0.01±0.01	251	53	18
Obese, COPD	-0.08±0.01	1470	75	17
Asian, Heart disease	0.09±0.02	890	86	17
Asian, Depression	0.32±0.05	958	54	14
Asian, Back pain	0.09±0.02	944	78	14
Asian, Rh. Arthritis	0.0±0.0	130	48	8
Asian, Fibromyalgia	0.04±0.01	294	44	9
Asian, Chronic kidney disease	0.01±0.02	4	33	9
Asian, Endometriosis	0.02±0.01	108	67	21

Table F6: All 34 valid UK Biobank implicit bias effect estimates with the size of the minimum highly-influential set such that its removal leads to an estimate $\theta_{\text{inf}} = 0$. We also show the number of Z (out of 112) and X (out of 65) proxy features selected by the proxy removal algorithm in producing a valid effect estimate. COPD = Chronic obstructive pulmonary disease. IBD = Inflammatory bowel disease. p.s. = post-secondary. insr. = insurance.

Fig. F6: Heatmap of the Szymkiewicz-Simpson coefficient across all (D, Y) pairs' high-influence sets for which $|\theta| > 0.01$. The set size of the high-influence set is reported in parentheses. COPD = Chronic obstructive pulmonary disease. IBD = Inflammatory bowel disease. p.s. = post-secondary. insr. = insurance.
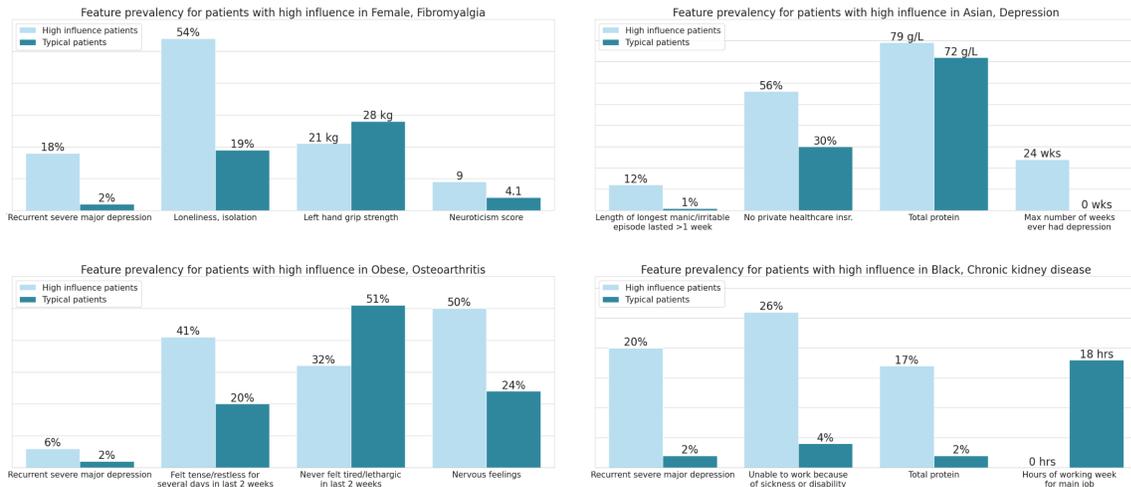
Fig. F7: Feature prevalency between high-influence patients and the rest of the data for four (D, Y) pairs. Four features were selected randomly for visualization. Categorical or binary features are aggregated by mean, continuous features are by median.

(D, Y)	$\theta \pm 95\%$ **CI** $\mathbb{E}[\cdot \mid W] = $ **Lasso**	$\theta \pm 95\%$ **CI** $\mathbb{E}[\cdot \mid W] = $ **XGBoost**
Low income, Depression	0.03±0.02	0.03±0.01
Disability insurance, Rh. Arthritis	0.06±0.00	0.04±0.00
Female, Heart disease	-0.19±0.06	-0.17±0.04
Black, Chronic kidney disease	0.14±0.03	0.10±0.03
Obese, Osteoarthritis	0.09±0.02	0.08±0.02
Asian, Osteoarthritis	-0.06±0.03	-0.06±0.03

Table F7: Comparing Lasso versus XGBoost regression for estimating the residuals of confounders W for six UK Biobank (D, Y) pairs. We use the same X, Z proxy selection sets.
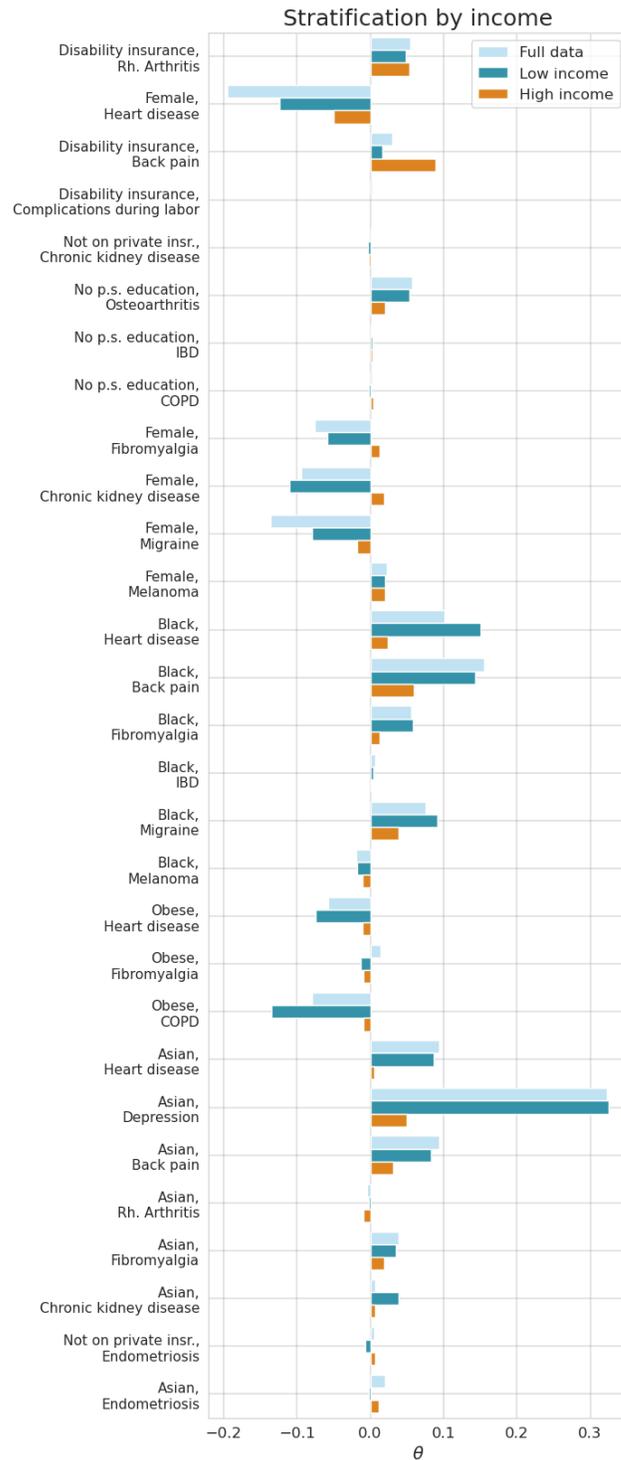
Fig. F8: Point re-estimates from stage 2 after stratification by income. COPD = Chronic obstructive pulmonary disease. IBD = Inflammatory bowel disease. p.s. = post-secondary. insr. = insurance.